

IEEE

control systems

APRIL 2026 VOLUME 46 NUMBER 2



Special Issue on Extremum Seeking

IEEE
Control Systems
Society™

IEEE




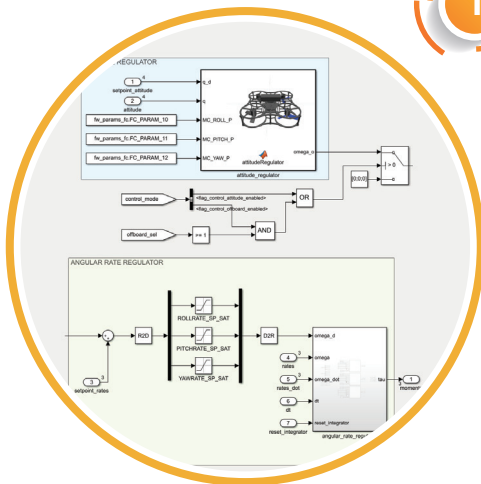
Drone Lab

The ready-to-use solution for drone simulation and flight testing



DEVELOP

Design custom controllers using MATLAB and Simulink. Automatically generate controller code and integrate in PX4. No C++ knowledge required!



SIMULATE

Validate and test your controller using PX4, Gazebo and ROS in a simulation environment that mirrors real-world conditions.



FLY!

Easily deploy your custom controller to the drone and conduct flight tests in a safe environment, in a repeatable way.



Open source



Powered by PX4[®] and ROS[™]



PX4 is a registered trademark of Lorenz Meier. The ROS logo is a trademark of Open Source Robotics Foundation.

Reach out to us to learn more on our turnkey laboratory solutions for research and education on UAVs.


www.antx.it | info@antx.it

IEEE control systems

April 2026 Volume 46 Number 2

WWW.IEEECSS.ORG/PUBLICATIONS/CSM

» FEATURES



COVER ARTWORK WAS GENERATED USING CHATGPT FROM A PROMPT AUTHORED BY TIAGO ROUX OLIVEIRA.

36 Into the Second Century of Extremum Seeking Control

An Introduction to the Special Issue

TIAGO ROUX OLIVEIRA, JORGE I. POVEDA, and MARTIN GUAY

39 Extremum and Nash Equilibrium Seeking With Delays and PDEs: Designs & Applications

Into the Second Century of Extremum-Seeking Control

TIAGO ROUX OLIVEIRA, MIROSLAV KRSTIĆ, and TAMER BAŞAR

88 Hybrid Set-Seeking Systems

Model-Free Feedback Optimization via Hybrid Inclusions

JORGE I. POVEDA and ANDREW R. TEEL

130 Examples of Extremum Seeking Control in Industry

Autotuning Applications From Electromagnetic Brakes to Fluid Dynamics

MOUHACINE BENOSMAN

150 Adaptive Control Meets Deep Learning

Extremum Seeking in Latent Space With Particle Accelerator Applications

ALEXANDER SCHEINKER

» DEPARTMENTS



3 FROM THE EDITOR

Thinking, Fast and Slow, and Eureka!

6 ABOUT THIS ISSUE

Special Issue on Extremum Seeking

8 PRESIDENT'S MESSAGE

Peer-to-Peer: Reviewing and Publishing in the IEEE Control Systems Society

10 25 YEARS AGO

Putting Energy Back in Control

14 THE LIGHTER SIDE OF CONTROL SYSTEMS

15 MEMBER ACTIVITIES

CSS Participates in IEEE Tech Summit 2025

18 TECHNICAL ACTIVITIES

Technical Committee on Healthcare and Medical Systems

Some Activities of the Technical Committee on Stochastic Systems and Control in 2025

- 27 Diversity, Outreach, and Development Activities**
Expanding Horizons in Systems and Control for Indonesia's Young Scientists
- 31 PEOPLE IN CONTROL**
Mouhacine Benosman
Martin Guay
- 177 HISTORIES OF CONTROL**
Extremum Seeking: The Simple Feedback That Will Outlive Us All
- 180 BOOK REVIEW**
Data-Driven Methods for Dynamic Systems
- 182 CONFERENCE REPORTS**
Riotous Assembly in San Diego: Ninth IEEE Conference on Control Technology and Applications
- 186 CSS AWARDS**
2025 IEEE CSS Awards
- 191 CONFERENCE REPORTS**
European Control Conference 2025
- 196 OBITUARY**
Dr. Neng (Eva) Wu.
- 198 CSS BUSINESS**
IEEE CSS Board of Governors Meeting December 9, 2025
Minutes of the Control Systems Society Board of Governors Meeting December 9, 2025
- 216 CONFERENCE CALENDAR**

IEEE PUBLISHING OPERATIONS
445 Hoes Lane, Piscataway, NJ 08854 USA

IEEE OFFICERS

Mary Ellen Randall, *IEEE President and CEO*
 Jill I. Gostin, *IEEE President-Elect*
 Kathleen A. Kramer, *IEEE Past President*
 Jenifer P. Castillo Rodriguez, *Director & Secretary*
 Gerardo Barbosa, *Director & Treasurer*
 Timothy P. Kurzweg, *Director & Vice President, Educational Activities*
 Andrew D. Lowery, *Director & Vice President, Member & Geographic Activities*
 W. Clem Karl, *Director & Vice President, Publication Services and Products*
 Gary R. Hoffman, *Director & President, Standards Association*
 F D Tan, *Director & Vice President, Technical Activities*
 Barry C. Tilton, *Director & President IEEE-USA*

IEEE EXECUTIVE STAFF

Sophia Muirhead, *Executive Director and COO*
 Kelly Armstrong, *Chief Financial Officer*
 Mojdeh Bahar, *Managing Director, Technical Activities*
 Ahsaki E. Benion, *General Counsel and Chief Compliance Officer*
 Russell Harrison, *Managing Director, IEEE-USA*
 Steven Heffner, *Chief Publication Officer and Managing Director*
 Donna Hourican, *Staff Executive, Corporate Activities*
 Marie Hunter, *Managing Director, Conferences, Events and Experiences*
 Cecelia Jankowski, *Managing Director, Member and Geographic Activities*
 Jamie Moesch, *Managing Director, Educational Activities*
 Jayne O'Brien, *IEEE Chief Marketing Officer (CMO)*
 Alpesh Shah, *IEEE Standards Association Managing Director*
 Allison Taylor, *IEEE Chief of Staff*
 Jeff Strohschein, *Chief Information Digital Officer*
 Cheri N. Wideman, *Chief Human Resources Officer*

IEEE PUBLISHING OPERATIONS

Dawn Melley, *Senior Director, Publishing Operations*
 Kevin Lisankie, *Director, Editorial Services*
 Peter M. Tuohy, *Director, Production Services*
 Neelam Khinvasara, *Associate Director, Digital Assets & Editorial Support*
 Felicia Spagnoli, *Advertising Production Manager*
 Shannon Campos, *Manager, Periodicals Production*

ADVERTISING SALES

Timothy Warder
Director of New Product and Audience Development

IEEE prohibits discrimination, harassment, and bullying. For more information, visit <https://www.ieee.org/nondiscrimination>.

MISSION STATEMENT AND SCOPE: As the official means of communication for the IEEE Control Systems Society, *IEEE Control Systems* publishes interesting, useful, and informative material on all aspects of control system technology for the benefit of control educators, practitioners, and researchers. With this mission statement in mind, *IEEE Control Systems* encourages submissions, both feature articles and columns, on all aspects of control system technology.

SUBMISSION OF MANUSCRIPTS: A feature article typically provides an in-depth treatment of either an application of control technology, a tutorial on some area of control theory, or an innovation in control education.

IEEE Control Systems publishes a variety of columns. "Applications of Control" columns are industrially oriented summaries of innovations in control technology. "Focus on Education" typically describes some aspect of education, such as novel control experiments. "Lecture Notes" can be theoretical in nature as long as they have clear tutorial value and intent. See recent issues for examples of these and other types of columns. Authors are encouraged to contact the editor-in-chief about the suitability of potential columns.

A detailed Author's Guide, a sample formatted manuscript, and LATEX template can be found at <http://ieeecss.org/publication/ieee-control-systems-magazine>. The specifications in this guide should be followed by all submissions.

All manuscripts should be submitted electronically to the *IEEE Control Systems* website, <https://css.paperplaza.net/conferences/scripts/start.pl>, with inquiries on appropriateness of content e-mailed to aanna@mit.edu.

SPECIAL ISSUES: *IEEE Control Systems* encourages proposals for special issues. Proposers are encouraged to contact the editor-in-chief to discuss potential topics.

BOOKS AND CONFERENCES: Submit information about recently published books to the associate editor for book reviews. Submit information about past and future conferences to the corresponding editor for conferences.

ADVERTISING: *IEEE Control Systems* accepts advertising for educational products, books, software, conferences, employment, and control-related technology. For information about advertising, contact Timothy Warder, t.warder@ieee.org, +1 732-562-6596.

IEEE CONTROL SYSTEMS—(ISSN 1066-033X) (ISMAD7) is published bi-monthly by The Institute of Electrical and Electronics Engineers, Inc. Headquarters: 3 Park Avenue, 17th Floor, New York, NY 10016, U.S.A. +1 212 419 7900. Responsibility for the contents rests upon the authors and not upon the IEEE, the Society, or its members. To order individual copies for members and nonmembers, please e-mail the IEEE Contact Center at contactcenter@ieee.org. Member and nonmember subscription prices available on request. Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of the U.S. Copyright law for private use of patrons: 1) those post-1977 articles that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01970, U.S.A.; and 2) pre-1978 articles without fee. For other copying, reprint, or republication permission, write to: Copyrights and Permissions Department, IEEE Service Center, 445 Hoes Lane, Piscataway NJ 08854, U.S.A. Copyright © 2026 by The Institute of Electrical and Electronics Engineers, Inc. All rights reserved. Periodicals postage paid at New York, NY, and additional mailing offices. Postmaster: Send address changes to *IEEE Control Systems*, IEEE, 445 Hoes Lane, Piscataway, NJ 08854 U.S.A. Canadian GST #125634188 Printed in U.S.A.

Thinking, Fast and Slow, and Eureka!

This column was inspired by the awards season, a time when many organizations recognize their top performers through peer-nominated awards. Elsewhere in this issue, an article from our Awards Chair, Thomas Parisini, provides a comprehensive list of these awards and their recipients. As I sat through the awards ceremony, I began to wonder whether there is a shared experience or a glide path, a common trajectory in the thinking processes of the awardees: a “Eureka” moment underlying their contributions. To explore this, I reached out to several awardees and other luminaries in the field, posing the following questions: What were the key steps in your thought process that led to this moment of breakthrough? Could you walk us through this process, ideally in a nontechnical, personal reflection, though technical details are welcome if necessary? Nearly all responded enthusiastically, and I present their exact responses in what follows.

TAMER BAŞAR

I have several “Eureka moments,” but the earliest one goes way back to the days (or nights) when I was working on my doctoral thesis at Yale, which was on minimax estimation for stochastic dynamical systems, where the noise distribution was unknown with some soft power constraints and the goal was to determine the “best” estimate of the state (the “min” part) under worst-case (so-called “least favorable”) noise distributions (the “max” part). I had solved it for the scalar case, but extending it to multidimensional systems presented a real challenge.

Thinking about this kept me (actually my mind) busy for several days (and nights), and finally one night during half-sleep it dawned on me that a subspace approach would be the right way to go; I immediately got up, found a piece of paper and a pencil (no laptops at the time) and jotted down an outline of the proof, and then went back to sleep with a piece of mind. The next morning, I went back to the notes I had taken in the middle of the night. The outline seemed correct, and I finalized the proof. This single “Eureka moment” kicked off the writing and publication of several journal papers in Annals of Statistics, International J. Control, and Stochastics, and a number of conference papers, including CDC. The structural result of this “Eureka moment” motivated, about 10 years later, an information-theoretic solution to encoder–decoder design under worst-case channel conditions, which appeared in IEEE Transactions on Information Theory.

My second Eureka moment, in brief, is the following: This arose in the study of uniqueness or otherwise of Nash equilibrium solutions in dynamic games, an issue of pertinent importance in this area, as acknowledged by many in multiple fields. Thinking about this kept me busy for several weeks and months, and one afternoon in the mid-1970s, it dawned on me that use of memory strategies would be the answer to nonuniqueness, which I called informational nonuniqueness, and it would also open the path toward establishing uniqueness by introducing (and requiring insensitivity to) modeling uncertainty. This led to a paper that appeared in International Journal of Game Theory in 1976 (and another one for the continuous-time version in SIAM Journal of Control in 1977), around the time Selten (1994 Nobel Laureate along with Nash and Harsanyi) published in the same journal his “trembling hand”



The awards ceremony at the CDC2025.

Although frustrating, repeatedly going down blind alleys unconsciously built the familiarity needed to truly understand the problem and eventually solve it.

approach for finite games, serving a similar purpose.

SIMONE GARRATI

One of my most meaningful Eureka moments came a few years ago while working with Marco Campi on generalizing the results of the scenario approach. After many blind alleys and failed attempts, we found ourselves facing infinite-dimensional generalized moment problems whose solution persistently eluded us. I remember living with the problem and developing an almost visual intuition, to the point where I could “see” probabilistic mass flowing between constraints. Then, quite suddenly, a small but crucial shift appeared: The introduction of a refined dominance condition immediately revealed to me a dual formulation equivalent to one I knew would lead to the solution. This moment was deeply exciting and marked the beginning of a theory for risk certification that applies to very general decision schemes under minimal assumptions. I especially value this moment because it was a truly collective Eureka moment: Marco and I were working face to face, with ideas bouncing back and forth in a sustained exchange that ultimately carried us to the final result.

Another remarkable Eureka moment came much earlier in my career, just a couple of years after completing my Ph.D. I was trying to characterize the distribution of risk in robust scenario optimization problems, guided at the time only by a conjecture and a loose intuition based on counterfactual calculations in specific cases. As often happens, progress required living with the problem for a long time: Although frustrating, repeatedly going down blind alleys unconsciously built the familiarity needed to truly understand the problem and eventually solve it. That is

exactly what happened. While reading a seemingly unrelated piece in combinatorial analysis, a sudden intuition emerged. Although the conceptual step was substantial, I could immediately see the entire path forward, and a beautiful moment problem crystallized, yielding the beta distribution as the solution. That insight marked the beginning of moment-problem methods in scenario theory and opened a line of research that has gone far, thanks also to a long-standing collaboration with exceptional coauthors.

JULIAN BERBERICH

I had been searching for the right angle to analyze robustness properties of quantum algorithms for months. One day, almost by accident, I was attending a talk from the quantum community that showed a specific block diagram. It suddenly became obvious what I needed to do: define the correct input–output map and analyze its gain. With the right viewpoint, the mathematical analysis fell into place within hours. This showed me, once again, the power of thinking in systems, which I see as one of the (if not THE) most valuable tools of our community.

SERGIO GRAMMATICO

The Eureka moment for our Transactions on Control of Network Systems Best Paper Award arrived when we stopped doing what everyone (including us) had been doing: endlessly refining distributed optimization algorithms to cope with whatever information network was given. It felt like we were driving cars collaboratively, almost perfectly ignoring that the road network might be redesigned! That realization flipped the problem on its head. Instead of asking “How can algorithms survive limited information?” we asked, “What if we design the information network so the algo-

rithms can thrive?” Once that inversion clicked, the entire framework fell into place with surprising simplicity; it was the kind of moment where the conversations gained scientific momentum, and the whiteboard suddenly became highly populated with formulas. The key ingredients were curiosity in questioning the usual background assumptions and a novel integration of operator theory and consensus theory applied to computational game theory and optimization

REZA MOHEIMANI

For months, I was trying to understand why scanning tunneling microscope (STM) tips frequently crashed during imaging and lithography. The Eureka moment occurred while examining a set of Bode plots obtained by my graduate student through closed-loop system identification of an STM. I noticed that the plots exhibited markedly different low-frequency gains, despite the measurements being taken with the same tip and sample. I immediately recognized that the underlying cause was the local barrier height—a quantum mechanical property of the tip–surface interaction—that acts as a variable component in the gain of the STM transfer function. This insight enabled us to formulate a control strategy to protect the tip, sparked a cascade of subsequent innovations in scanning tunneling microscopy, and, *inter alia*, resulted in the Transition to Practice Award.

ROBERT GREGG

After completing my Ph.D. on the control of dynamic walking robots, I pivoted into a postdoc in rehabilitation engineering at a time when there was growing skepticism for robot-assisted rehabilitation, especially from clinicians. The first clinical trials on robot-assisted stroke gait training were disappointing; patients could literally fall asleep as the robot walked them. While studying the results, I thought to myself that the problem may not be fundamental to the robot, but rather how it is controlled. Initially, I tried thinking of ways to adapt the reference trajectories, such as making them time-invariant, but the user would still be tied to a predefined pattern that may not help them recover motor

function. My Eureka moment was the day it occurred to me that the person's voluntary walking gait is analogous to a passive dynamic walking gait, which can be shaped by external actuators to exhibit different dynamics. Instead of assigning joint trajectories for the person, we could augment their perceived dynamics to make it easier for them to practice their walking gait patterns. This inspired me to translate my Ph.D. work on energy shaping control into trajectory-free controllers for exoskeletons that assist the user's voluntary gait patterns. I remember the first time that I experienced modified gravitational dynamics through an exoskeleton while remaining in full control of my limb movements, which convinced me this was the right path forward for the field. This ultimately became the topic of my National Science Foundation CAREER award, led to subsequent National Institutes of Health-funded clinical trials, and earned me the Transactions on Control Systems Technology Outstanding Paper Award.

ANASTASIA BIZYAEVA

I find that my research is frequently pushed forward by small Eureka moments, usually in the form of making unexpected mental connections. Recently, I attended a seminar talk on fault detection methods in industrial factory automation. It was an interesting topic that I had no familiarity with a priori. In the middle of the talk, the speaker mentioned an idea that reminded me a lot of a machine learning problem I have been stuck on for months. Following that thread, I discovered that the exact mathematical problem I have been struggling to understand has been extensively studied in the fault detection literature, under a set of keywords I never would have known to look up had it not been for this chance attendance of a seemingly unrelated seminar.

MASIH HASELI

In 2019, we showed that a necessary and sufficient condition for finding Koopman eigenfunctions in a finite-dimensional space is to apply extended dynamic mode decomposition forward and backward in time and compare the eigen decompositions

I find that my research is frequently pushed forward by small Eureka moments, usually in the form of making unexpected mental connections.

of the resulting matrices. Since then, I have been interested in extending the result to approximations. Finally, in 2022, I found some time to work on the problem, which led to a tight error bound on all function predictions (not just the approximated eigenfunctions) in the subspace. There was no Eureka moment; it was more of a steady progression.

ANU ANNASWAMY

I was chasing the trajectories of a nonlinear dynamic system that I was sure was a global property, not a local one. I was totally living on the manifold, trying to impose all of the properties that I knew to be true. I obsessed over this for months. And one day, after a nap, I knew exactly what was happening with these trajectories. That led to a proof (one of my best) and the George Axelby Award.

JOHN DOYLE

These are interesting stories but completely different from my comparable recollections. Maybe it is just that I have a poor memory.

Is there anything to be gleaned from these (definitely not statistical) sample responses? An $n = 10$, and that too a very select 10 set of responses, is by no means an encapsulation of the process of creativity. With this important caveat, I attempt to relate our thought processes to the framework described in Daniel Kahneman's *Thinking, Fast and Slow*, particularly the concepts of "System 1" and "System 2" modes of thinking.

Kahneman posits in his book [1] that our cognitive processes can be categorized into a System 1 and a System 2. System 1 corresponds to *fast thinking*, replete with characteristics that are quick, effortless, unconscious, intuitive, and emotional responses that rely on patterns and experience.

In contrast, System 2 corresponds to *slow thinking*, which is deliberate, slow, effortful, conscious, logical, analytical, and requires focus and attention. My hypothesis, certainly not the first, as noted in [2] and likely elsewhere, is that scientific creativity emerges from the interplay between System 1 and System 2. System 1 enables rapid connections between disparate concepts and the identification of hidden patterns, sparking moments of insight. System 2 then provides the structure and rigor necessary to develop these insights, for the spark to ignite and forge an edifice through methodical analysis, technical excavations and explorations, and judicious integration of nuanced mathematical tools. I suggest that the responses above provide evidence of this combined application resulting in their impactful contributions.

Whether or not you agree with my hypothesis, I invite you to analyze your own thought processes that you have engaged in in your research and see where you are with respect to such thinking, or whether you totally break the mold, and have your own unique glide path that simply prevents you from being pinned in any such box. It would be doubly wonderful if you are able to share your thoughts. Who knows? Your responses themselves could form the subject matter for a future column!

—Anuradha Annaswamy^{ID}

REFERENCES

- [1] D. Kahneman, *Thinking, Fast and Slow*. Macmillan, 2011.
- [2] R. Khalil and M. Brüne, "Adaptive decision-making 'fast' and 'slow': A model of creative thinking," *Eur. J. Neurosci.*, vol. 61, no. 5, 2025, Art. no. e70024, doi: 10.1111/ejn.70024.



Special Issue on Extremum Seeking

Volume 46, issue 2, corresponds to a Special Issue on Extremum Seeking. The guest editors are Tiago Roux Oliveira, Jorge Ivan Poveda, and Martin Guay, assisted by the magazine’s Associate Editor Alexander Bazanella. The issue presents the topic through four articles, all of which help provide a snapshot of this energetic topic. The “From the Editor” column [A1], inspired by the awards season that typically occurs at the end of each year in various spheres, talks about our collective Aha moments in our research. Finally, the President’s column explores our extensive peer-review enterprise [A2].


The “Member Activities” column celebrates the role of local chapters

in the IEEE Control Systems Society (CSS), with a focus on the IEEE Italy Chapter and the IEEE Nanjing Chapter [A3]. The two technical committees showcased in the “Technical Activities” section of this issue are Healthcare and Medical Systems and Stochastic Systems and Control [A4]. The recently introduced Diversity, Outreach, and Development Activities (DODA) spotlights another success story stemming from a DODA-backed initiative [A5].

The “People in Control” segment showcases Dr. Mouhacine Benosman, principal research scientist at Amazon Robotics, and Dr. Martin Guay, professor at Queen’s University, Can-

ada [A6]. In this issue, we also return to a second edition of the “Histories in Control” [A7] segment that we introduced in volume 45, issue 3. This segment is where we talk to a leading researcher on a specific topic in control systems, who provides their perspectives and point of view of how the field evolved, various milestones, bottlenecks, challenges, surprises, and overall

impact. This is meant to be an interview with their musings as the focus rather than an objective exposition of the field. In this issue, we speak with Prof. Miroslav Krstic, who talks about the topic of extremum seeking.



Digital Object Identifier 10.1109/MCS.2026.3659228

Date of current version: 26 March 2026

Contributors



Tamer Başar



Mouhacine Benosman



Martin Guay



Miroslav Krstić



Tiago Roux Oliveira



Jorge I. Poveda



Alexander Scheinker




Andrew R. Teel

This issue also covers several awards conferred annually to individuals and teams for technical achievements in the various areas of interest of the Society, as well as awards for service to the Society. In addition, we feature an “In Memoriam” tribute to Prof. Eva Wu of Binghamton University, recognizing her contributions and commemorating her passing [A8].

In the “Book Review” segment, Yorie Nakahira presents a book review on *Data-Driven Methods for Dynamic Systems* by Jason Bramburger [A9].

Among the regular columns, “25 Years Ago” revisits the article “Putting energy back in control” authored by Ortega et al. [A10]. The “Conference Reports” section provides reports on two conferences, the Conference on Control Technology and Applications 2025, which is a CSS conference, and European Control Conference 2025, which is technically cosponsored by the CSS [A11]. The “Conference Calendar” [A12] lists upcoming conferences sponsored or

cosponsored by the CSS. We cover CSS business that pertains to the minutes and consent agenda of the Board of Governors meeting that occurred at the 2025 IEEE Conference on Decision and Control in Rio in December. Finally, we have a brand new comic strip from Brian Douglas, which is now a recurring segment, put together exclusively for the readers of the magazine! Enjoy!

—Anuradha Annaswamy 

APPENDIX: RELATED ARTICLES

- [A1] A. Annaswamy “Thinking, fast and slow, and Eureka! [From the Editor],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 3–5, Apr. 2026, doi: [10.1109/MCS.2026.3659227](https://doi.org/10.1109/MCS.2026.3659227).
- [A2] C. Beck and A. Serrani, “Peer-to-peer: Reviewing and publishing in the IEEE Control Systems Society [President’s Column],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 8–9, Apr. 2026, doi: [10.1109/MCS.2026.3659160](https://doi.org/10.1109/MCS.2026.3659160).
- [A3] N. Ozay, “CSS participates in IEEE Tech Summit 2025 [Member Activities],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 15–17, Apr. 2026, doi: [10.1109/MCS.2026.3659287](https://doi.org/10.1109/MCS.2026.3659287).
- [A4] S. D. G. M. Pequito, “Technical Committee on Healthcare and Medical Systems [Technical Activities],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 18–24, Apr. 2026, doi: [10.1109/MCS.2026.3659230](https://doi.org/10.1109/MCS.2026.3659230).

[A5] S. Mastellone, “Report on diversity, outreach, and development activities [Diversity, Outreach, and Development Activities],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 27–30, Apr. 2026, doi: [10.1109/MCS.2026.3659232](https://doi.org/10.1109/MCS.2026.3659232).

[A6] A. Annaswamy, “People in Control,” *IEEE Control Syst.*, vol. 46, no. 2, p. 31, Apr. 2026, doi: [10.1109/MCS.2026.3659289](https://doi.org/10.1109/MCS.2026.3659289).

[A7] A. Annaswamy, “Extremum seeking: The simple feedback that will outlive Us all [Histories of Control],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 177–179, Apr. 2026, doi: [10.1109/MCS.2026.3659233](https://doi.org/10.1109/MCS.2026.3659233).

[A8] J. Sun, “In Memoriam: Neng (Eva) Wu, 1956–2025 [Obituary],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 196–197, Apr. 2026, doi: [10.1109/MCS.2026.3659282](https://doi.org/10.1109/MCS.2026.3659282).

[A9] Y. Nakahira, “Data-driven methods for dynamic systems [Book Review],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 180–181, Apr. 2026, doi: [10.1109/MCS.2026.3659234](https://doi.org/10.1109/MCS.2026.3659234).

[A10] R. Ortega, A. J. van der Schaft, I. Mareels, and B. Maschke, “Putting energy back in control [25 Years Ago],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 10–13, Apr. 2026, doi: [10.1109/MCS.2026.3659161](https://doi.org/10.1109/MCS.2026.3659161).

[A11] R. Bitmead, “Riotous assembly in San Diego: Ninth IEEE conference on control technology and applications [Conference Reports],” *IEEE Control Syst.*, vol. 46, no. 2, pp. 182–185, Apr. 2026, doi: [10.1109/MCS.2026.3659280](https://doi.org/10.1109/MCS.2026.3659280).

[A12] “[Conference Calendar]” *IEEE Control Syst.*, p. 216, Apr. 2026, doi: [10.1109/MCS.2026.3659163](https://doi.org/10.1109/MCS.2026.3659163).



Peer-to-Peer: Reviewing and Publishing in the IEEE Control Systems Society

The scale of our peer-review enterprise is both impressive and sobering. As of December 2025, approximately 7,300 papers were submitted to IEEE Control Systems Society (CSS) journals over the preceding year. These submissions were handled by roughly 60 senior editors (SEs) and 300 associate editors (AEs) (see Table 1), numbers that will need to grow substantially in 2026. Given IEEE guidelines requiring a minimum of three reviews per paper, this translates to more than 22,000 journal reviews completed by our community in a single year, just for IEEE publications.

One way to interpret this number is through the idea of being *review-neutral*: for every paper you submit, you and your coauthors collectively contribute three reviews in return. In principle, this represents a fair and sustainable system: each author gives back what they expect to receive. In practice, however, anyone who has served as an AE knows that securing timely, high-quality reviews (or sometimes, any reviews at all) is increasingly difficult. We are all busy and seemingly getting busier. Still, it is worth emphasizing that submitting a paper implicitly carries a professional obligation to review the work of others. While CSS does not currently enforce this as a norm, several other research communities have begun moving in that direction. For example, major machine learning conferences now require author participation in the review process. At the 2025 International Conference on Machine Learning, at least one “qualified” author per submission was required to

register as a reviewer, with qualifications defined by having prior publications in recognized venues; exemptions were possible but required justification. Similar policies are now in place for the Annual Conference on Neural Information Processing Systems.

Reviewing demands comparable to those of CSS journals exist for CSS conferences as well. For the Conference on Decision and Control (CDC), just over 250 AEs serving on the Conference Editorial Board (CEB) oversaw approximately 8,000 reviews for around 2,200 papers (with some papers handled through *IEEE Control Systems Letters* and tutorial papers managed separately). For the IEEE Conference on Control Technology and Applications, about 110 AEs serving on the Technical Conference Editorial Board (TCEB) handled roughly 250 papers in 2025.

Given these demands, a common question we receive is how these AEs and SEs are selected and whether there is a systematic process for joining the journal and conference editorial review boards. In brief, the selection process

is as follows, and a review of every nomination for SE and AE is conducted at the Spring Executive Committee meeting each year (see Figures 1 and 2) according to these guidelines:

- 1) SEs are selected from experienced AEs (typically those with multiple years of service) who have demonstrated expertise in areas of need for the journal.
- 2) Journal AEs (for *IEEE Transactions on Automatic Control*, *IEEE Transactions on Control Systems Technology*, *IEEE Transactions on Control of Network Systems*, *IEEE Control Systems Letters*, or *IEEE Open Journal of Control Systems*) are generally required to have at least three years of service on the CEB or TCEB, with strong records of responsiveness and attention to detail, or comparable service elsewhere. AEs are typically assumed to be at the associate professor level, to have roughly five years of industry experience, or to be at a similar career stage within other systems.
- 3) CEB and TCEB members are individuals who typically are at least

TABLE 1 Paper submissions and editor numbers for CSS journals.

Journal	SEs	AEs	2024 Submissions	2025 Submissions
<i>IEEE Control Systems</i>	0	18	52	39
<i>IEEE Control Systems Letters</i>	13	80	1,463	1,474
<i>IEEE Open Journal of Control Systems</i>	2	25	54	107
<i>IEEE Transactions on Automatic Control</i>	7	60	2,982	3,384
<i>IEEE Transactions on Control of Network Systems</i>	9	40	915	958
<i>IEEE Transactions on Control Systems Technology</i>	10	50	1,207	1,486

Digital Object Identifier 10.1109/MCS.2026.3659160
Date of current version: 26 March 2026



FIGURE 1 CSS editors-in-chief, Conference Editorial Board chairs, and Executive Committee members gather for the 2025 spring meeting, where SE and AE nominations are reviewed.




FIGURE 2 CSS editors-in-chief, CEB chair, and TCEB chair (Miroslav Krstic, Anu Annaswamy, Amir Aghdam, Alessandro Astolfi, Ilya Kolmanovsky, Laura Menini, Lacre Pavel, and Stefano di Cairano) at the spring meeting.



Andrea Serrani, coauthor and CSS vice-president for publication activities, at the CDC 2025 Awards Ceremony.

three years past their Ph.D. and are nominated as in guideline 4.

- 4) All AE appointments require a referral or nomination from someone at a similar or higher level in the review system, and applications should be submitted through PaperPlaza.

A crucial point bears repeating: all SEs, AEs, and editors-in-chief serve entirely as volunteers, donating substantial time each week to support the CSS publication and conference ecosystem. The success of this system depends fundamentally on their professionalism, dedication, and generosity, for which the community should be genuinely grateful.

As a final topic, we would like to invite discussion on the structure of the review process itself: specifically, should CSS consider introducing some level of *blindness* in the review process? Currently, both AEs and reviewers know the authors' identities, and the authors know their assigned AEs but, of course, not their reviewers. Possible alternatives include delaying disclosure of AE identities until the time of publication or adopting a form of double-blind review in which reviewers do not know the authors' identities until a final decision is reached. Other research communities have implemented such approaches with some degree of success, although we recognize that maintaining anonymity is increasingly challenging in the era of arXiv and widespread online preprints.

We welcome your thoughts on whether CSS should explore partial or full blind review mechanisms, as well as any questions or ideas you may have regarding review processes, reviewer expectations, and the required editorial service. Sustaining a fair, efficient, and high-quality review system ultimately depends on shared responsibility across our entire community.

Carolyn Beck 

CSS President

Andrea Serrani 

CSS VP Publication Activities



Putting Energy Back in Control

ROMEO ORTEGA, ARJAN J. VAN DER SCHAFT, IVEN MAREELS, AND BERNHARD MASCHKE

In this issue, “25 Years Ago” revisits the article “[Putting Energy Back in Control](#)” by Romeo Ortega, Arjan J. van der Schaft, Iven Mareels, and Bernhard Maschke, in *IEEE Control Systems Magazine*, vol. 21, no. 2, pp. 18–33. Below is an excerpt from the article.

Energy is one of the fundamental concepts in science and engineering practice, where it is common to view dynamical systems as energy-transformation devices. This perspective is particularly useful in studying complex nonlinear systems by decomposing them into simpler subsystems that,

upon interconnection, add up their energies to determine the full system’s behavior. The action of a controller may also be understood in energy terms as another dynamical system—typically implemented in a computer—interconnected with the process to modify its behavior. The control problem can then be recast as finding a dynamical system and an interconnection pattern such that the overall energy function takes the desired form. This “energy-shaping” approach is the essence of passivity-based control (PBC), a controller design technique that is very well known in mechanical systems.

Our objectives in this article are threefold:

- » First, to call attention to the fact that PBC does not rely on some

particular structural properties of mechanical systems, but hinges on the more fundamental (and universal) property of energy balancing.

- » Second, to identify the physical obstacles that hamper the use of “standard” PBC in applications other than mechanical systems. In particular, we will show that “standard” PBC is stymied by the presence of unbounded energy dissipation, hence it is applicable only to systems that are stabilizable with passive controllers.
- » Third, to revisit a PBC theory that has been recently developed to overcome the dissipation obstacle as well as to make the incorporation of process

Digital Object Identifier 10.1109/MCS.2026.3659161

Date of current version: 26 March 2026

» IEEE CONTROL SYSTEMS EDITORIAL BOARD

EDITOR-IN-CHIEF

Anuradha Annaswamy
Massachusetts Institute of Technology
Department of Mechanical Engineering
Cambridge, MA 02139
USA
aanna@mit.edu

PAST EDITORS-IN-CHIEF

Mo Jamshidi (1981–1984)
Herbert E. Rauch (1985–1992)
Stephen Yurkovich (1993–1998)
Tariq Samad (1999–2003)
Dennis S. Bernstein (2003–2011)
Richard D. Braatz (2012–2014)
Jonathan P. How (2015–2019)
Rodolphe Sepulchre (2020–2024)

ASSOCIATE EDITORS, BOOK REVIEWS

Tryphon Georgiou
University of California at Irvine, USA

Scott R. Ploen
Jet Propulsion Laboratory, USA

Somayeh Soujoudi
University of California at Berkeley, USA

CSN ASSOCIATE EDITOR, EDUCATION

Richard Pates
Lund University, Sweden

CORRESPONDING EDITOR, CONFERENCE ACTIVITIES

S. Zeeshan Rizvi
Coming, Inc., USA

TECHNICAL ASSOCIATE EDITORS

Daniel Abramovitch
Agilent Technologies, USA

Mahnoosh Alizadeh
University of California, Santa Barbara, USA

Ravi Banavar
IIT Bombay, India

Alexandre Bazanella
Universidade Federal do Rio Grande do Sul, Brazil

Silvère Bonnabel
University of New Caledonia, France

Jie Chen
Beijing Institute of Technology, China

Shlomo Engelberg
Jerusalem College of Technology, Israel

Kingsley Fregene
Lockheed Martin, USA

Mihailo Jovanovic
University of Southern California, USA

Cameron Nowzari
George Mason University, USA

Rodolphe Sepulchre
University of Cambridge, UK

Gang Wang
Beijing Institute of Technology, China

Lihua Xie
Nanyang Technological University, Singapore

Melanie Zeilinger
ETH, Switzerland

Digital Object Identifier 10.1109/MCS.2026.3659379

Date of current version: 26 March 2026

prior knowledge more systematic. These two important features allow us to design energy-based controllers for a wide range of physical systems.

INTELLIGENT CONTROL PARADIGM

Control design problems have traditionally been approached from a signal-processing viewpoint; that is, the plant to be controlled and the controller are viewed as signal-processing devices that transform certain input signals into outputs. The control objectives are expressed in terms of keeping some error signals small and reducing the effect of certain disturbance inputs on the given regulated outputs, despite the presence of some unmodeled dynamics. To make the problem mathematically tractable, the admissible disturbances and unmodeled dynamics are assumed to be norm bounded, and consequently, the indicators of performance are the size of the gains of the operators that map these various signals. In the case of linear time-invariant systems, this “intelligent control paradigm” (paraphrasing Willems [1]) has been very successful, essentially because disturbances and unmodeled dynamics can be discriminated, via filtering, using frequency-domain considerations. The problem of reducing the gains of nonlinear operators can also be expressed in a clear, analytical way [2]. There are two fundamental differences, however, with respect to the linear time-invariant case: first, the solution involves some far from obvious computations. Second, and perhaps more important, since nonlinear systems “mix” the frequencies, it is not clear how to select the most likely disturbances, and we have to “crank up” the gain to quench the (large set of) undesirable signals and meet the specifications. Injecting high gains in the loop, besides being intrinsically conservative—hence yielding below-par performance—brings along a series of well-known undesirable features (e.g., noise amplification, actuator wear, and high energy consumption).

Our first control objective is to regulate the static behavior (i.e., the equilibria), which is determined by the shape of the energy function.

There are many practical control problems where we have available structural information about the plant. In these cases, it is reasonable to expect that the conservatism mentioned above could be reduced if we could incorporate this prior information in the controller design. Unfortunately, a procedure to systematically carry out this objective does not seem to be available. (The typical approach is to classify the nonlinearities according to the role they play in the derivative of a Lyapunov function candidate. This test has very little to do with the physics of the system. It is obviously tied up with the particular choice of the Lyapunov function, which, stemming from our linear inheritance, is systematically taken to be a quadratic function in the “errors.”) It is our contention that the inability to incorporate prior knowledge is inherent to the signal-processing viewpoint of the intelligent control paradigm and is therefore independent of the particular design technique. In the authors’ opinion, this situation has stymied communication between practitioners and control theorists, seriously jeopardizing the future of modern model-based nonlinear control systems design.

The purpose of this article is to contribute, if modestly, to the reversal of this trend by calling attention to the importance of incorporating energy principles in control. To achieve our objective, we propose to abandon the intelligent control paradigm and instead adopt the behavioral framework proposed by Willems [1]. In Willems’s far-reaching interpretation of control, we start from a mathematical model obtained from first principles, say, a set of higher order differential equations and some algebraic equations.

Among the vector of time trajectories satisfying these equations are components that are available for interconnection. The controller design then reduces to defining an additional set of equations for these interconnection variables to impose a desired behavior on the controlled system. We are interested here in the incorporation into this paradigm of the essential energy component. Therefore, we view dynamical systems (plants and controllers) as energy-transformation devices, which we interconnect (in a power-preserving manner) to achieve the desired behavior. More precisely, we are interested in lumped-parameter systems that satisfy an energy-balancing principle, where the interconnection with the environment is established through power port variables. The power port variables are conjugated, and their product has units of power, for instance, currents and voltages in electrical circuits or forces and velocities in mechanical systems. This is the scenario that arises from any form of physical network modeling.

Our first control objective is to regulate the static behavior (i.e., the equilibria), which is determined by the shape of the energy function. It is therefore natural to recast our control problem in terms of finding a dynamical system and an interconnection pattern such that the overall energy function takes the desired form. There are at least two important advantages of adopting such an “energy-shaping” perspective of control:

- 1) The energy function determines not just the static behavior, but also, via the energy transfer between subsystems, its transient behavior. Focusing our attention on the system energy, we can then aim, not just

at stabilization, but also at performance objectives that can, in principle, be expressed in terms of “optimal” energy transfer. Performance, not stability, is, of course, the main concern in applications.

- 2) Practitioners are familiar with energy concepts, which can serve as a *lingua franca* to facilitate communication with control theorists, incorporating prior knowledge and providing physical interpretations of the control action.

BACKGROUND

The idea of energy shaping has its roots in the work of Takegaki and Arimoto [3] in robot manipulator control, a field in which it is well known and highly successful. Simultaneously and independently of [3], the utilization of these ideas for a large class of Euler-Lagrange systems was suggested in [4]. (See also Slotine’s innovative paper [5] and the related view on the control of physical systems by Hogan [6].) Using the fundamental notion of passivity, the principle was later formalized in [7], where the term PBC was coined to define a controller design methodology whose aim is to render the closed-loop system passive with a given storage function. The importance of linking passivity to energy shaping can hardly be overestimated. On the one hand, viewing the control action in terms of interconnections of passive systems provides an energy-balancing interpretation of the stabilization mechanism. More precisely, we have defined in [8] a class of systems (which includes mechanical systems) such that the application of PBC yields a closed-loop energy that is equal to the difference between the stored and the supplied energies. This special class of PBC energy is called energy-balancing PBC. On the other hand, showing that the approach does not rely on some particular structural properties of mechanical systems, but hinges instead on the more fundamental (and universal) property of passivity, it can be extended to cover a wide range of applications.

In carrying out this extension, two approaches have been pursued:

- » The first approach is similar to classical Lyapunov-based design, where we first select the storage function to be assigned and then design the controller that ensures this objective. Extensive applications of this line of research may be found in [9] (see also [10]-[15]) and are not reviewed here. (It should be noted that in this approach, the desired storage function—typically quadratic in the increments—does not qualify as an energy function in any meaningful physical sense. Actually, it has been shown that the stabilization mechanism is akin to system inversion instead of energy shaping [9], hence a stable invertibility assumption is usually required).
- » The second, newer approach stems from the energy-balancing view of mechanical systems discussed above. The closed-loop storage function—which is now a *bona fide* energy function—is not postulated *a priori*, but is instead obtained as a result of our choice of desired subsystems interconnections and damping. This idea was first advanced for stability analysis in [16]; the extension for controller design was then reported in [17] and [8]; since then many successful applications, including mass-balance systems [18], electrical machines [19], power systems [20], magnetic levitation systems [21], and power converters [22], have been reported.


The aim of this article is to provide a new energy-balancing perspective of PBC that embraces and unifies its classical and modern versions. To enhance readability and widen our target audience, we strip away as much as possible the mathematical details and concentrate instead on the basic underlying principles and limitations. To underscore the fact that the principles are universal, we pres-

ent them in a very general circuit-theoretic framework, without any additional mathematical structure attached to the system models. Particular emphasis is given to exhibiting the physical interpretation of the concepts, for instance, the central role played by dissipation. Toward this end, we illustrate our main points with simple physical examples.

The remainder of the article is organized as follows. First, we review the basic notions of passivity and stabilization via energy shaping. Next, we define the concept of energy-balancing PBC and prove that this principle is applicable to all mechanical systems. Later we show that systems which extract an infinite amount of energy from the controller (i.e., systems with unbounded dissipation) cannot be stabilized with energy-balancing PBCs. To characterize the class of systems that are stabilizable with energy-balancing PBCs and eventually extend PBC to systems with unbounded dissipation,

GP-6

ANALOG COMPUTER



Forever the best experiments to introduce Classic Control Theory, analog circuitry, embedded simulations for design testing and analysis. Still the only way to program linear circuits. For details...

visit our web-site:

www.comdyna.com

COMDYNA, INC.

Barrington, IL tel & fax 847/381-7560

we propose to adopt Willems's "control-as-inter-connection" viewpoint, a perspective that naturally provides a geometric interpretation to the notion of energy shaping. Then, after identifying a class of "admissible dissipations," we view the control action as the interconnection of the system with a passive controller. To stabilize systems with unbounded dissipations, we propose to model the action of the control as a state-modulated power-preserving interconnection of the plant with an infinite energy source system. These developments, which lead to the definition of a new class of PBCs called interconnection and damping assignment PBC, are presented for the so-called port-controlled Hamiltonian systems. Finally, we detail the application of interconnection and damping assignment PBC to a physical example, and then we present some concluding remarks.

CONCLUDING REMARKS

We have given a tutorial presentation of a control design approach for physical systems based on energy considerations that has been developed by the authors of the present article, as well as by some other researchers cited in the references, in the last few years. The main premise of this approach is that the fundamental concept of energy is lost in the signal processing perspective of most modern control techniques; hence, we present an alternative viewpoint that focuses on interconnection. The choice of a suitable description of the system is essential for this research; thus, we have adopted port-controlled Hamiltonian models that provide a classification of the variables and the equations into those associated with phenomenological properties and those defining the interconnection structure related to the exchanges of energy.

There are many possible extensions and refinements to the theory we have presented in this article.

Central among the various open issues that need to be clarified, one finds, of course, the solvability of the PDE. Although we have shown that the added degrees of freedom can help us in its solution, it would be desirable to have a better understanding of their effect, which would lead to a more systematic procedure in their design. For general port-controlled Hamiltonian systems, this is, we believe, a far-reaching problem. Hence, we might want to study it first for specific classes of physically motivated systems.

Solving new problems is, of course, the final test for the usefulness of a new theory. Our list of references is a testament to the breadth of application of our approach; hence, we tend to believe that this aspect has been amply covered by our work.



**TAP.
CONNECT.
NETWORK.
SHARE.**

Connect to IEEE—no matter where you are—with the IEEE App.



Stay up-to-date with the latest news



Schedule, manage, or join meetups virtually



Get geo and interest-based recommendations



Read and download your IEEE magazines



Create a personalized experience

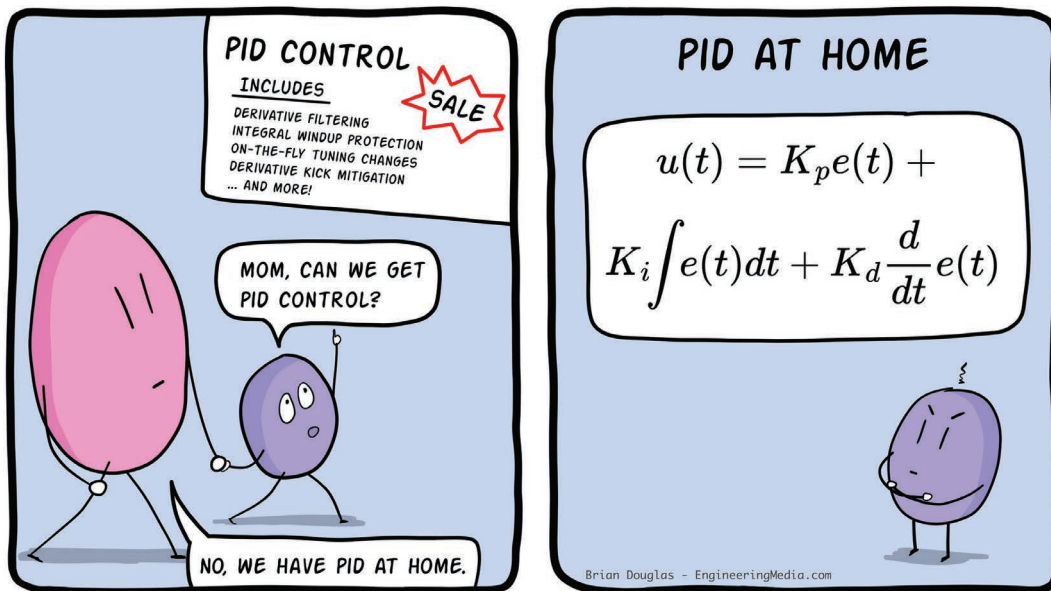


Locate IEEE members by location, interests, and affiliations

Download Today!



We have a brand new comic strip, "The Lighter Side of Control Systems," penned by Brian Douglas, exclusively for the magazine! Enjoy!



Digital Object Identifier 10.1109/MCS.2026.3659162
 Date of current version: 26 March 2026



Member Activities

This column features a contribution from Tariq Samad, Pramod Khargonekar, Peter Fogh Odgaard, Lucy Pao, and Jakob Stoustrup describing their participation in the recent IEEE Clean-Tech Summit held in October 2025 in Glasgow. Tariq, Pramod, Peter, Lucy, and Jakob represented the IEEE Control Systems Society (CSS) at the Summit as panelists addressing the role of controls in mitigating and adapting to climate change. Tariq and Pramod organized the panel; the discussions and presentations

by the panel members were based in part on a previous article in *Control Systems Magazine* (volume 44, issue 3), for which most of the panel participants were coauthors. Many thanks to Tariq, Pramod, and all of the panel members for providing an excellent representation of control technologies and CSS at the Summit.

—Necmiye Ozay ^{ID}

(Vice President Member Activities)

CSS Participates in IEEE Tech Summit 2025

IEEE has launched an annual event series, the IEEE Technology Summit. Different topics will be featured each year. The 2025 inaugural event, held in Glasgow from 2 to 3 October, focused on climate change and clean tech innovation (<https://techsummit.ieee.org/cleantech-2025/>; Figure 1). The Tech Summit attracted about 200 attendees. Participation was global, but Scotland and the United Kingdom were particularly well represented. IEEE leadership, including the president, past president, and president-elect, were present. Several IEEE societies were cosponsors/partners: Aerospace and Electronic Systems, Computer, Geophysical and Remote Sensing, Industry Applications, Product Safety Engineering, and Vehicular Technology.

The IEEE Control Systems Society (CSS) participated in a panel titled “Climate Change Mitigation,

Adaptation, and Resilience—Control Systems as a Critical Enabler.” The panel was organized by Tariq Samad (University of Minnesota, USA) and Pramod Khargonekar (University of California, Irvine, USA); participants also included Peter Fogh Odgaard (Goldwind Energy, Denmark), Lucy

Pao (University of Colorado, Boulder, USA), and Jakob Stoustrup (Aalborg University, Denmark).

Climate, of course, is a complex dynamical system, and if we are to address the adverse impacts resulting from climate change, control principles will need to be applied to myriad



FIGURE 1 The Glasgow Science Center, the venue for the opening reception (with a bag-piper silhouetted against the display).

facets of this problem: mitigation (reducing greenhouse gas emissions), adaptation (making changes to deal with the effects of a warming planet), and resilience (planning and responding to floods, heatwaves, wildfires, etc.). Not only do we need to encourage control engineers and scientists to focus on the topic, but the broader science and technology community and policymakers need to better appreciate the relevance of control systems to deal with these challenges. The panel session was targeted to the IEEE technology leaders, and we were pleased with the level of interest and the significant engagement from our audience.

The topic was recently addressed at some length in an article in *Control Systems* magazine [1] (an outcome of the IEEE CSS Roadmap 2030 report [2]), which the majority of the panel participants helped co-author (Figure 2). To help promote the message, the CSS offered this article free of charge from Xplore for October 2025.

The program included plenary sessions (several of which highlighted sponsoring societies) and breakout sessions. Our session was included in one of the latter (Figure 3). (We would be happy to share our slides on request.)

The agenda for our session was:

- » Tariq Samad offered motivating remarks on climate change as a control systems application, stating that control science and engineering provides a rigorous approach for managing complex dynamical systems. He also highlighted [1] and showed the 14 “targets of opportunity” detailed therein.
- » Pramod Khargonekar presented several clean tech solutions leveraging control systems technology. He also discussed broader themes; for example, the need for cross-cutting transformations across power, energy, social, economic, and Earth systems.
- » Jakob Stoustrup reviewed the Danish Climate Act, targeting a climate-neutral society by 2050 at the latest. Aalborg University is a leading part of that effort, taking a mission-oriented approach based on a broad national partnership. The approach features, in part, a new control architecture for a new energy system.
- » Lucy Pao focused on floating offshore wind turbines and farms. Globally, the majority of offshore wind resources are in deeper waters, where floating support structures are likely to be needed. The challenge is economic, and controls is a key enabler for driving down the cost of floating wind energy.
- » Peter Fogh Odgaard gave an industrial perspective from the wind turbine sector. He noted a number of areas where control has, or could have, a business impact, such as integrating multiobjective optimization, balancing design criteria and business cases, and implementing site-specific adaptive and learning control schemes.

The panel session generated many questions from the audience, more than we had time to handle. It also led to an invitation for Khargonekar to give a presentation, scheduled for




FIGURE 2 Article in *IEEE Control Systems Magazine*, June 2024, that was highlighted in the CSS panel. The article was made open access for the month of October 2025 [1].



FIGURE 3 The session participants, left to right: P. Khargonekar, J. Stoustrup, T. Samad, L. Pao, P.F. Odgaard.

December 2025, jointly with Samad at the IEEE Online Forum on Climate Change Technologies, which attracts an audience of 500 people.

Overall, the summit sessions covered a broad spectrum of technologies to address climate change. Policy, market, and geopolitical considerations were also discussed. Below are a few highlights that should interest CSS members. For more on these and other topics covered, check out the program on the website.

- » Wind energy was especially prominent. Scotland has one of the highest potentials worldwide for offshore wind generation (as also noted in the control session). We also heard that Glasgow has more people working in clean energy than any other city in the United Kingdom. Glasgow, incidentally, is where James Watt invented his steam engine and adapted a centrifugal governor for feedback control.
- » Of course, artificial intelligence was featured in many sessions. One was titled, “The Double-

Edged Sword of Artificial Intelligence in Climate Solutions.” The proliferation of data centers was also discussed, and we heard that 25% of Ireland’s electricity consumption is now for data centers. In our session, we also noted that managing data centers is a significant opportunity for control technologies.

- » The greenhouse gas impact of commercial aviation was discussed in a talk on “Envisioning Near-Zero Air Travel.” A factoid from this presentation: 25%–50% of global warming impact from aviation is from contrails, not fuel burn. A couple of us discussed issues related to sensing and control with the speaker (Dr. Florian Allroggen, Massachusetts Institute of Technology).
- » Nuclear energy had a brief mention. Siemens is working with Rolls-Royce to develop a control system for small modular reactors. There is also research underway on even smaller-scale

nuclear generators: micro modular reactors.

- » Another session featured a deep discussion of the increasing energy efficiency of semiconductor and related technologies.

The summit also included an entrepreneurial pitch competition, with selected participants from around the world presenting their innovations related to climate change in two tracks: “Ideation” and “Startup.” Topics included water-quality sensing (Agri-NOVA Tech, Malaysia), transforming agricultural waste into building materials (Agreka Build, U.K.), atmospheric water extraction (Uravu Labs, India), nutrient mining from food waste (BillionCarbon, India), and transforming agricultural waste into high-value chemicals (P-Vita, Egypt).

Next year’s IEEE Tech Summit is still in the planning stages; it will likely be held in Europe, with artificial intelligence and ethics as a theme. This is another topic with synergies with control, and we encourage CSS participation.

ACKNOWLEDGMENT

We thank the IEEE Control Systems Society officers, especially Carolyn Beck, for their encouragement and support of the session.

—Tariq Samad

University of Minnesota, USA

—Pramod Khargonekar

University of California, Irvine, USA

—Peter Fogh Odgaard

Goldwind Energy, Denmark

—Lucy Pao

University of Colorado, Boulder, USA

—Jakob Stoustrup

Aalborg University, Denmark 

REFERENCES

- [1] P. P. Khargonekar et al., “Climate change mitigation, adaptation, and resilience: Challenges and opportunities for the control systems community,” *IEEE Control Syst. Mag.*, vol. 44, no. 3, pp. 33–51, Jun. 2024, doi: [10.1109/MCS.2024.3382377](https://doi.org/10.1109/MCS.2024.3382377).
- [2] A. M. Annaswamy, K. H. Johansson, and G. J. Pappas, “Control for societal-scale challenges: Road map 2030,” *IEEE Control Syst. Mag.*, vol. 44, no. 3, pp. 30–32, Jun. 2024, doi: [10.1109/MCS.2024.3382376](https://doi.org/10.1109/MCS.2024.3382376).



Technical Committee on Healthcare and Medical Systems

The Technical Committee on Healthcare and Medical Systems (TC-HMS) of the IEEE Control Systems Society (CSS) brings together researchers working at the intersection of control, cyberphysical systems, and medicine. Our portfolio spans acute and chronic care, from anesthesia and critical-care drug delivery to diabetes and artificial pancreas systems, neurotechnology and epilepsy, stroke rehabilitation and assistive devices, oncology, and data-driven clinical decision support. In last year's *IEEE Control Systems Magazine* column, we framed this work in terms of the broader cyberphysical systems and health-care "roadmaps," highlighting how control can help deliver safe, equitable, and trustworthy automation in health and medicine.

Since then, a new wave of artificial intelligence (AI)-in-health "meta-roadmaps" has appeared and is already shaping how TC-HMS thinks about its mission. The National Academy of Medicine's special publication *An Artificial Intelligence Code of Conduct for Health and Medicine: Essential Guidance for Aligned Action* [1] proposes a unifying AI Code of Conduct (AICC) framework and a set of "Code Commitments" (for example, to advance humanity, ensure equity, engage impacted communities, improve workforce well-being, monitor performance, and foster continuous learning). These commitments are intended as simple, interoperable rules for developers, health systems, and payers that reduce the current fragmentation of AI guidance across institutions and

jurisdictions and promote aligned governance across the AI lifecycle.

Complementing this, Canada's Drug Agency 2025 Watch List: Artificial Intelligence in Health Care [2] takes a horizon-scan view: It identifies the top five emerging AI technologies in health care—such as AI for notetaking, clinical training, disease detection and diagnosis, treatment optimization, and remote monitoring—alongside the top five issues, including privacy and data security, liability and accountability, data quality and bias, data sovereignty and governance, and environmental costs. Together with the AICC, this Watch List identifies where AI is most likely to enter clinical workflows in the near term and the oversight, data infrastructure, and workforce capacity needed to deploy these tools responsibly.

On the engineering side, the National Academies' report *Machine Learning for Safety-Critical Applications: Opportunities, Challenges, and a Research Agenda* [3] outlines a research program to integrate machine learning (ML) into systems where failure is unacceptable, explicitly including domains, such as autonomous medical devices, robotic surgery, and high-risk decision support. It highlights open problems in verification, testing, monitoring, and assurance for ML-enabled components embedded in safety-critical control loops. In parallel, the President's Council of Advisors on Science and Technology report, *Strategy for Cyber-Physical Resilience: Fortifying Our Critical Infrastructure for a Digital World* [4], emphasizes the system-level resilience of cyberphysical infrastructures, with clear implications for health-care delivery networks.

Taken together, these documents create a coherent policy and research landscape: The AICC and Watch List articulate what responsible AI in health should look like and where it is likely to be deployed, while the safety-critical ML and cyberphysical resilience reports sketch how we might engineer and assure such systems in practice. TC-HMS naturally sits at this intersection of AI, control, and safety. Our current and planned activities—ranging from conference initiatives to cross-TC collaborations—can be read as attempts to operationalize these external "roadmaps," bringing control-theoretic modeling, analysis, and closed-loop design to bear on concrete health and medical applications.

In 2025, the committee's activities were guided by the theme of "closing the loop to the clinic." On the scientific side, this included organizing an International Federation of Automatic Control (IFAC) workshop on diabetes technology and automated insulin-delivery systems and coleading a major joint set of invited sessions at the American Control Conference (ACC) focused on AI- and control-enabled health care. On the community side, TC-HMS helped launch a new joint student paper award with the CSS Technical Committee on Systems Biology and Biological Systems (TC-SBB) and experimented with lightweight outreach formats, such as a short "Meet & Greet" video series profiling members' work.

The mailing list currently includes 64 active members, with four identified student members and approximately 11% women (7/64). While this represents a vibrant and diverse community in terms of geography and research topics, increasing participation

from women and other historically underrepresented groups remains a clear priority. To that end, the TC aims to use its invited activities, awards, and communication channels to improve visibility and foster a more inclusive pipeline of young researchers.

ACTIVITIES

Workshops and Conferences

A key highlight of 2025 was the first IFAC Workshop on Engineering Diabetes Technology (EDT'25), held in Valencia, Spain, on 8–9 May 2025. The workshop focused on engineering and control for diabetes management, with topics such as automated insulin-delivery algorithms, physiological modeling, multiday therapy design, adjuvant therapies, patient education, and embedded computation for closed-loop systems. In addition to its technical program, the event included young-author and innovation-oriented recognition awards to encourage early-career researchers working at the interface of control, devices, and clinical translation.

TC-HMS also co-organized a major cross-society initiative at the ACC 2025, held in Denver, CO, on 8–10 July 2025. Under the banner of American Society of Mechanical Engineers (ASME)–IEEE Joint Invited Sessions on Healthcare and Medical Systems, these sessions were jointly coordinated by the ASME Dynamic Systems and Control Division Bio-Systems and Healthcare Technical Committee, the IEEE CSS Technical Committee on Process Control, and IEEE CSS TC-HMS. The program spanned modeling, analysis, estimation, and control for biomedical systems, and highlighted contributions on AI and ML for health, brain–computer and neural control interfaces, medical devices and wearables, emergency-response systems, and individualized and personalized treatments. This collaboration strengthened ties between communities that traditionally meet under different umbrellas and provided a broad, systems-level view of health-care engineering.

TC Meetings

TC-HMS continues to hold its main in-person meeting in December, co-located with the IEEE Conference on Decision and Control (CDC). The most recent meeting took place at CDC 2025 in Rio de Janeiro, Brazil (9–12 December 2025), with a focus on AI governance, safety-critical ML, and clinical translation. Building on the strong applications focus of the committee, we are also planning a dedicated TC-HMS meeting at the IEEE Conference on Control Technology and Applications (CCTA 2026) in Vancouver, Canada (12–14 August 2026), which has become one of the best venues for showcasing work that moves control and AI methods from theory into practice in health-care and medical systems.

Other Activities and Outreach

To improve internal visibility and make it easier for members to discover potential collaborators, TC-HMS launched a “Meet & Greet” video series in February 2025. Each short episode combines a brief interview with a 3- to 4-min self-introduction and an accessible summary of the featured member’s research vision. The first seven installments highlighted:

- » Erfan Nozari (Episode #1, 19 February 2025)
- » Emily (Reed) Pereira (Episode #2, 19 February 2025)
- » Alexander Medvedev (Episode #3, 17 March 2025)
- » Yuzhen Qin (Episode #4, 1 April 2025)
- » Marzia Cescon (Episode #5, 29 April 2025)
- » Clara Mihaela Ionescu (Episode #6, 23 May 2025)

Since then, a new wave of artificial intelligence (AI)-in-health “meta-roadmaps” has appeared and is already shaping how TC-HMS thinks about its mission.

- » Corentin Briat (Episode #7, 6 July 2025).

The series is distributed via TC-HMS channels (including LinkedIn), with the dual goal of making the committee more welcoming to newcomers and highlighting the breadth of research being carried out by its members.

MEMBERS’ RECOGNITIONS

TC-HMS members received a number of distinctions in 2025, reflecting their leadership in research, teaching, and service across both control and biomedical engineering:

- » Naira Hovakimyan (University of Illinois Urbana-Champaign) received the Grainger College of Engineering Award for Excellence in Graduate Student Mentoring for 2025, recognizing her sustained contributions to training the next generation of researchers.
- » Warren Dixon (University of Florida) was appointed Interim Dean of the Herbert Wertheim College of Engineering in July 2025, underscoring both his scientific standing and his leadership within the broader engineering community.
- » Elisa Franco (University of California, Los Angeles) received the Mechanical and Aerospace Engineering Faculty Teaching Award in 2025 and served as a plenary speaker at ACC 2025, highlighting her impact in both research and education at the interface of control, synthetic biology, and molecular systems.
- » Juergen Hahn (Rensselaer Polytechnic Institute) was elected Fellow of the Asia-Pacific Artificial Intelligence Association and

TC-HMS naturally sits at this intersection of AI, control, and safety.

named a Member of the European Academy of Sciences and Arts in 2025, recognizing his pioneering work in data-driven modeling and analysis for biomedical applications.

- » Alexander Medvedev (Uppsala University) was awarded a Swedish Research Council project grant (2025–2027) for the project “Pulse-Modulated Feedback Control of Nonlinear Infinite-Dimensional Positive Systems,” supporting fundamental research with potential implications for advanced therapeutic and monitoring technologies.
- » Robert D. Gregg (University of Michigan) delivered a keynote lecture at the International Conference on Robotics and Automation 2025, reflecting his international leadership in the development of control and robotics methods for rehabilitation and assistive devices.

In addition to these individual honors, TC-HMS and the TC-SBB jointly established the IEEE CSS TC-HMS & TC-SBB Outstanding Student Paper Award 2025. The inaugural award was presented to Fabian Flürenbrock for the paper “Model-Based Estimation of Ventricular Cerebrospinal Fluid Volume,” coauthored with Simon Muntwiler, Leonie Korn, Marianne Schmid Daners, and Melanie N. Zeilinger. Nominated by Prof. Zeilinger, the paper was selected by a jury composed of Sérgio Pequito (chair), Jürgen Hahn, Elisa Franco, Alexander Medvedev, and Meriem Laleg. The committee’s decision highlighted the work’s scientific merit, methodological rigor, and clear relevance to healthcare and medical systems.

Together, these recognitions illustrate the depth and breadth of the

TC-HMS community—ranging from foundational control theory to translational work in clinical and industrial settings—and reinforce the committee’s mission to advance healthcare through control-inspired thinking and technology.

MEMBERS’ PROJECT HIGHLIGHTS

Every year, TC-HMS showcases a small set of flagship projects that capture the breadth of research conducted by our members at the interface of control, AI, and medicine. This year’s “Members’ Project Highlights” section focuses on three complementary stories along the patient journey: from safe, intelligent drug delivery in the operating room, through lifestyle-based prevention of type 2 diabetes, to model-based monitoring and treatment of hydrocephalus. Together, they illustrate how principled modeling, control, and data-driven methods can be turned into clinically meaningful tools in high-stakes settings.

The three vignettes that follow were chosen deliberately to span scales and modalities. Clara Mihaela Ionescu’s European Research Council Consolidator project AMICAS (Adaptive Multi-drug Infusion Control System for General Anesthesia in Major Surgery) tackles the real-time coordination of multiple anesthetic drugs during major surgery; Alessandro Borri and collaborators develop a multitimescale model that links physical activity, inflammation, and long-term diabetes progression; and the 2025 TC-HMS & TC-SBB Outstanding Student Paper Award recognizes a model-based estimator for ventricular cerebrospinal fluid (CSF) volume, a key building block for “smart” shunt systems in hydrocephalus care. Each highlight is written by the team behind the work and is introduced below.

Adaptive Multidrug Infusion Control for General Anesthesia

Modern anesthesia is, in practice, a demanding multivariable control problem. During major surgery, anesthesiologists must simultaneously regulate unconsciousness, analgesia, and hemodynamics in patients whose responses are highly individual and time-varying. Clara Mihaela Ionescu’s European Research Council Consolidator project AMICAS tackles this challenge by bringing adaptive, feasibility-aware control into the operating room [5]. Instead of relying solely on fixed population models and expert intuition, AMICAS uses closed-loop algorithms that adapt drug infusion rates in real time to each patient’s physiological response, explicitly accounting for variability in drug sensitivity, metabolism, comorbidities, and surgical stress. A conceptual overview of this computer-guided control framework is shown in Figure 1.

The AMICAS architecture combines multivariable patient models, online parameter adaptation, and constrained optimal control to individualize anesthesia while guaranteeing safety and robustness under uncertainty. A virtual-patient simulation platform supports large-scale *in silico* testing and personalization prior to clinical deployment, enabling exhaustive exploration of drug combinations, disturbance scenarios, and controller tuning that would not be feasible at the bedside. The long-term vision is a “shared-control” paradigm in which the anesthesiologist supervises and overrides a decision-support system that handles much of the low-level, continuous control.

Beyond improving performance in well-resourced centers, AMICAS is motivated by global disparities in access to anesthesia expertise. By embedding advanced feedback and predictive control into intelligent support tools, the project aims to narrow gaps in training and experience across health systems and geographies, offering consistent, data-driven assistance

to clinicians regardless of local resources. This vision extends to related work in Clara Ionescu’s group on objective, wearable pain monitoring, which seeks to move from episodic, subjective pain scores to continuous, sensor-based indices using smart-textile technologies. This evolution from discrete, self-reported assessments to unobtrusive, continuous tracking of pain dynamics is illustrated in Figure 2. Together, these efforts show how control, modeling, and sensing can turn anesthesia and pain management into transparent, quantifiable, and ultimately more equitable components of modern health care.

Turning Physical Activity Into a Long-Term Control Input for Diabetes Prevention

Type 2 diabetes is often framed as an almost unavoidable progression from impaired glucose regulation to overt disease, yet decades of clinical evidence show that lifestyle—particularly physical activity—can substantially delay or even prevent onset. What has been missing is a quantitative framework that links realistic activity patterns to long-term glucose regulation in a form usable for control design. In their IEEE *Control Systems Letters* article “The Long-Term Effects of Physical Activity on Blood Glucose Regulation: A Model to Unravel Diabetes Progression”—which received the 2024 IEEE CSS TC-HMS Outstanding Student Paper Award [6]—Alessandro Borri and collaborators develop a multitimescale model where short-term exercise bouts interact with slow metabolic changes over years via antiinflammatory mediators, such as interleukin-6 (IL-6). This model captures how repeated physical-activity “pulses” can reshape β -cell function and insulin sensitivity over the long term, offering a mechanistic explanation for why regular exercise can meaningfully delay the onset of diabetes.

In their more recent work “Long-Term Diabetes Prevention via Physical Activity: An Output-Feedback MPC Approach” [7], the team takes a decisive step from modeling to control

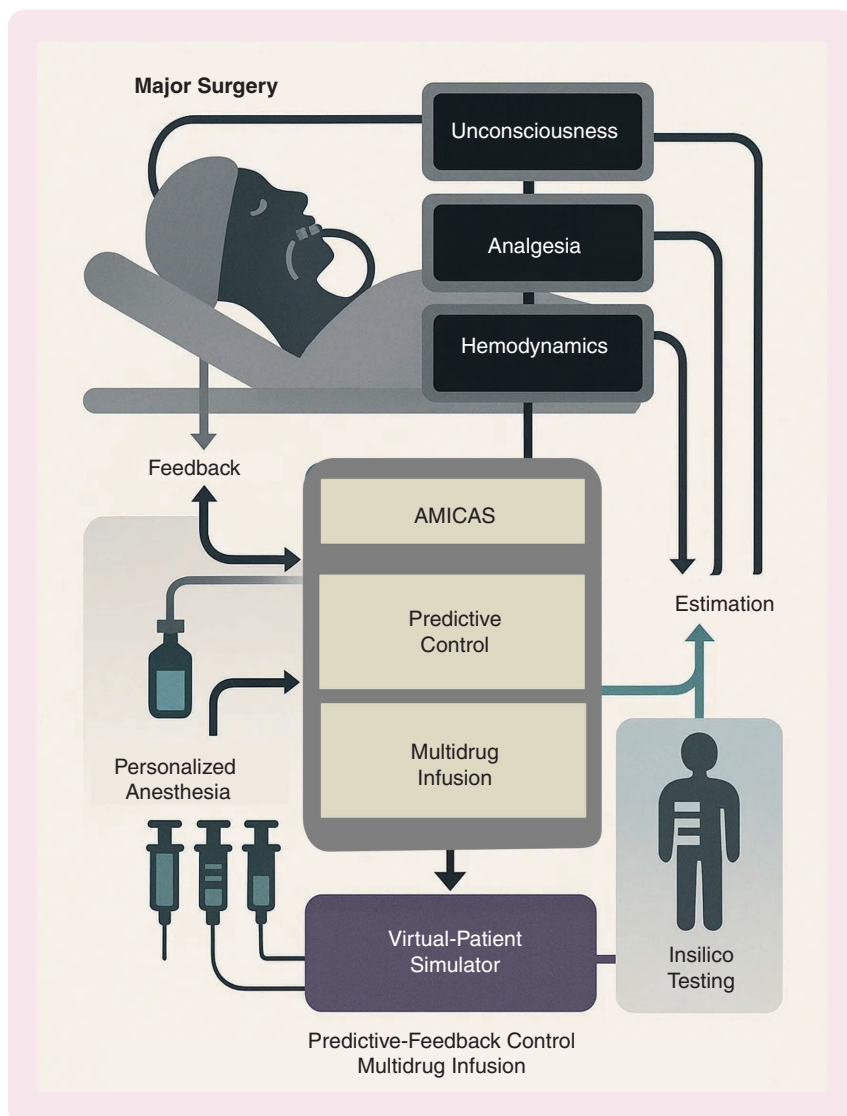


FIGURE 1 Computer-guided control framework for multidrug general anesthesia. The anesthesiologist supervises a closed-loop controller that adjusts multiple drug infusions based on real-time patient monitoring. A virtual-patient simulation environment supports large-scale testing and patient-specific personalization before clinical use.

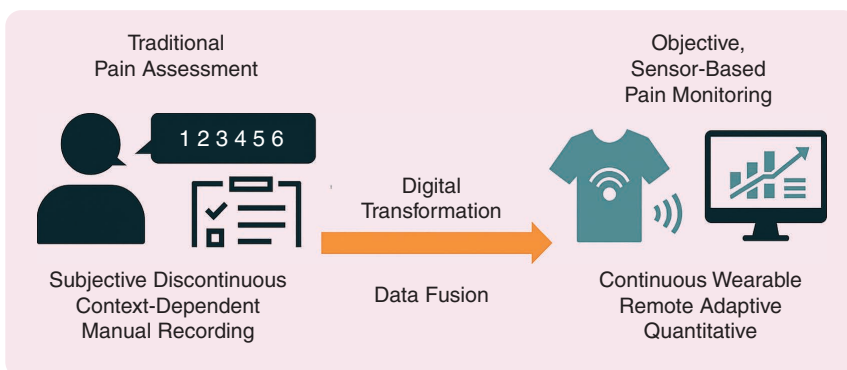


FIGURE 2 From episodic to continuous pain assessment. Conceptual illustration of the transition from traditional, subjective pain scoring during clinical visits to wearable, objective, and continuous monitoring, enabling unobtrusive tracking of pain dynamics across clinical and everyday settings.

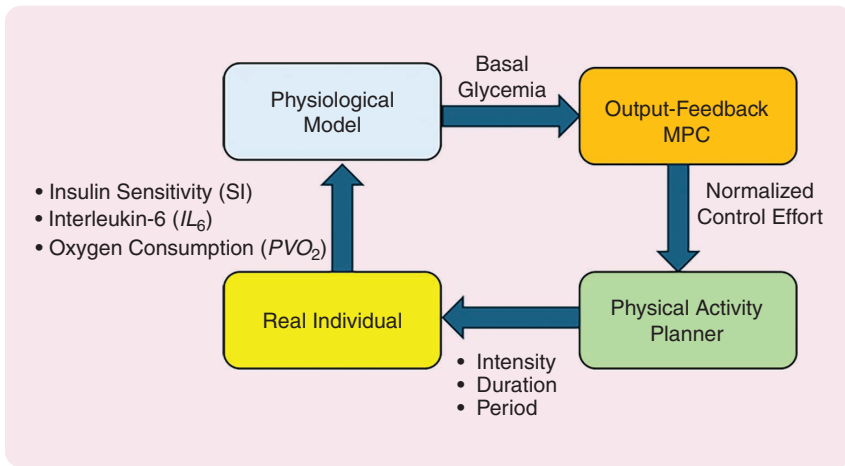


FIGURE 3 Conceptual feedback loop for lifestyle-based diabetes prevention. Physical activity serves as a control input that triggers IL-6 release and other exercise-induced mechanisms, which in turn influence β -cell function, insulin sensitivity, and long-term glucose regulation. The multimescale model links short-term exercise sessions to slow metabolic progression, providing the prediction engine for output-feedback MPC design.

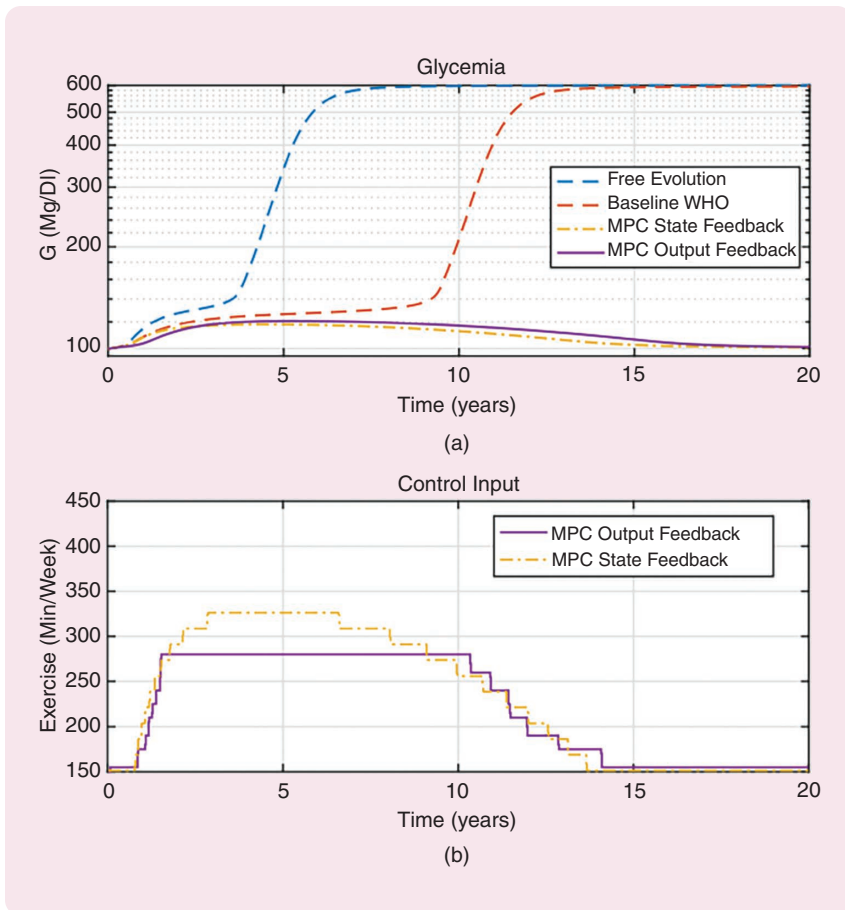


FIGURE 4 Example of closed-loop simulation for an average virtual prediabetic patient. (a) Glycemia trajectories for free evolution (no intervention), a baseline WHO-inspired constant-activity recommendation, and the proposed MPC state- and output-feedback controllers. The MPC-based strategies maintain normoglycemia and delay diabetes onset over 20 years. (b) Corresponding control input, expressed as prescribed exercise (minutes per week), for the MPC state- and output-feedback controllers, which gradually taper the activity level toward a long-term maintenance regime.

by treating physical activity itself as a long-term control input. Using the multimescale physiological model as a prediction engine, they design output-feedback model predictive controllers (MPC) that recommend weekly activity targets—expressed in terms of frequency, duration, and intensity of exercise sessions—tailored to an individual’s metabolic state and history. Rather than searching for a single “optimal” schedule, the framework explicitly accommodates interindividual variability and day-to-day fluctuations in adherence, working within guideline-compatible ranges of physical activity. The controller observes a clinically accessible output, for example, fasting glucose, updates internal state estimates, and proposes feasible activity plans to keep an individual in a normoglycemic regimen over a multiyear horizon. A schematic of the resulting closed-loop architecture, from physical activity to IL-6 dynamics and long-term glucose regulation, is shown in Figure 3.

Large-scale in silico experiments on virtual prediabetic populations suggest that this closed-loop, lifestyle-based strategy can substantially reduce the fraction of individuals who progress to diabetes under realistic activity regimens. An example simulation for an “average” virtual patient under the proposed controllers is depicted in Figure 4: Figure 4(a) shows glycemia trajectories under free evolution, a baseline World Health Organization (WHO)-inspired recommendation, and the MPC state- and output-feedback strategies, while Figure 4(b) displays the corresponding prescribed exercise (minutes per week) for the two MPC controllers. Beyond the specific numerical results, the concept is powerful: Physical activity is no longer just a generic “do more exercise” recommendation but becomes a dosage-controllable therapy, governed by the same predictive, feedback-oriented principles that underlie modern drug-delivery and process-control systems. Borri and colleagues envision such controllers as the algorithmic core of future

decision-support tools and digital companions that help clinicians and patients codesign sustainable, personalized physical-activity prescriptions for long-term diabetes prevention.

Model-Based Estimation of Ventricular CSF Volume for Smart Hydrocephalus Shunts

The 2025 IEEE CSS TC-HMS & TC-SBB Outstanding Student Paper Award was given to “Model-based Estimation of Ventricular Cerebrospinal Fluid Volume” by Fabian Flürenbrock, Simon Muntwiler, Leonie Korn, Marianne Schmid Daners, and Melanie N. Zeilinger [8]. Hydrocephalus is a medical condition characterized by disturbed CSF dynamics and an excessive accumulation of CSF in the brain’s ventricles. Without appropriate therapy, such as a shunt system that drains excessive CSF from the ventricles to another body compartment like the abdomen, patients suffer from alterations in intracranial pressure that can lead to severe neurological and physical impairments [9]. While assessing ventricular CSF volume is key to diagnosing and monitoring hydrocephalus, current clinical practice relies on medical imaging technologies, such as magnetic resonance imaging and computed tomography to obtain this assessment.

To improve continuous monitoring for hydrocephalus patients, researchers from the Intelligent Control Systems Group at ETH Zurich developed a model-based estimator for real-time estimation of ventricular CSF volume [8]. In contrast to other approaches, the presented estimator enables patient-specific calibration and sensor fusion across measurement modalities, including intracranial pressure, CSF flowthrough a shunt system, and intraventricular bioimpedance.

A comprehensive three-step pipeline is used to derive the model-based CSF volume estimator. 1) *Optimal experiment design*: A medical intervention is designed that maximizes the information about a patient’s CSF dynamics while maintaining safety. 2) *System*




FIGURE 5 Schematic representation of the experimental test setup, comprising the mechatronic test bench developed in [11] and an illustration of the tetrapolar measurement principle to acquire intraventricular bioimpedance (BI). The intracranial pressure (ICP) and BI measurement catheter (shown in red) is inserted into the ventricle of the silicone-carbon brain phantom (shown in black) and connected to the custom-built measurement system. Changes in the ventricular CSF volume are induced by compressing the connected bellows using a feedback-controlled linear voice coil motor. The BI measurement principle relies on an excitation current i_{BI} (shown as dashed orange arrows) injected at the two outer electrodes of the measurement catheter, which induces a voltage drop v_{BI} across the two inner electrodes. Since the electrical conductivity of CSF (σ_{CSE}) is about one order of magnitude higher than that of brain tissue (σ_{Brain}), increases in ventricular CSF volume cause decreases in BI.

identification: The unknown parameters of a state-space model describing the patient’s CSF dynamics are identified using maximum likelihood estimation, utilizing expectation-maximization and nonlinear programming. 3) *State estimation*: A Kalman filter is designed that combines the identified patient model with multimodal measurement data for real time estimation of the CSF volume.

The proposed model-based estimator was validated through experimental testing with a custom-built measurement system and a mechatronic test bench that replicates CSF dynamics in vitro (see Figure 5). Achieving a normalized root-mean-square estimation error below 6%, the CSF volume estimator demonstrates its potential as a clinical tool to improve hydrocephalus patient monitoring. Integration into mechatronic shunt systems capable of active CSF drainage control [10] could further enable closed-loop CSF volume regulation, thereby laying the foundation for a more physiological hydrocephalus shunt therapy.

In summary, the work featured in this column captures where TC-HMS

is heading: from policy-aware, safety-critical AI and adaptive anesthesia, through lifestyle-based diabetes prevention, to smart hydrocephalus shunts that regulate ventricular volume directly. Across very different clinical contexts, the common thread is the use of modeling, control, and learning to build health-care systems that are safer, more transparent, and more equitable in practice.

Membership in the TC-HMS is free. If your work touches on control, optimization, AI/ML, or cyberphysical systems in health and medicine, we warmly invite you to join. To do so, simply visit the TC-HMS page on the CSS website and subscribe to the mailing list, or contact the TC chair by e-mail (sergio.pequito@tecnico.ulisboa.pt) to be added. Students, early-career researchers, clinicians, and industry colleagues are all welcome. We look forward to seeing you at future TC-HMS meetings at CDC and CCTA, as well as related events.

Sérgio Daniel Gonçalves
Melo Pequito

REFERENCES

- [1] National Academy of Medicine, "An artificial intelligence code of conduct for health and medicine: Essential guidance for aligned action," The Nat. Academies Press, Washington, DC, USA, 2025. [Online]. Available: <https://www.nationalacademies.org/publications/29087>
- [2] "2025 watch list: Artificial intelligence in health care: Health technologies," Can. Agency for Drugs and Technologies in Health, Ottawa, ON, Canada, Report No.: ER0015, Mar. 2025.
- [3] National Academies of Sciences, Engineering, and Medicine, "Machine learning for safety-critical applications: Opportunities, challenges, and a research agenda," The Nat. Academies Press, Washington, DC, USA, 2025, doi: [10.17226/27970](https://doi.org/10.17226/27970). [Online]. Available: <https://www.nationalacademies.org/publications/27970>
- [4] President's Council of Advisors on Science and Technology, "Strategy for cyber-physical resilience: Fortifying our critical infrastructure for a digital world," Executive Office of the President, Washington, DC, USA, Feb. 2024. [Online]. Available: https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/02/PCAST_Cyber-Physical-Resilience-Report_Feb2024.pdf
- [5] C. M. Ionescu, "Adaptive multi-drug infusion control system for general anesthesia in major surgery (AMICAS)," Eur. Commission, Brussels, Belgium, ERC Consolidator Grant 101043225, 2022–2027, project record on CORDIS. [Online]. Available: <https://cordis.europa.eu/project/id/101043225>
- [6] P. F. De Paola, A. Paglialonga, P. Palumbo, K. Keshavjee, F. Dabbene, and A. Borri, "The long-term effects of physical activity on blood glucose regulation: A model to unravel diabetes progression," *IEEE Contr. Syst. Lett.*, vol. 7, pp. 2916–2921, 2023, doi: [10.1109/LCSYS.2023.3290774](https://doi.org/10.1109/LCSYS.2023.3290774).
- [7] P. F. De Paola, A. Borri, A. Paglialonga, F. Dabbene, and P. Palumbo, "Long-term diabetes prevention via physical activity: An output-feedback MPC approach," *IEEE Contr. Syst. Lett.*, vol. 9, pp. 679–684, 2025, doi: [10.1109/LCSYS.2025.3578031](https://doi.org/10.1109/LCSYS.2025.3578031).
- [8] F. Flürenbrock, S. Muntwiler, L. Korn, M. Schmid Daners, and M. N. Zeilinger, "Model-based estimation of ventricular cerebrospinal fluid volume," in *Proc. IEEE Conf. Control Technol. Appl. (CCTA)*, 2024, pp. 302–309, doi: [10.1109/CCTA60707.2024.10666541](https://doi.org/10.1109/CCTA60707.2024.10666541).
- [9] D. Rigamonti, *Adult Hydrocephalus*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [10] F. Flürenbrock et al., "VIEshunt: Towards a ventricular intelligent and electromechanical shunt for hydrocephalus therapy," *Fluids Barriers CNS*, vol. 22, Mar. 2025, Art. no. 28, doi: [10.1186/s12987-025-00629-w](https://doi.org/10.1186/s12987-025-00629-w).
- [11] C. C. Wemmers, F. Flürenbrock, B. Maurer, A. Benninghaus, K. Radermacher, and S. Leonhardt, "A mechatronic test-bench to investigate the impact of ventricular pulsation in hydrocephalus," *Biomed. Signal Process. Control*, vol. 75, May 2022, Art. no. 103579, doi: [10.1016/j.bspc.2022.103579](https://doi.org/10.1016/j.bspc.2022.103579).

Some Activities of the Technical Committee on Stochastic Systems and Control in 2025

THE 19TH IEEE INTERNATIONAL SYSTEMS CONFERENCE

On 7 April 2025, during the 19th IEEE International Systems Conference (SysCon 2025), in Montréal, Canada, Tyrone Duncan, Bozenna Pasik-Duncan, and Hamidou Tembine organized the session titled "Mastering Uncertainty: Stochastic Systems and Rosenblatt Noise."

The study of stochastic systems is inherently interdisciplinary. It bridges mathematics, engineering, computer science, economics, and biomedical sciences. These systems are pivotal. They help us understand brain disorders. They model volatile financial markets. They are essential for designing resilient autonomous vehicles and smart networks. However, a fundamental challenge remains. We must manage randomness in interactive environments. Standard models often fall short. Real-world disturbances are rarely simple.

Digital Object Identifier 10.1109/MCS.2026.3659231
Date of current version: 26 March 2026

The Tutorial

On the afternoon of 7 April 2025, experts gathered to address this challenge. The tutorial, titled "Stochastic Systems, Control, and Game Theory with Rosenblatt Noise," offered an introduction to advanced modeling. The session focused on Rosenblatt noise, which is a specific type of stochastic disturbance that is non-Gaussian, non-Poissonian, and non-Markovian. Capturing the nature of this noise is vital. [Figure 1](#) exposes the failure of the "approximation-first" strategy to maintain moment structures, contrasting it with the "optimization-first" approach that succeeds by respecting the system's complex temporal dynamics.

The organizers guided participants to discuss the following key topics through three core areas, namely: 1) foundations of systems influenced by nonstandard noise processes, 2) an overview of stochastic optimal control and mean-field-type game theory, and 3) techniques to apply these theories to real data.

The tutorial was not purely abstract, but emphasized applications. Speakers explored how advanced stochastic methods inform design strategies. Key discussions included: 1) "Smart Energy Systems: Managing Uncertainties in Renewable Energy Grids"; 2) "Blockchain Economics: Modeling Token Velocity and Decentralized Finance Behaviors"; and 3) "Decision Making: Improving Control Strategies Under Deep Uncertainty."

The session aimed to empower the next generation of researchers. The speakers brought extensive mentoring experience. They provided a solid foundation for professional careers. The primary goals were to introduce advanced stochastic modeling tools, highlight applications in energy and finance, and equip attendees for interdisciplinary research.

Tyrone Duncan, Bozenna Pasik-Duncan, and Hamidou Tembine led the session ([Figure 2](#)). The atmosphere was collaborative. Participants engaged in hands-on exercises. The

tutorial successfully bridged the gap between mean-field-type game theory and pressing engineering problems.

INVITED SESSION ORGANIZED AT THE CONFERENCE ON CONTROL, DECISION, AND INFORMATION TECHNOLOGIES 2025

Thanks to the great opportunities provided by the IEEE 11th International Conference on Control, Decision, and Information Technologies (CoDIT 2025, held in Malta, 15–18 July 2025), this Technical Committee (TC) organized a special session (organized by Bozenna Pasik-Duncan and George Yin) devoted to stochastic systems, control, optimization, and applications. We brought together researchers from multidisciplinary communities in control and systems theory, applied stochastic processes, financial economics, actuarial sciences, applied mathematics, applied probability, biology, electrical engineering, ecology, and networked science, to review and substantially update the most recent progress. For broader impacts, we invited researchers from different areas that span a wide range of subfields. The talks were theoretically oriented to application-intensive presentations.

The motivations for putting together the proposed special session stem from emerging and existing applications in control systems, communication networks, signal processing, queueing networks, production planning, biological systems, ecosystems, financial engineering, large-scale systems involving mean-field interactions, hybrid systems under the influence of random environments, and machine learning and neural network applications to control and systems theory.

Coverage of the Special Session

The talks reflected some of the most interesting topics in a focused area, namely, stochastic control. Yet, they touched upon many cutting-edge research topics with a wide range of applications. The proposed session

proved to be invaluable to the scientific community. It could contribute to the subsequent development of

stochastic systems theory and practice and have broader impacts on a wide variety of interdisciplinary fronts.

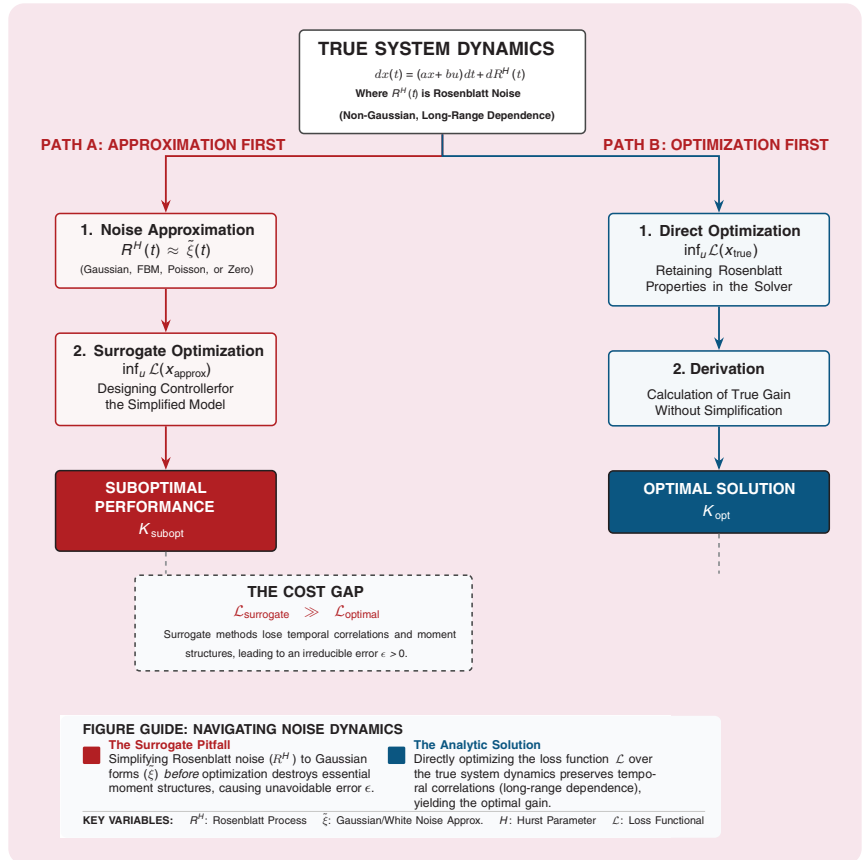


FIGURE 1 Understanding this noise is critical for systems with long-range dependence. This figure contrasts the inevitable performance loss of the “approximation-first” strategy against the “optimization-first” framework that yields the true solution.



FIGURE 2 During the IEEE Conference on Decision and Control (CDC) 2025 held in Rio de Janeiro, Brazil. From left to right: Dr. Ali Pakniyat from the University of Alabama, USA; Prof. Bozenna Pasik-Duncan from University of Kansas, USA; Dr. Julian Barreiro-Gomez from Khalifa University, United Arab Emirates; and Dr. Ricardo de Lima Ribeiro from King Abdullah University of Science and Technology, Kingdom of Saudi Arabia.

The primary goals were to introduce advanced stochastic modeling tools, highlight applications in energy and finance, and equip attendees for interdisciplinary research.

More Details

Speaker Pawel Domanski (of Warsaw University of Technology, Warsaw, Poland) presented “The Need for Non-Gaussian Noise in Control System Models, Why Non-Gaussian Noise Matters” (joint work with Tyrone Duncan and Bozenna Pasik-Duncan of University of Kansas, Lawrence, Kansas, USA). Lukasz Stettner (of Polish Academy of Sciences) presented “Markov Control of Continuous Time Markov Processes with Long Run Functional by Time Discretization.” Marcin Pitera (Jagiellonian University, Kraków) presented “Blackwell Optimality in Risk-Sensitive Stochastic Control.” Chao Zhu (of University of Wisconsin, Milwaukee) presented “Optimal Harvesting Problems with Mean Field Interactions” [joint work with Kurt L. Helmes (Humboldt University of Berlin, Germany) and Richard H. Stockbridge (University of Wisconsin-Milwaukee)]. Nhu Nguyen (of University of Rhode Island) presented “An Application Stochastic Approximation to Decision-Making Problems.” Zhuo Jin (Macquarie University, Australia) presented “Optimal Risk Mitigation Strategies for Cyber Contagion in Networks: A Hybrid Deep Learning Method.”

The session was well received. All of the speakers and attendees indicated that we need to continue this trend of research. In fact, work has been done to organize the subsequent invited session for CoDIT 2026.

INTEGRATING RESEARCH AND TEACHING WITH ARTIFICIAL INTELLIGENCE AND ETHICS





As a pilot program for TC activities, we launched a new series of talks for students at the University of Kansas. A number of researchers from other institutions were invited to give talks in the series. The objective is to bridge research, classroom practice, and control theory applications. Among the talks, we mention in particular the talk given by Tembine Hamidou on ethical responsibility in artificial intelligence (AI)-enabled control systems. The focus is on critical thinking, not automation. AI is presented as an accelerator of insight, not a substitute for reasoning. Not yet.

A flagship contribution to this series was Prof. Tembine Hamidou’s invited lecture, “Control and AI for Graduate Students,” delivered within the course on Stochastic Adaptive Control. The talk combined theory, practice, and ethics. It addressed how AI tools reshape research workflows, learning habits, and decision making in control. The lecture emphasized a central message: “Overreliance on AI weakens critical thinking.”

When answers appear instantly, reflection can disappear. Students were encouraged to treat large learning models as tools, not crutches. AI must amplify human reasoning and insight, rather than supplanting professional judgment and responsibility. Blind copying was explicitly discouraged. Responsible use was framed as a professional obligation.

The session was highly interactive. Students worked through short exercises on AI-assisted debugging, data analysis, and scientific writing. Each exercise required human verification and theoretical justification. This design reinforced accountability. It connected AI usage directly to graduate-level research practice. Student feedback highlighted several takeaways: 1) The talk raised awareness of societal dependence on AI. It prompted reflection on personal research habits. 2) It demonstrated practical, disciplined, and ethical usage of AI tools. The examples were concrete and immediately applicable. The material was well aligned with stochastic adaptive control, yet broadly relevant across disciplines.

We will continue to make these efforts. Ethical risk awareness will be integrated into research activities and homework assignments. Workshops and tutorials are planned for students and young professionals. These activities are aimed at the 2026 International Federation of Automatic Control (IFAC) World Congress, IEEE Conference on Decision and Control 2026, and other IEEE and IFAC partner conferences. As probability and statistics become foundational in modern control and AI, we also plan a webinar series. This series will be developed in collaboration with world-class statisticians and probabilists. The talk series on “Ethical Concerns in Data Science” was inaugurated by Prof. Richard De Veaux. He is the author of widely used textbooks on data modeling and introductory statistics. His lecture set a high standard. It emphasized responsibility, transparency, and ethical reasoning in data-driven research. Prof. De Veaux will continue to inspire and motivate students through future contributions to this series.

Bozenna Pasik-Duncan 
Hamidou Tembine 
George Yin 


Report on Diversity, Outreach, and Development Activities

The IEEE Control Systems Society, through its Diversity, Outreach, and Development Activities (DODA), plays a pivotal role in fostering a vibrant and inclusive global community in automatic control. At the heart of these efforts lies the Outreach Fund, a program that annually supports more than 10 innovative initiatives worldwide.

These projects cover a wide spectrum: from inspiring the next generation of students and researchers to explore control engineering, to organizing scientific workshops and summer schools in regions where access to such opportunities is limited.

This article spotlights one recent success story: a DODA-backed initiative that partnered with the Instrumentation, Control, and Automation Research Group at the Institut Teknologi Bandung in Indonesia to engage and empower Indonesian

students and young professionals in the field of control systems. The Workshop Series on Instrumentation, Control, and Automation for Future Indonesian Scientists and Engineers was held in two parts, providing the opportunity to expand exposure to systems and control engineering across Indonesia.

Future issues of this magazine will showcase additional impactful initiatives supported by the Outreach Fund. If this inspires you to design and lead your own DODA activity, detailed guidelines and application information are available at: <https://ieeecss.org/activities/control-systems-society-outreach-fund>

Silvia Mastellone 

*Vice-President, Diversity, Outreach,
and Development Activities*

Expanding Horizons in Systems and Control for Indonesia's Young Scientists

The Workshop Series on Instrumentation, Control, and Automation (ICA) for Future Indonesian Scientists and Engineers was held in two parts in 2025, supported by the IEEE Control Systems Society (CSS) Diversity, Outreach, and Development Activities (DODA) Program. The first workshop took place in Bali on 3 July 2025 in conjunction with the Asian Control Association (ACA) Forum, while the second was orga-

nized in Bandung on 27–29 August 2025 alongside the ninth International Conference on Instrumentation, Control, and Automation (ICA 2025). The series was initiated by the ICA Research Group of the Faculty of Industrial Technology, Institut Teknologi Bandung (ITB), with the aim of expanding exposure to systems and control engineering across Indonesia, particularly among students and young researchers.

Although systems and control underpin many aspects of modern engineering, awareness of the field in

Indonesia remains uneven. Students and early-career academics from institutions located outside major urban and academic hubs, especially in the eastern part of the Indonesian archipelago, often have limited access to advanced courses, research groups, or professional networks in this area. Participation by Indonesian-affiliated researchers in major international events, such as the IEEE Conference on Decision and Control (CDC), IEEE Conference on Control Technology and Applications (CCTA), and the International Federation of Automatic

Although systems and control underpin many aspects of modern engineering, awareness of the field in Indonesia remains uneven.

Control (IFAC) World Congress also remains modest. National-level activities in control, while steadily growing, are still developing in visibility and reach. These conditions motivated the creation of this workshop series to introduce aspiring Indonesian scientists and engineers to fundamental perspectives, emerging research directions, and global developments in systems and control.

The workshop series was designed to broaden exposure to state-of-the-art research and industrial practice in control systems, strengthen regional academic communities in instrumentation and automation, and create opportunities for young Indonesian

researchers to engage with the ACA and the CSS community. It also aimed to inspire the next generation of control engineering leaders, particularly those from eastern and western regions of Indonesia who are geographically distant from the capital.

BALI WORKSHOP (SERIES 1)

The first workshop was organized as part of the ACA Forum, a major regional gathering for academic leaders in the field of control systems and a pre-event of the 15th Asian Control Conference, which Indonesia will host in 2026. Bali was chosen to facilitate participation by young scientists and students residing in eastern Indonesia,

where opportunities to engage with the international control community are relatively limited. The morning session featured a dedicated workshop series delivered by national experts from ITB, including Prof. Yul Yunazwin Nazaruddin, Dr. Ayu G. Risangtuni, Dr. Vebi Nadhira, Dr. Azka Muji Burohman, Dr. Augie Widyotriatmo, and Prof. Endra Joelianto. Their presentations covered a wide range of emerging themes, artificial intelligence (AI) integration in instrumentation and automation, intelligent and assistive control systems, image-based measurement technologies, data-informative approaches for system analysis and model reduction, autonomous robotic decision-making, and cyber-resilient Internet of Things and automation, providing participants with exposure to current developments in modern control engineering. The session concluded with a hands-on, interactive discussion among participants.

In the afternoon, the workshop continued with the ACA Forum 2025: “Future Technologies in a Rapidly Changing World,” featuring invited lectures from prominent researchers across Asia, who are also steering committee members of ACA, including Prof. Lei Guo (Bode Lecture Prize Recipient 2019), Prof. Shuzhi Sam Ge (former Board of Governors CSS), Prof. Keum-Shik Hong, Prof. Zhixin Liu, and Prof. Qing-Long Han (IEEE Fellow under CSS), along with contributions from other leading regional speakers. The workshop attracted more than 30 participants from five universities in Bali and East Java. Beyond the technical program, the event also provided a valuable networking platform for participants from eastern and western Indonesia and members of the broader Asian control community (Figure 1).

Structured into three thematic forums, “Extending Control Horizon,” “Aiming Autonomous Agents,” and “Having Intelligent Systems,” the ACA Forum showcased a diverse spectrum of frontier research topics, ranging from adaptive and learning-based control, stochastic and game-theoretic



FIGURE 1 Group photo of workshop participants with speakers from ACA professors.



FIGURE 2 Group photo of Indonesian workshop-speakers. From left: Augie Widyotriatmo, Ayu Gareta, Vebi Nadhira, Endang Juliastuti, Yul Yunazwin, Azka Muji Burohman, and two participants.

approaches, and quantum estimation, to resilient multiagent coordination, autonomous mobility technologies, neuromodulation, and data-driven intelligent systems. Through these rich discussions, participants gained meaningful exposure to current research directions and emerging applications that will shape the future of control engineering. Overall, the workshop series strengthened regional academic capacity and inspired young researchers to engage more actively with the Asian control community and pursue interdisciplinary collaborations (Figure 2).

The sessions introduced participants to modern control topics and showcased applications being developed in Indonesia, including autonomous ground vehicles and intelligent instrumentation systems. Many attendees enthusiastically engaged in discussions on how advanced control techniques relate to emerging technologies and national needs. For several participants, this workshop represented their first opportunity to interact directly with internationally recognized researchers in the field (Figure 3).

BANDUNG WORKSHOP (SERIES 2)

The second workshop was held in Bandung as an associated event of the ninth International Conference on Instrumentation, Control, and Automation 2025, a biennial conference organized by the ICA Research Group since 2009 (Figure 4). ICA has consistently served as a national focal point for Indonesian researchers in control, instrumentation, and related fields, providing a venue for exchanging ideas with international experts.

The Bandung program began with opening remarks by Prof. Karl H. Johansson, who presented an introduction to CSS and the DODA program. The technical program continued with Augie Widyotriatmo, Ph.D., delivering “Transforming the World Through Learning and Control,” followed by Prof. Mustafa Khammash, who discussed “What Makes a Good Theory?” Prof. Yul

It also aimed to inspire the next generation of control engineering leaders, particularly those from eastern and western regions of Indonesia who are geographically distant from the capital.

Yunazwin Nazaruddin then presented “Integrating AI into Instrumentation, Control and Automation.” After the lunch break, Prof. Deddy Kurniadi gave a talk on “Multipath Ultrasonic Flowmeter in Industry,” and Tua Agustinus Tamba, Ph.D., shared his work on “End-to-End Design and Validation of Safety-Preserving Control Methods on Low-Cost Robotics Platforms.” The program concluded with an industry session by Albert Ryan of Schneider Electric Indonesia, who presented “New Generation of

Industrial Automation Technology,” offering participants insights into current industrial challenges and emerging directions in automation. The day closed with a networking session that encouraged further exchange among participants.

More than 50 participants attended the Bandung workshop, representing more than 20 Indonesian institutions, including universities, research centers, government agencies, and industry. Participants were invited to join ICA plenary sessions, featuring talks by Karl



FIGURE 3 Participants engaged enthusiastically in Dr. Ayu Garetta’s workshop on soft robotic gloves.



FIGURE 4 Bandung Workshop opening address by Prof. Karl H. Johansson (IEEE CSS Vice President DODA)



FIGURE 5 Active engagement of a woman-in-control participant during the workshop.



FIGURE 6 Group photo of the participants and speakers at the Bandung workshop.



FIGURE 7 Plenary speakers at ICA 2025: Mustafa Khammash, Rubita Sudirman, Karl H. Johansson, and Augie Widyotriatmo.



FIGURE 8 ICA 2025 banquet, cultural showcase, and city sightseeing.

H. Johansson on “Observer Synthesis for Nonlinear Systems via Physics-Informed Learning,” Mustafa Khammash on “Genetic Control Systems for Cell Therapy,” Rubita Sudirman from Universiti Teknologi Malaysia on “Bio-signal Monitoring Devices for Cognitive Function Evaluation,” and Dr. Augie Widyotriatmo on “Robust Autonomous Systems: Sensing, Controls, State Estimation, Digital Twins, Cyber-Resilience, and Complex Interactions.” ICA 2025 itself brought together more than 100 participants from across Indonesia and abroad. The workshop and conference conclude with a banquet and cultural showcase, followed by a city tour the next day (Figures 5–8).

Together, the two workshop events strengthened awareness of systems and control engineering across diverse regions of Indonesia. The series provided young researchers with early exposure to global developments, meaningful interactions with leading scientists, and an understanding of how the field connects to national industrial needs. The workshops also underscored the importance of building stronger research communities and increasing national participation in international control forums.

ACKNOWLEDGMENT

The organizers gratefully acknowledge the support of the IEEE Control Systems Society Diversity, Outreach, and Development Activities Program, the Instrumentation, Control, and Automation Research Group at Institut Teknologi Bandung, the Asian Control Association, and the Organizing Committee of International Conference on Instrumentation, Control, and Automation 2025.

Augie Widyotriatmo 
and Azka Muji Burohman 




In this issue of *IEEE Control Systems*, we speak with Mouhacine Benosman, principal research scientist at Amazon Robotics in Boston, USA, and Martin Guay, a professor at Queen's University in Kingston, Canada.

Before joining Amazon Robotics, Mouhacine Benosman worked at Reims University, France; Strathclyde University, Scotland; the National University of Singapore; and Mitsubishi Electric Research Labs in Cambridge, USA. He received a Ph.D. in applied mathematics in 2002 from Ecole Centrale de Nantes, France, and an MA in psychology from Harvard University, USA, in 2025. His research interests include multiagent distributed control with applications to robotics; physics-informed machine learning for modeling, estimation, and control of infinite dimensional systems with applications to fluid dynamics; data-driven robust reinforcement learning control; and applying psychology theory to artificial intelligence systems design. He served as an associate editor for *IEEE Control*

Systems Letters and for *Journal of Optimization Theory and Applications* and as senior editor of the *International Journal of Adaptive Control and Signal Processing*.

Martin Guay is a full professor at Queen's University, where he received the Queen's University Chancellor Research Award and the Premier Research Excellence Award. He also received the Syncrude Innovation Award and the D. G. Fisher Award from the Canadian Society of Chemical Engineers. In 2011, he received, with Dr. Veronica Adetola, the Best Paper Award (Theory) from *Journal of Process Control* (2008-1011). He is a Fellow of the Chemical Institute of Canada. He is the editor-in-chief of *Journal of Process Control* and senior editor for *IEEE Transactions on Automatic Control*. He is an associate editor for *Automatica* and *The Canadian Chemical Engineering Journal*. His research interests are in control theory and process control.

Anuradha Annaswamy 

Digital Object Identifier 10.1109/MCS.2026.3659289

Date of current version: 26 March 2026

MOUHACINE BENOSMAN

Q. How did your education and early career lead to your initial and continuing interest in the control field?

Mouhacine: My early academic background was in applied mathematics, which gave me a strong foundation in analytical thinking. However, I soon realized that I wanted to apply mathematics to real-world systems rather than study it purely abstractly. Control theory struck me as the natural bridge between mathematical rigor and engineering practice. It not only satisfied my interest in dynamical systems theory but also opened pathways to-



Mouhacine Benosman.

ward multidisciplinary research in areas such as mechatronics, robotics, and autonomous systems.

Q. What are some of your research interests?

Mouhacine: My early research focused on nonlinear dynamical systems and flexible multibody modeling, primarily motivated by applications in space robotics. This naturally led me to explore robust and adaptive nonlinear control, particularly for fault-tolerant systems where resilience is critical. Over time, my curiosity expanded toward cooperative and distributed control of multiagent systems.

After more than a decade immersed in model-based "classical" control theory, I became increasingly interested in how learning methods could enhance adaptability and performance. This shift inspired my

Digital Object Identifier 10.1109/MCS.2026.3659290

Date of current version: 26 March 2026

We are entering a transformative era defined by large-scale computing and the rapid evolution of artificial intelligence.

current line of research on hybrid control frameworks that integrate model-based and data-driven approaches. In recent years, my work has centered on robustness and adaptation in reinforcement learning control, as well as physics-informed and operator-learning methods for infinite-dimensional systems, with applications to robotics and fluid dynamical systems.

Most recently, my research journey took an interdisciplinary turn: I returned to graduate school to study psychology, where I started investigating psychometric and cognitive frameworks for evaluating artificial intelligence systems. I aim to better understand and design safe human-aligned artificial intelligence (AI) systems in autonomous decision-making environments.

Q. What are some of the most promising opportunities you see in the control field?

Mouhacine: We are entering a transformative era defined by large-scale computing and the rapid evolution of artificial intelligence. I believe control theory will play a pivotal role in shaping this revolution, particularly in ensuring that AI-based systems are not only efficient but also provably safe and reliable.

One especially promising direction lies in the emerging connection between operator learning and dynamical systems theory. This bridge provides a natural framework for control theorists to apply their deep understanding of stability, robustness, and safety to the design of learning-based and AI-driven systems. In many ways, the next frontier for control sci-

ence is to embed these classical guarantees into the very fabric of modern machine learning.

Q. You are the author of a textbook in the control field. What topics does your book cover?

Mouhacine: Working in industrial research doesn't leave much room for writing textbooks, so completing one sometimes feels like a small miracle. My book focuses on the emerging field of hybrid model-based and data-driven learning control. It examines how classical model-based control can provide the essential guarantees of safety and robustness, while data-driven and extremum-seeking methods bring adaptability and learning. The central idea is that these two paradigms, model-based and data-driven, are not competing



An encounter with a bison in Arizona.



Mouhacine Benosman next to Amazon Robotics' mobile robots display.



At the top of Mount Major in New Hampshire.



Mouhacine Benosman next to a package delivery drone at the Amazon Machine Learning Conference.

philosophies but complementary ones, capable of producing controllers that are both safe and agile.

Profile of Mouhacine Benosman

- *Current position:* Principal research scientist at Amazon Robotics (AR).
- *Contact information:* Amazon Robotics, North Reading, MA 01864 USA. E-mail: m_benosman@ieee.org, <https://www.linkedin.com/in/mbenos>.
- *IEEE Control Systems Society experience highlights:* Associate editor, *IEEE Control Systems Letters*, 2019–2023.

Q. What are some of your interests and activities outside of your professional career?

Mouhacine: I've always loved being on or in the water. When time allows, I volunteer as a certified swimming coach for kids and row with a crew team in the Boston area. Beyond the water, I like hiking and don't mind getting *moderately* lost in the forest; it's

a good way to reset and think away from equations and code.

Q. Thank you for your comments.

Mouhacine: It's a real privilege to represent control practitioners from industry in this prestigious feature. My sincere thanks to the editors for the invitation and for highlighting the diverse voices shaping the control community.

MARTIN GUAY

Q. What are some of your research interests?

Martin: My current research sits at the intersection of nonlinear control, optimization, and data-driven methods, with a particular emphasis on control strategies that can operate reliably when accurate models are unavailable or costly to obtain. A central theme of my work is extremum-seeking and real-time optimization, motivated by industrial process control problems where performance objectives must be optimized online in the presence of disturbances, uncertainty, and limited measurements. I am especially interested in model-free and learning-based approaches that retain strong theoretical guarantees while remaining practical for real-world implementation.

I am highly interested in design methods for nonlinear systems. A primary focus has been nonlinear model predictive control (MPC) design techniques and their adaptive and robust extensions. MPC remains a very powerful technique that facilitates performance focused designs for general nonlinear systems subject to constraints. I am also interested in differential geometric control

methods. My work in this area has favored the application of exterior differential systems to feedback linearization, observer error linearization, and homotopy operator decomposition methods.

Applications provide an important source of inspiration for my work. Process systems, energy-related applications, and large-scale interconnected systems have all shaped my



Martin Guay with Dr. Telema Harry and Jeremy Henderson, Stratotegic Founder, at the Canadian Space Agency.

I am particularly motivated by problems where control theory can directly impact efficiency, sustainability, and operational reliability.

perspective on control research. I am particularly motivated by problems where control theory can directly impact efficiency, sustainability, and operational reliability.

Ultimately, my goal is to contribute control methodologies that are mathematically sound, intuitively grounded, and deployable in practice, while mentoring students to appreciate both the elegance of theory and the discipline required for real-world control design.

Q. What courses do you teach relating to control? Do you have a favorite course? How would you describe your teaching style?

Martin: Teaching has always been more than just transferring knowledge for me; it's about inviting students into the world of *real-time decision making and control intuition*. At Queen's University, I've had the pleasure of teaching core chemical engineering and control-oriented courses that sit at the heart of process dynamics and optimization. Whether it's introducing undergraduates to the fundamentals of feedback and process control or men-

toring graduate seminars on nonlinear methods and extremum-seeking strategies, my classes are structured to challenge students to think dynamically, to question models, and to connect theory to practice. I've always tried to bring content alive with real industrial examples and hands-on reasoning, the same way I approach research, and it was truly an honor to be recognized by students with the Golden Apple Award for Teaching for my work in the classroom at Queen's Engineering.

If I had to choose a favorite, it would have to be the graduate-level control courses focusing on nonlinear systems analysis, differential geometric control, and adaptive control. These classes allow students to tackle the technical aspects of nonlinear control while understanding the impact of model uncertainty and explore model-free and adaptive control methods.

My teaching philosophy leans toward an interactive, question-driven style. I encourage students to actively participate in the class. Their engage-

ment motivates a curiosity-driven approach to frame problems, and to connect analytical tools with physical intuition. The interplay between theory and student discovery remains one of the most rewarding aspects of academic life.

Q. What are some of the most promising opportunities you see in the control field?

Martin: One of the most promising opportunities in control lies in the growing convergence between data-driven methods and classical feedback theory. Learning, optimization, and control are no longer separate activities; they are increasingly intertwined in systems that must adapt and improve performance in real time. The challenge is to integrate these ideas while preserving the robustness, stability, and interpretability that have always been central to control.

Another exciting direction is real-time optimization and decision making in large-scale, networked systems. Energy systems, process networks, and autonomous platforms all demand control strategies that operate across multiple time and spatial scales, often with limited models and incomplete information. Control theory is uniquely positioned to provide principled solutions to these problems.

I believe the control community has an important role to play in shaping sustainable and resilient technologies.



Martin Guay at the suspended temple in Taiyuan.



Martin Guay (and Bonnie) with Denis Dochain (and Martine) and Jaime Moreno in Vienna.

Profile of Martin Guay

- *Current position:* Professor, Queen's University.
- *Contact information:* Queen's University, Kingston, ON K7L 3N6, Canada. E-mail: guaym@queensu.ca, <https://smithengineering.queensu.ca/directory/faculty/martin-guay.html>.
- *IEEE Control Systems Society experience highlights:* Chair, Technical Committee on Industrial Process Control, 2011–2016; associate editor, *IEEE Transactions on Control Systems Technology*, 2010–2016; senior editor, *IEEE Control Systems Letters*, 2018–2024; associate editor, *IEEE Transactions on Automatic Control*, 2022–2024; senior editor, *IEEE Transactions on Automatic Control*, 2025–present.



Martin Guay at the cottage.

From improving efficiency in process industries to enabling autonomous and energy-aware systems, control theory offers tools that can have lasting societal impact. The most exciting opportunities, in my view, lie at the intersection of strong theoretical foundations and meaningful real-world applications.

I also see tremendous opportunity in model-free and optimization-based control methods, such as extremum seeking, where feedback is used not just to regulate optimal behavior. These approaches resonate strongly with industrial needs, where uncertainty and changing operating conditions are the norm rather than the exception.

An especially important opportunity lies in the integration of deep learning with control theory. Deep learning has demonstrated remarkable ability to extract structure from high-dimensional data, but without feedback and dynamical constraints, learning alone is often brittle when deployed in real systems. Control theory provides the language of stability, robustness, and safety that is essential for turning learned models and policies into dependable closed-loop systems. Conversely, learning offers control new tools for representation, adaptation, and scalability in complex environments.

There is significant opportunity in the integration of differential geometric control with modern deep learning and computational methods for nonlin-

ear systems. Differential geometry has given us powerful insights into structure, invariance, and feedback design for nonlinear dynamics, but its practical deployment has often been limited by modeling complexity and computational burden. Advances in learning and numerical methods now make it possible to revisit these ideas in a more constructive way, using data-driven representations and computational tools to identify structure, approximate manifolds, and synthesize feedback laws in systems where analytic models are incomplete or unavailable. This synthesis offers a promising path toward nonlinear control designs that are both theoretically principled and computationally scalable.

Q. You are the author of two books in the control field. What topics do these books cover?

Martin: The first book, titled *Robust and Adaptive Model Predictive Control of Nonlinear Systems*, coauthored with collaborators D. DeHaan and V. Adetola, focuses on how MPC methods can be extended and adapted to handle nonlinear dynamics with uncertainties, including robust and adaptive approaches that update parameters online and address constraint handling and estimation issues in real applications. It covers theory, design methodology, and computational aspects relevant to adaptive robust nonlinear MPC. This text aligns closely with my longstanding research interests in nonlinear control, adaptive estimation, and

real-time optimization and is part of a broader literature connecting model-based control with practical deployment in uncertain environments.

The second book, titled *Robust Viability of Hybrid Systems*, coauthored by Gina Labinaz, offers a systematic treatment of hybrid systems, focusing on the interaction between continuous-time dynamics and discrete-event behavior that arises naturally in many engineered systems. The book emphasizes hybrid automata and state-based representations, which allow continuous dynamics, discrete modes, and logic-based switching to be treated in a coherent fashion. The book also explores control design for hybrid systems, including supervisory control, mode-dependent feedback laws, and the prevention of undesirable behaviors such as chattering and Zeno phenomena. Attention is paid to how control objectives and constraints propagate across discrete transitions, a topic of critical importance in safety-critical and event-driven applications.

Q. What are some of your interests and activities outside of your professional career?

Martin: Music has always been my main interest outside of work. I play guitar and bass (as often as I can). My influences include fusion jazz and progressive rock. As a proud Canadian, I enjoy playing hockey and fly fishing.

Q. Thank you for your comments.

Martin: Thank you for the opportunity.

Into the Second Century of Extremum Seeking Control

TIAGO ROUX OLIVEIRA^{id}, JORGE I. POVEDA^{id}, and MARTIN GUAY^{id}



GRAPHIC CREATED FROM AUTHOR-GENERATED ARTWORK USING OPENAI CHATGPT. FOR IMAGE CREDITS, PLEASE SEE THE "ACKNOWLEDGMENT" SECTION OF THIS ARTICLE.

The early promise of extremum seeking control (ESC) was alarmingly simple: perturb, observe, and adjust. Yet, behind this simplicity lay a subtle question that would occupy the community for decades: *Why does it work?* The emergence of rigorous stability and convergence proofs around the turn of

the millennium transformed ESC from an ingenious trick into a reliable method. With that transformation came confidence, and with confidence came adoption. Today, ESC is deployed across chip manufacturing, wind and photovoltaic energy systems, mobile robotics, bioengineering, acoustic control, noncooperative games, neuromuscular electrical stimulation, biological reactors, oil drilling, additive manufacturing, and traffic flow optimization for urban mobility, among many others.

Digital Object Identifier 10.1109/MCS.2025.3649967
Date of current version: 26 March 2026

In an era captivated by machine learning, ESC reminds us that learning, optimization, and stability need not be traded off against one another.

What distinguishes ESC from many contemporary data-driven and learning-based approaches is not merely its model-free character but the nature of its guarantees. ESC does not rely on large offline datasets nor does it defer performance assurances until after extensive retraining. Instead, it operates online, in a closed loop, and comes equipped with quantifiable convergence and robustness properties that speak directly to the language of control. In an era captivated by machine learning, ESC reminds us that learning, optimization, and stability need not be traded off against one another.

Recent years have witnessed an acceleration of theoretical and practical developments that have expanded the scope of ESC far beyond its original formulations. Extensions to systems with delays, infinite-dimensional dynamics governed by partial differential equations, stochastic and nonsmooth dynamics, and hybrid architectures combining continuous-time physics with discrete-time decision making have opened new territories for real-time model-free optimization. These advances are not incremental; they represent qualitative shifts that position ESC as a central tool for complex cyberphysical systems.

Against this backdrop, the present special issue has an ambitious but deliberately focused goal: to provide the *IEEE Control Systems* readership with a coherent snapshot of high-momentum directions in extremum seeking research and application. Rather than attempting a comprehensive survey—a task that would be both monumental and quickly outdated—we have curated a set of tutorial-style contributions that emphasize insight, perspective, and connection. Our hope is that readers new to ESC will find clear entry points into the field, while experts will discover resonances across theory and practice.

The first two articles in this special issue focus on theory. The first [A1], by Tiago Roux Oliveira, Miroslav Krstić, and Tamer Başar, surveys extremum and Nash equilibrium seeking for systems with delays and partial differential equations, tracing the evolution of ESC from finite-dimensional settings to networks of distributed infinite-dimensional agents. The second [A2], by Jorge I. Poveda and Andrew R. Teel, addresses extremum seeking in hybrid dynamical systems, where continuous-time dynamics intertwine with discrete events, logic, and decision making.

Complementing these theoretical contributions are two application-oriented articles. One [A3], by Mouhacine Benosman, illustrates how ESC has been translated into robust industrial solutions, from electromagnetic brakes to fluid dynamic systems, demonstrating the

method's ability to survive—and thrive—in demanding operational environments. The other [A4], by Alexander Scheinker, explores the emerging synergy between extremum seeking and deep generative learning, showing how ESC can enhance the robustness of machine learning-based controllers for time-varying systems, with particle accelerators providing a compelling case study.

Mirroring the four technical articles, the cover artwork of this special issue follows a Warhol-meets-Kurosawa (Rashomon) rationale: four colorful variations of a single image, symbolizing four distinct perspectives on ESC—different viewpoints on the same underlying idea. In essence, extremum seeking is not about knowing the mountain but about learning how to climb to its peak.

Taken together, the articles in this special issue portray a field that is both mature and restless: mature in its analytical foundations yet restless in its pursuit of new problem classes and technological frontiers. If the first century of extremum seeking was devoted to discovering and justifying its core ideas, the second century promises integration, with hybrid systems, games, infinite-dimensional dynamics, and modern learning architectures.

We hope that this special issue will serve not only as a reference but also as a catalyst. By gathering diverse perspectives in one place, we aim to encourage cross fertilization among communities and inspire new collaborations across disciplines and continents. Extremum seeking began as a practical response to uncertainty. Its future, we believe, lies in embracing complexity without relinquishing the rigor that has brought it this far.

ACKNOWLEDGMENT

The guest editors would like to thank Miroslav Krstic for reading the draft of this editorial and providing valuable feedback.

A century after its invention and a quarter century after the first rigorous convergence proofs, ESC stands as one of the most enduring and quietly transformative ideas in control. Born from the practical need to optimize systems whose models are uncertain, incomplete, or altogether unavailable, ESC has matured from heuristic beginnings into a principled framework for real-time optimization of *dynamic* systems. This special issue of *IEEE Control Systems* marks not only a historical milestone but also a turning point: an invitation to look forward, into the second century of extremum seeking.

The emergence of rigorous stability and convergence proofs around the turn of the millennium transformed ESC from an ingenious trick into a reliable method.

They are also grateful for the support of the current and former editors-in-chief, Anuradha Annaswamy and Rodolphe Sepulchre, as well as associate editors Mihailo Jovanovic and Alexandre Bazanella. The opening artwork was generated using ChatGPT. Prompt by Tiago Roux Oliveira: “Provide a mosaic of four images with a goat searching for the peak of a mountain. The aesthetic must follow Andy Warhol’s art style.”

AUTHOR INFORMATION

Tiago Roux Oliveira (tiagoroux@uerj.br) joined the State University of Rio de Janeiro, Rio de Janeiro-RJ 20550-900, Brazil, as an associate professor in the Department of Electronics and Telecommunication Engineering in 2010. In 2014, he was a visiting scholar with the University of California, San Diego, USA. He has published about 300 refereed journal articles, book chapters, and conference papers. In 2017, he was nominated as an affiliate member of the Brazilian Academy of Sciences. He has served as an associate editor on the editorial boards of *IEEE Transactions on Automatic Control*, *IEEE Open Journal of Control Systems*, *IEEE Control Systems Letters*, *Automatica*, and *Systems & Control Letters*, among others. In 2020, he was elected and nominated chair of Technical Committee Adaptive and Learning Systems of the International Federation of Automatic Control for the triennium 2020–2023 (being reelected to the triennium 2023–2026). In 2021, he received the IEEE Transactions on Control Systems Technology Outstanding Paper Award from the IEEE Control Systems Society. In 2022, he was elected president of the Brazilian Society of Automatics for the biennial 2023–2025. He is a coauthor of the book *Extremum Seeking Through Delays and PDEs*, published by the Society for Industrial and Applied Mathematics in 2022. He is a Senior Member of IEEE.

Jorge I. Poveda received dual B.S. degrees in electronics and mechanical engineering from the University of Los Andes, Colombia, as well as the M.S. degree in electrical engineering from the same institution. He then received M.Sc. and Ph.D. degrees in electrical and computer engineering from the University of California (UC), Santa Barbara. He is an associate professor in the Department of Electrical and Computer Engineering at UC San Diego, La Jolla, CA 92093 USA. After receiving the Ph.D. degree, he was a postdoctoral scholar at Harvard University and an assistant professor at the University of Colorado Boulder. He is the recipient of the Computer and Information Science and Engineering Research Initiation Initiative and CAREER Awards from the National Science Foundation, Young Investigator

Awards from the Air Force Office of Scientific Research and Society of Hispanic Professional Engineers, the American Automatic Control Council Donald P. Eckman Award, the CCSD Outstanding Scholar Fellowship and Best Ph.D. Dissertation Awards from UC Santa Barbara, and the UC San Diego Engineering Early Career Faculty Development Award. He has also received several best paper and best paper finalist awards with his students and colleagues, including recognitions from *IEEE Transactions on Control of Network Systems* (2023), the American Control Conference (2024), the IEEE Conference on Decision and Control (2017 and 2021), and the IFAC Conference on Analysis and Design of Hybrid Systems (2024). His group develops foundational frameworks and algorithms for the design, control, and operation of autonomous and cyberphysical systems, with an emphasis on hybrid dynamics, adaptation, game theory, artificial intelligence, and real-time decision making across multiagent engineering, societal, and biological domains.

Martin Guay is a professor in the Department of Chemical Engineering, Queen’s University, Kingston, ON K7L 4V1, Canada. He received the Queen’s University Chancellor Research Award and the Premier Research Excellence Award. He also received the Syncrude Innovation Award and the D. G. Fisher Award from the Canadian Society of Chemical Engineers. In 2011, he received, with Dr. Veronica Adetola, the Best Paper Award (Theory) from *Journal of Process Control* (2008–2011). He is the editor-in-chief of *Journal of Process Control* and a senior editor of *IEEE Control Systems Letters*. He is an associate editor of *Automatica* and *IEEE Transactions on Automatic Control*. He is also a review editor for *Canadian Chemical Engineering Journal*. His research interests include process control, control theory, and applied statistics. He is a Senior Member of IEEE and a fellow of the Chemical Institute of Canada.

APPENDIX: RELATED ARTICLES

- [A1] T. R. Oliveira, M. Krstić, and T. Başar, “Extremum and Nash equilibrium seeking with delays and PDEs: Designs & applications: Into the second century of extremum-seeking control,” *IEEE Control Syst.*, vol. 46, no. 2, pp. 39–87, Apr. 2026, doi: [10.1109/MCS.2025.3647235](https://doi.org/10.1109/MCS.2025.3647235).
- [A2] J. I. Poveda and A. R. Teel, “Hybrid set-seeking systems: Model-free feedback optimization via hybrid inclusions,” *IEEE Control Syst.*, vol. 46, no. 2, pp. 88–129, Apr. 2026, doi: [10.1109/MCS.2025.3646104](https://doi.org/10.1109/MCS.2025.3646104).
- [A3] M. Benosman, “Examples of extremum seeking control in industry: Auto-tuning applications from electromagnetic brakes to fluid dynamics,” *IEEE Control Syst.*, vol. 46, no. 2, pp. 130–149, Apr. 2026, doi: [10.1109/MCS.2025.3646338](https://doi.org/10.1109/MCS.2025.3646338).
- [A4] A. Scheinker, “Adaptive control meets deep learning: Extremum seeking in latent space with particle accelerator applications,” *IEEE Control Syst.*, vol. 46, no. 2, pp. 150–176, Apr. 2026, doi: [10.1109/MCS.2025.3646850](https://doi.org/10.1109/MCS.2025.3646850).



Extremum and Nash Equilibrium Seeking With Delays and PDEs: Designs & Applications

INTO THE SECOND CENTURY OF
EXTREMUM-SEEKING CONTROL



TIAGO ROUX OLIVEIRA^{ID},
MIROSLAV KRSTIĆ^{ID},
and TAMER BAŞAR^{ID}

Classical extremum seeking (ES) deals with model-free online optimization problems in which, regardless of whether there is a single input or multiple inputs, only a single payoff is maximized (or cost minimized). Optimization problems become more interesting when there are multiple payoffs being maximized by the respective multiple inputs/actors/players. When the different payoffs are distinct functions of the inputs of different players, it is clear that conflicting objectives may arise among the players. This scenario leads to what is called a

Digital Object Identifier 10.1109/MCS.2025.3647235
Date of current version: 26 March 2026

Strong motivation exists for designs that ensure convergence to Nash equilibria in the presence of delays and PDEs.

noncooperative game. A game is noncooperative if players cannot form alliances or if all agreements need to be self-enforcing (for example, through credible threats), focusing on predicting individual players' actions and payoffs and analyzing *Nash equilibria* [50]. A Nash equilibrium is an outcome that, once achieved or reached, means no player can increase its payoff by changing decisions unilaterally [9].

The realization that ES is applicable not only to single-player optimization (possibly with a vector of inputs maximizing a single payoff), in an online model-free setting, but also to multiple competing players in the noncooperative game setting emerged in [21]. Each player employing an ES algorithm maximizes its own payoff, irrespective of the other players' actions. It is shown in [21] that, if all players employ ES algorithms, they collectively converge to a Nash equilibrium. In other words, each of the players finds an optimal strategy in an online fashion, even though they do not know the analytical forms of the payoff functions (neither the other players' nor their own)

and neither have access to the actions applied by the other players nor to the payoffs achieved by the other players. Recent versions of the Nash equilibrium-seeking (NES) algorithm originally proposed in [21] can be found in [66], where fixed-time convergence is incorporated into time-varying networks.

Delays and partial differential equations (PDEs) [40], [41] in a game context arise in applications like network virtualization, software-defined networks, cloud computing, the Internet of Things, context-aware networks, green communications, and security [4], [5], [26]. In particular, differential games with delays are dealt with (in a partly or fully model-based fashion) in [12], [13], [17], [24], [34], [49], and [64]. PDE dynamics arise in the Black–Scholes model of behavior in financial markets [45], [68]. Hence, strong motivation exists for designs that ensure convergence to Nash equilibria in the presence of delays and PDEs.

The development of ES has progressed from static maps to finite-dimensional dynamic systems to networks of static and dynamic players (see “[Summary](#)”). Extensions from ordinary differential equation (ODE) dynamics to maps and players that incorporate delays, or even PDEs, were the next natural step in that progression, driven by increasing research challenges [58]. In [59] and in [19] and [53], we dealt with classical ES with single payoff functions, where the actuation dynamics is governed by hyperbolic and parabolic PDEs, respectively. In [60], we have considered noncooperative games in which players act through arbitrarily long delays or transport hyperbolic PDEs. The delays may be distinct, and, in general, each player knows only the length of its own delay. To compensate for distinct delays at the inputs of the players, we have employed predictor feedback.

In the first half of this article, we consider noncooperative games where players act through a more complex and complete scenario of PDE dynamics.

First, following a natural transition from the NES design through delays to design through parabolic PDEs, we progress to developing a nonmodel-based strategy for locally stable convergence to Nash equilibria in quadratic noncooperative games with player actions subject to diffusion (heat) PDE dynamics with distinct diffusion coefficients and each player having access only to its own payoff value. The proposed approach employs ES with sinusoidal perturbation signals employed to estimate the gradient (first derivative) and Hessian (second derivative) of unknown quadratic functions. To compensate for distinct diffusion processes in the

Summary

The development of ES has progressed over the past 100 years, from static maps to finite-dimensional dynamic systems to networks of static and dynamic agents. Extensions from ordinary differential equation dynamics to maps and agents that incorporate delays or even PDEs are the next natural step in that progression through ascending research challenges. This article provides an overview of results on the algorithm design and theory of ES for such infinite-dimensional systems. Both hyperbolic and parabolic dynamics are presented: delays or transport equations, heat-dominated equations, wave equations, and reaction–advection–diffusion equations. NES methods are introduced for noncooperative game scenarios of the model-free kind and then specialized to single-agent optimization. Even heterogeneous PDE games, such as a duopoly with one parabolic and one hyperbolic agent, are considered. Several engineering applications are touched upon for illustration, including flow-traffic control for urban mobility, oil-drilling systems, deep-sea cable-actuated source seeking, additive manufacturing modeled by the Stefan PDE, biological reactors, light-source seeking with flexible-beam structures, and neuromuscular electrical stimulation.

The applications featured in the article are selected not only for their physical variety but also for their mathematical typicality.

inputs of the N players, we employ boundary control with averaging-based estimates.

Second, when players are competing in a noncooperative game, there is no reason to assume that players are pursuing even remotely related objectives. One could be maximizing its profit, while another could maximize some social good. Likewise, there is no reason to assume that different players are subject to the same type of physics. One might be propagating an input through some social opinion dynamics, while another may be propagating it through some epidemiological dynamics. This means that, while first we provide useful designs and guarantees of Nash equilibrium attainment in the presence of hyperbolic and parabolic PDE dynamics (examples of which are delays or heat dynamics), it is of interest to see when different players who—unaware of the competing optimization pursuits in different physical domains, are employing ES algorithms with compensation of their own specific PDEs—will be able to achieve Nash optimality under the interference of other players who are compensating their own PDE dynamics, which are from different classes.

In this sense, one could formulate the problem of a noncooperative game with dynamics from different PDE classes (hyperbolic, parabolic, Korteweg–De Vries, Schrödinger, Kuramoto–Sivashinsky, etc.), but we do not pursue all such possibilities here. We only demonstrate the achievement of convergence to a Nash equilibrium in a game with two “heterogeneous” players (a duopoly), where one player compensates a transport PDE and the other one a heat PDE. While not coupled directly, the coupling of the players in the payoff functions and the coupling that results from the ES algorithms acting at the boundary conditions of the PDEs result in a coupled heterogeneous pair of PDEs in the closed-loop system, which also happens to be time varying and nonlinear.

In both scenarios (homogeneous or heterogeneous games), we apply an adequate small-gain analysis for the resulting input-to-state stable (ISS) parabolic PDE–ODE loop as well as averaging theory in infinite dimensions, because of the infinite-dimensional state of the heat PDEs and delays, to obtain local convergence to a small neighborhood of the Nash equilibrium. We quantify the size of these residual sets for this distinct class of parabolic–hyperbolic PDEs.

The material in this article and the related results [61], [62], [63], presented only in conferences but not published in journals, constitute the first instance of noncooperative

games being tackled in a model-free fashion with ES in the presence of PDE dynamics, represented by heat PDEs and delays. The employment of PDE motion planning (for obtaining the probing signals) combined with the PDE–ISS formulation, the distributed boundary control feedback design, and the closed-loop analysis via small-gain theorems for PDEs are the main bits of new knowledge that a reader already acquainted with ES will gain from this article.

We start the second part of the article by highlighting the relation between NES and classical ES for infinite-dimensional systems. In this discussion, classes of PDEs more general than transport and heat equations are addressed.

The second half of the article is dedicated to a number of engineering applications, presenting traffic control, oil-drilling control, deep-sea cable-actuated control, additive manufacturing, bioreactors, light-source seeking with cantilever beams, and neuromuscular electrical stimulation (NMES) under delays. These applications may attract interest on the basis of either their physical nature or their mathematical nature.

Physically, the applications selected to introduce the reader to some of the technological possibilities of ES are grouped as follows. First, we present an application representative of flow: traffic. Next, we show applications entailing mechanical vibrations: drill-string instabilities in oil drilling and cable oscillations in undersea manipulation and construction. Third, we turn to thermal diffusion in additive manufacturing and the diffusion of chemical or biological matter in tubular reactors. Fourth, in a light-seeking application, we consider cantilever beams (vibrating systems qualitatively different from the torsional vibrations in drill strings and translational vibrations in cables). Fifth, we consider NMES in patient rehabilitation.

The applications featured in the article are selected not only for their physical variety but also for their mathematical typicality. The traffic and NMES applications feature a pure transport PDE or delay. The drilling and cable applications feature second-order hyperbolic PDEs. The additive manufacturing and bioreactor applications feature parabolic PDEs. Through these applications, the reader is exposed to the full spectrum of basic PDE types in which ES can be successfully pursued.

Another key motivation for this article is for it to serve as a tutorial-style guide to ES and NES in the context of systems with delays and PDEs—an area that has seen rapid development but remains technically challenging and

fragmented across various subdomains. In this sense, the article is also constructed to

- » integrate and systematize the core theoretical tools and results from recent developments, including our own [58]
- » provide intuitive explanations, design guidelines, and application examples to lower the entry barrier for students and researchers new to this topic
- » clarify the interplay between finite- and infinite-dimensional systems as well as delay effects in ES and NES algorithms
- » motivate future research by outlining open problems and application domains in which these techniques can have a strong impact.

The article is organized as follows. The section “*N-Player Game With Quadratic Payoffs for NES*” introduces the standard assumptions (such as diagonal dominance for the Hessian matrix) for noncooperative games with N players and quadratic payoff functions. The setup taking delays into account is shown in the section “*Noncooperative Scenario With Delays*,” and we consider the inclusion of heat PDEs in noncooperative games in the section “*Noncooperative Games With Heat PDEs*.” In these two sections, we consider delays and heat PDEs separately, each player handling a distinct heat PDE or transport PDE. The extension from homogeneous noncooperative games to heterogeneous ones is carried out in the section “*Noncooperative Duopoly With Heterogeneous Transport-Heat PDE Dynamics*.” In that section, we seek Nash equilibria in a quadratic noncooperative duopoly game ($N = 2$) with player actions concomitantly subject to transport-heat PDE dynamics—one player compensating for a delay (transport PDE) and the other one for a heat (diffusion)

PDE—and each player having access only to its own payoff value. Stability analyses for each specific case are conducted in the corresponding sections “*Heat PDEs Case*,” “*Delay Case*,” and “*Stability Analysis*,” with the proofs of the main theorems made available. We illustrate the theoretical results numerically on an example combining hyperbolic and parabolic dynamics in a two-player setting, given in the section “*Simulations*.” The section “*From Nash Equilibrium to ES*” brings further results on ES for infinite-dimensional systems governed by a wider class of PDE systems. A series of application sections follows, covering traffic control, oil-drilling control, deep-sea cable-actuated control, additive manufacturing, bioreactors, light-source seeking with cantilever beams, and NMES under delays, showing the vast space of possibilities of designing ES algorithms for real-world applications. Concluding remarks are included in the section “*Conclusions*.”

An outline of the article’s extension of ES results is provided in “*Extensions and Synthesis*.” The notation and terminology used in the article are listed in “*Notation, Norms, and Terminology*.”

N-PLAYER GAME WITH QUADRATIC PAYOFFS FOR NES

Game theory provides an important framework for the mathematical modeling and analysis of scenarios involving different players where there is coupling in their actions, in the sense that their respective outcomes (outputs) $y_i(t) \in \mathbb{R}$ do not depend exclusively on their own actions/strategies (input signals) $\theta_i(t) \in \mathbb{R}$, with $i = 1, \dots, N$, but at least on a subset of others’. Moreover, defining $\theta := [\theta_1, \dots, \theta_N]^T$, each player’s payoff function $J_i(\theta): \mathbb{R}^N \rightarrow \mathbb{R}$ depends on the

Extensions and Synthesis

We provide a consolidated and generalized treatment of both ES and NES in the presence of delays and heterogeneous PDE dynamics. We introduce a unified controller design framework that accommodates both transport and diffusion as well as expanded settings to other classes of PDEs.

The article also introduces new convergence proofs tailored to the integrated framework and general convergence results for (N)ES under mixed-type PDE dynamics, which were not presented in earlier articles.

Beyond such a more complete and modular treatment of delay and PDE compensation in both ES and NES interactions, application domains and numerical simulations are developed, highlighting the broad applicability of the proposed methods.

In summary, the article extends ES results in several key directions, as follows:

- *Multiplayer generalization*: This work develops a generalized framework for *N-player games*, possibly accommo-

dating heterogeneous dynamics and delay profiles. This requires extending the controller structure and convergence analysis beyond what was previously established in the existing literature.

- *Heterogeneous infinite-dimensional dynamics*: The article deals with *mixed-type PDE dynamics* (for example, combinations of transport and diffusion equations) across different agents.
- *Unified design and analysis*: We provide a *systematic controller design methodology*, applicable to a wide range of PDE-governed agents, and a *unified convergence analysis* that covers both homogeneous and heterogeneous systems. This allows the reader to understand the underlying structure behind various NES schemes and apply them in broader contexts.
- *Tutorial exposition*: We integrate previously scattered results, enhance accessibility, and motivate future developments.

action θ_j of at least one other Player j , $j \neq i$. An N -tuple of actions, θ^* , is said to be in Nash equilibrium if no Player i can improve its payoff by unilaterally deviating from θ_i^* , this being so for all i [9].

Consider games where the payoff function of each player is quadratic, expressed as a strictly concave¹ combination of their actions

$$J_i(\theta(t)) = \frac{1}{2} \sum_{j=1}^N \sum_{k=1}^N \epsilon_{jk}^i H_{jk}^i \theta_j(t) \theta_k(t) + \sum_{j=1}^N h_j^i \theta_j(t) + c_i \quad (1)$$

where $\theta_j(t) \in \mathbb{R}$ is the decision variable (action) of Player j , H_{jk}^i , h_j^i , and $c_i \in \mathbb{R}$ are constants, $H_{ii}^i < 0$, $H_{jk}^i = H_{kj}^i$, and $\epsilon_{jk}^i = \epsilon_{kj}^i > 0$, $\forall i, j, k$.

¹By strict concavity, we mean that $J_i(\theta)$ is strictly concave in θ_i for all θ_{-i} , this being so for each $i = 1, \dots, N$.

Quadratic payoff functions of the previous type are of particular interest in game theory, first because they constitute second-order approximations to other types of nonquadratic payoff functions, and second because they are analytically tractable, leading in general to closed-form equilibrium solutions that provide insight into the properties and features of the equilibrium solution concept under consideration [9].

For the sake of completeness, we provide here, in mathematical terms, the definition of a Nash equilibrium $\theta^* = [\theta_1^*, \dots, \theta_N^*]^T$ in an N -player game:

$$J_i(\theta_i^*, \theta_{-i}^*) \geq J_i(\theta_i, \theta_{-i}^*), \quad \forall \theta_i \in \mathcal{U}_i, \quad i \in \{1, \dots, N\} \quad (2)$$

where J_i is the payoff function of Player i , and the variable θ_i corresponds to its action, while \mathcal{U}_i is its action set, and θ_{-i} denotes the actions of the other players. Hence, no player has an incentive to unilaterally deviate its action from θ^* . In a

Notation, Norms, and Terminology

The 2-norm of the state vector $X(t)$ for a finite-dimensional system described by an ODE is denoted by single bars $|X(t)|$. By contrast, norms of functions (of x) are denoted by double bars. By default, $\|\cdot\|$ denotes the spatial $L_2[0, D]$ -norm, that is, $\|\cdot\| = \|\cdot\|_{L_2[0, D]}$, where we drop the index $L_2[0, D]$ if not otherwise specified. Since the state variable $u(x, t)$ of the infinite-dimensional system governed by a PDE is a function of two arguments, we should emphasize that taking a norm in one of the variables makes the norm a function of the other variable, as adopted in [40] and [41]. For example, the $L_2[0, D]$ -norm of $u(x, t)$ in $x \in [0, D]$ is $\|u(t)\| = (\int_0^D u^2(x, t) dx)^{1/2}$, while the $L_\infty[0, D]$ -norm is defined by $\|u(t)\|_{L_\infty[0, D]} = \|u(t)\|_\infty = \sup_{x \in [0, D]} |u(x, t)|$. Moreover, the \mathcal{H}_1 -norm is given by $\|u(t)\|_{\mathcal{H}_1} = \|u(t)\|_{L_2}^2 + \|u_x(t)\|_{L_2}^2$.

We denote the partial derivatives of a function $u(x, t)$ as $\partial_x u(x, t) = \partial u(x, t) / \partial x$, $\partial_t u(x, t) = \partial u(x, t) / \partial t$. We conveniently use the compact form $u_x(x, t)$ and $u_t(x, t)$ for the former and the latter, respectively.

Consider a generic nonlinear system described by $\dot{x} = f(t, x, \epsilon)$, where $x \in \mathbb{R}^n$, and $f(t, x, \epsilon)$ is periodic in t with period T ; that is, $f(t + T, x, \epsilon) = f(t, x, \epsilon)$. Thus, for $\epsilon > 0$ sufficiently small, we can obtain its average model given by $\dot{x}_{av} = f_{av}(x)$, with $f_{av}(x) = 1/T \int_0^T f(\tau, x, 0) d\tau$, where $x_{av}(t)$ denotes the average version of the state $x(t)$ [S1].

As defined in [S1], a vector function $f(t, \epsilon) \in \mathbb{R}^n$ is said to be of order $O(\epsilon)$ over an interval $[t_1, t_2]$ if $\exists k, \bar{\epsilon} : |f(t, \epsilon)| \leq k\epsilon$, $\forall \epsilon \in [0, \bar{\epsilon}]$ and $\forall t \in [t_1, t_2]$. In most cases, we do not provide explicit values for the constants k and $\bar{\epsilon}$, in which case $O(\epsilon)$ can be interpreted as an order-of-magnitude relation for sufficiently small ϵ .

The term “s” stands for either the Laplace variable or the differential operator d/dt , depending on the context. For a transfer function $H_0(s)$ with a generic input u , the pure convolution $h_0(t) * u(t)$, with $h_0(t)$ being the impulse response of $H_0(s)$,

is also denoted by $H_0(s)u$, as done in [S2]. The maximum and minimum eigenvalues of a square matrix A are denoted by $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$, respectively.

The definitions of the ISS for ODE-based as well as PDE-based systems are assumed to be as provided in [S3] and [32], respectively.

Let $A \subseteq \mathbb{R}^n$ be an open set. By $C^0(A; \Omega)$, we denote the class of continuous functions on A , which take values in $\Omega \subseteq \mathbb{R}^m$. By $C^k(A; \Omega)$, where $k \geq 1$, is an integer, we denote the class of functions on $A \subseteq \mathbb{R}^n$ with continuous derivatives of order k , which take values in $\Omega \subseteq \mathbb{R}^m$. In addition, $C([a, b]; \mathbb{R}^n)$ is the Banach space of continuous functions mapping the interval $[a, b]$ into \mathbb{R}^n ; see [S4, Chapter 2]. Alternatively, $C^n(\mathcal{X})$ denotes an n -times continuously differentiable function on the domain \mathcal{X} . In addition, \mathbb{R}_+ stands for the domain of positive real numbers, including 0.

In line with [S5] and [S4], we assume the usual definitions for any delayed-system $\dot{x}(t) = f(t, x_t)$, $t \geq t_0$, and $x(t_0 + \Theta) = \xi(\Theta)$, $\Theta \in [-D_{\max}, 0]$, where t_0 is an arbitrary initial time instant $t_0 \geq 0$, $x(t) \in \mathbb{R}^n$ is the state vector, $D_{\max} > 0$ is the maximum time delay allowed, the history function of the delayed state is given by $x_t(\Theta) = x(t + \Theta) \in C([-D_{\max}, 0]; \mathbb{R}^n)$, and the functional initial condition ξ is also assumed to be continuous on $[-D_{\max}, 0]$. Without loss of generality, we take $t_0 = 0$ throughout the article.

REFERENCES

- [S1] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, NJ, USA: Prentice-Hall, 2002.
- [S2] P. A. Ioannou and J. Sun, *Robust Adaptive Control*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1996.
- [S3] E. D. Sontag and Y. Wang, “On characterizations of the input-to-state stability property,” *Syst. Control Lett.*, vol. 24, no. 5, pp. 351–359, 1995, doi: [10.1016/0167-6911\(94\)00050-6](https://doi.org/10.1016/0167-6911(94)00050-6).
- [S4] J. K. Hale and S. M. V. Lunel, *Introduction to Functional Differential Equations*. New York, NY, USA: Springer-Verlag, 1993.
- [S5] E. Fridman, *Introduction to Time-Delay Systems - Analysis and Control*. Basel, Switzerland: Birkhäuser, 2014.

duopoly example ($N = 2$) to be considered later, $\mathcal{U}_1 = \mathcal{U}_2 = \mathbb{R}$, where \mathbb{R} denotes the set of real numbers.

To determine the Nash equilibrium solution in strictly concave quadratic games with N players, where each action set is the entire real line, one should differentiate J_i with respect to $\theta_i(t)$, $\forall i = 1, \dots, N$, setting the resulting expressions equal to zero and solving the set of equations thus obtained.

This set of equations, which also provides a sufficient condition due to the strict concavity, is

$$\sum_{j=1}^N \epsilon_{ij}^i H_{ij}^i \theta_j^* + h_i^i = 0, \quad i = 1, \dots, N \quad (3)$$

which can be written in compact form as

$$\begin{bmatrix} \epsilon_{11}^1 H_{11}^1 & \epsilon_{12}^1 H_{12}^1 & \dots & \epsilon_{1N}^1 H_{1N}^1 \\ \epsilon_{21}^2 H_{21}^2 & \epsilon_{22}^2 H_{22}^2 & \dots & \epsilon_{2N}^2 H_{2N}^2 \\ \vdots & \vdots & \ddots & \vdots \\ \epsilon_{N1}^N H_{N1}^N & \epsilon_{N2}^N H_{N2}^N & \dots & \epsilon_{NN}^N H_{NN}^N \end{bmatrix} \begin{bmatrix} \theta_1^* \\ \theta_2^* \\ \vdots \\ \theta_N^* \end{bmatrix} = - \begin{bmatrix} h_1^1 \\ h_2^2 \\ \vdots \\ h_N^N \end{bmatrix}. \quad (4)$$

Defining the Hessian matrix H and vectors θ^* and h by

$$H := \begin{bmatrix} \epsilon_{11}^1 H_{11}^1 & \epsilon_{12}^1 H_{12}^1 & \dots & \epsilon_{1N}^1 H_{1N}^1 \\ \epsilon_{21}^2 H_{21}^2 & \epsilon_{22}^2 H_{22}^2 & \dots & \epsilon_{2N}^2 H_{2N}^2 \\ \vdots & \vdots & \ddots & \vdots \\ \epsilon_{N1}^N H_{N1}^N & \epsilon_{N2}^N H_{N2}^N & \dots & \epsilon_{NN}^N H_{NN}^N \end{bmatrix} \quad \theta^* := \begin{bmatrix} \theta_1^* \\ \theta_2^* \\ \vdots \\ \theta_N^* \end{bmatrix}, \quad h := \begin{bmatrix} h_1^1 \\ h_2^2 \\ \vdots \\ h_N^N \end{bmatrix} \quad (5)$$

there exists a unique Nash equilibrium at $\theta^* = -H^{-1}h$ if H is invertible:

$$\begin{bmatrix} \theta_1^* \\ \theta_2^* \\ \vdots \\ \theta_N^* \end{bmatrix} = - \begin{bmatrix} \epsilon_{11}^1 H_{11}^1 & \epsilon_{12}^1 H_{12}^1 & \dots & \epsilon_{1N}^1 H_{1N}^1 \\ \epsilon_{21}^2 H_{21}^2 & \epsilon_{22}^2 H_{22}^2 & \dots & \epsilon_{2N}^2 H_{2N}^2 \\ \vdots & \vdots & \ddots & \vdots \\ \epsilon_{N1}^N H_{N1}^N & \epsilon_{N2}^N H_{N2}^N & \dots & \epsilon_{NN}^N H_{NN}^N \end{bmatrix}^{-1} \begin{bmatrix} h_1^1 \\ h_2^2 \\ \vdots \\ h_N^N \end{bmatrix}. \quad (6)$$

For more details, see [9, Chapter 4].

The Basic Idea of NES in a Two-Player Game

Let players P1 and P2 represent two firms that produce the same good, have dominant control over a market, and compete for profit by setting their prices u_1 and u_2 , respectively. The profit of each firm is the product of the number of units sold and the profit per unit, which is the difference between the sale price and the marginal or manufacturing cost of the product. In mathematical terms, the profits are modeled by

$$J_i(t) = s_i(t)(u_i(t) - m_i) \quad (S1)$$

where s_i is the number of sales, m_i the marginal cost, and $i \in \{1, 2\}$ for P1 and P2. Intuitively, the profit of each firm will be low if it either sets the price very low since the profit per unit sold will be low, or if it sets the price too high since then consumers will buy the other firm's product. The maximum profit is expected to lie somewhere in the middle of the price range, and it crucially depends on the price level set by the other firm. (See Figure S1).

To model the market behavior, we assume a simple, but quite realistic model, where for whatever reason, the consumer prefers the product of P1 but is willing to buy the product of P2 if its price u_2 is sufficiently lower than the price u_1 . Hence, we model the sales for each firm as

$$s_1(t) = S_d - s_2(t), \quad s_2(t) = \frac{1}{p}(u_1(t) - u_2(t)) \quad (S2)$$

where the total consumer demand S_d is held fixed for simplicity, the preference of the consumer for P1 is quantified by $p > 0$, and the inequalities $u_1 > u_2$ and $(u_1 - u_2)/p < S_d$ are assumed to hold.

Substituting (S2) into (S1) yields expressions for the profits $J_1(u_1, u_2)$ and $J_2(u_1, u_2)$ that are both quadratic functions of the prices u_1 and u_2 , namely

$$J_1 = \frac{-u_1^2 + u_1 u_2 + (m_1 + S_d p)u_1 - m_1 u_2 - S_d p m_1}{p} \quad (S3)$$

$$J_2 = \frac{-u_2^2 + u_1 u_2 + m_2 u_1 + m_2 u_2}{p} \quad (S4)$$

and thus, the Nash equilibrium is easily determined to be

$$u_1^* = \frac{1}{3}(2m_1 + m_2 + 2S_d p) \quad (S5)$$

$$u_2^* = \frac{1}{3}(m_1 + 2m_2 + S_d p). \quad (S6)$$

To make sure the constraints $u_1 > u_2$, $(u_1 - u_2)/p < S_d$ are satisfied by the Nash equilibrium, we assume that $m_1 - m_2$ lies in the interval $(-S_d p, 2S_d p)$. If $m_1 = m_2$, this condition is automatically satisfied.

For completeness, we provide here the definition of a Nash equilibrium $u^* = [u_1^*, u_2^*]^T$ in a two-player game:

$$J_i(u_i^*, u_{-i}^*) \geq J_i(u_i, u_{-i}^*), \quad \forall u_i \in U_i, i \in \{1, 2\} \quad (S7)$$

where J_i is the payoff function of Player i , u_i its action, U_i its action set, and u_{-i} denotes the action of the other player. Hence, no player has an incentive to unilaterally deviate its action from u^* . In the duopoly example, $U_1 = U_2 = \mathbb{R}_+$, where \mathbb{R}_+ denotes the set of positive real numbers.

(Continued)

The Basic Idea of NES in a Two-Player Game (Continued)

To attain the Nash strategies (S5) and (S6) without any knowledge of modeling information, such as the consumer's preference p , the total demand S_d , or the other firm's marginal cost or price, the firms implement a nonmodel-based real-time optimization strategy, for example, deterministic ES with sinusoidal perturbations, to set their price levels. Specifically, P1 and P2 set their prices, u_1 and u_2 , respectively, according to the time-varying strategy (Figure S1):

$$\dot{\hat{u}}_i(t) = k_i \mu_i(t) J_i(t) \quad (\text{S8})$$

$$u_i(t) = \hat{u}_i(t) + \mu_i(t) \quad (\text{S9})$$

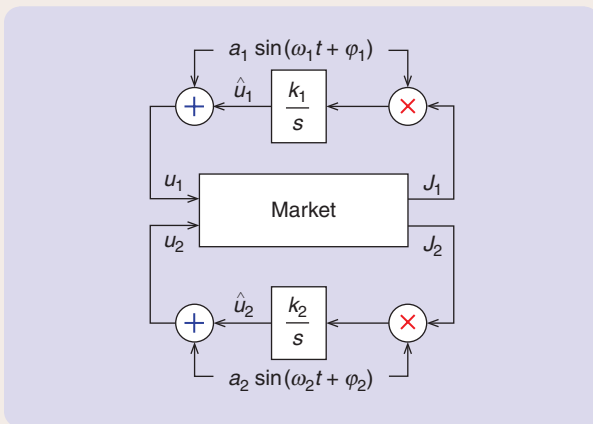


FIGURE S1 Deterministic Nash-seeking schemes applied by players in a duopoly market structure.

where $\mu_i(t) = a_i \sin(\omega_i t + \varphi_i)$, $k_i, a_i, \omega_i > 0$ and $i \in \{1, 2\}$. Further, the frequencies are of the form

$$\omega_i = \omega \bar{\omega}_i \quad (\text{S10})$$

where ω is a positive real number and $\bar{\omega}_i$ is a positive rational number.

By contrast, the firms are also guaranteed to converge to the Nash equilibrium when employing the standard parallel action update scheme [9, Proposition 4.1]

$$u_1^{(k+1)} = \frac{1}{2}(u_2^{(k)} + m_1 + S_d p) \quad (\text{S11})$$

$$u_2^{(k+1)} = \frac{1}{2}(u_1^{(k)} + m_2) \quad (\text{S12})$$

which requires each firm to know both its own marginal cost and the other firm's price at the previous step of the iteration, and also requires P1 to know the total demand S_d and the consumer preference parameter p . In essence, P1 must know nearly all of the relevant modeling information. When using the ES algorithm (S8) and (S9), the firms only need to measure the value of their own payoff functions, J_1 and J_2 . The convergence of (S11) and (S12) is global, whereas the convergence of the Nash-seeking strategy for this example can be proved to be semiglobal, following [S6], or local, by applying the theory of averaging [S1].

REFERENCE

[S6] Y. Tan, D. Nesić, and I. Mareels, "On non-local stability properties of extremum seeking control," *Automatica*, vol. 42, no. 6, pp. 889–903, 2006, doi: 10.1016/j.automatica.2006.01.014.

In addition to the following Assumption 1 (formulated in [21]), we further assume/formalize the subsequent Assumption 2 for noncooperative games.

Assumption 1

The Hessian matrix H given by (5) is strictly diagonal dominant, that is,

$$\sum_{j \neq i}^N |\epsilon_{ij}^i H_{ij}^i| < |\epsilon_{ii}^i H_{ii}^i|, \quad i \in \{1, \dots, N\}. \quad (7)$$

Assumption 2

The parameters ϵ_{jk}^i and ϵ_{kj}^i that appear in the Hessian matrix H given by (5) satisfy the following conditions:

$$\epsilon_{ii}^i = 1, \quad \epsilon_{jk}^i = \epsilon_{kj}^i = \epsilon, \quad \forall j \neq k \quad (8)$$

with $0 < \epsilon < 1$ in the payoff functions (1).

By Assumption 1, the Nash equilibrium θ^* exists and is unique since strictly diagonally dominant matrices are

nonsingular by the Levy–Desplanques theorem [29]. Assumption 2 could be relaxed to consider different values of the coupling parameters ϵ_i for each Player i . However, without loss of generality, we have assumed the same weights for the interconnection channels among the

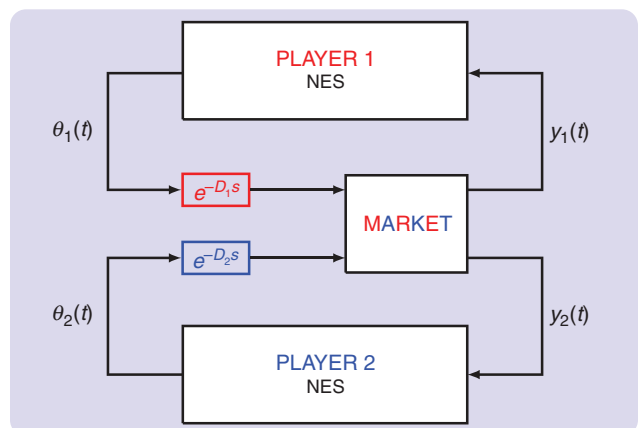


FIGURE 1 NES schemes applied by two players ($N = 2$) in a duopoly market structure with delayed players' actions.

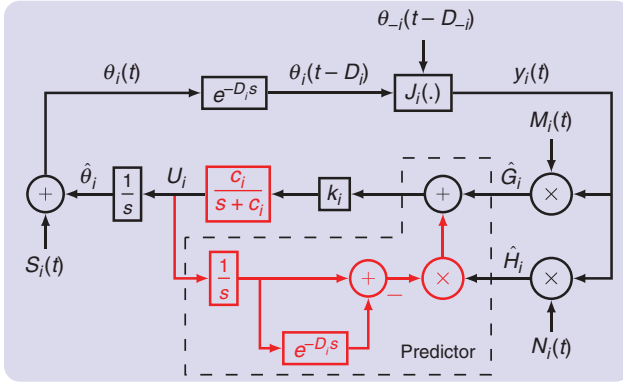


FIGURE 2 Block diagram illustrating the NES strategy performed for each player. In red is shown the predictor feedback used to compensate the individual delay D_i for the noncooperative case. With some abuse of notation, the constants c_i were chosen to denote the parameters of the filters $c_i/(s + c_i)$, but they are not necessarily the same constants that appear in the payoff functions given in (1).

players to facilitate the proofs of our theorems and also to guarantee that the considered noncooperative game is not favoring any specific player. See also “The Basic Idea of NES in a Two-Player Game.” In what follows, we introduce

NES in noncooperative games with model-free ES in the presence of PDE dynamics.

NONCOOPERATIVE SCENARIO WITH DELAYS

For the sake of completeness and to keep the material self-contained, we briefly review the case of noncooperative games subject to multiple and distinct delays originally addressed in [60]. In this scenario, the purpose of the ES is still to estimate the Nash equilibrium vector θ^* but without allowing any sharing of information among the players. As mentioned earlier, each player only needs to measure the value of its own payoff function described by

$$y_i(t) = J_i(\theta(t - D)) = \frac{1}{2} \sum_{j=1}^N \sum_{k=1}^N \epsilon_{jk}^i H_{jk}^i \theta_j(t - D_j) \theta_k(t - D_k) + \sum_{j=1}^N h_j^i \theta_j(t - D_j) + c_i \quad (9)$$

with J_i given by (1). In this sense, we are able to formulate the closed-loop system in a *decentralized* fashion, where no knowledge about the payoffs y_{-i} or actions θ_{-i} of the other players is required, as illustrated in Figure 1.

Motion Planning for PDEs

The field of boundary control of PDEs is almost entirely devoted to stabilization problems for PDEs, dealing with the design of *feedback* laws. However, since the PDEs we address in this article are linear, the same feedback laws designed for stabilization of zero equilibria are capable of also stabilizing any other feasible trajectories of these PDEs. To accomplish such tasks, we need to be able to first generate such trajectories as functions of t and x . In other words, we need to be able to generate not only the necessary feedback controls but also the *feedforward* controls.

In “Trajectory Generation Examples (Part I)” and “Trajectory Generation Examples (Part II),” we provide an introduction to some recently emerging ideas in *trajectory generation* [42, Chapter 12], *motion planning* [44], or, simply put, open-loop control for PDEs.

As we shall see, once we are able to generate a *reference* trajectory for the state of the PDE, we can easily combine our feedforward and feedback controls to stabilize this trajectory. For example, if the control input to the PDE is of Dirichlet type $u(1, t)$ and we have designed a feedback control

$$u(1, t) = \int_0^1 k(1, y) u(y, t) dy \quad (S13)$$

and if we have generated a state reference trajectory $u^r(1, t)$, then the overall “feedback + feedforward” control law

$$u(1, t) = u^r(1, t) + \int_0^1 k(1, y) (u(y, t) - u^r(y, t)) dy \quad (S14)$$

will (exponentially) stabilize the trajectory; namely, it will ensure that $u(x, t) - u^r(x, t) \rightarrow 0$ as $t \rightarrow \infty$ (for all x , in an appropriate sense).

While trajectory tracking for a *state* reference is a relevant problem, a truly important engineering problem is that of tracking an *output* reference. To solve such a problem, one starts from an output reference trajectory, for example, $u^r(0, t)$, as the desired temporal waveform of the system output $u(0, t)$ at $x = 0$; then generates the state trajectory $u^r(x, t)$ for all x (including $x = 1$, which produces the *control reference*); and then combines this result with a feedback control law to stabilize the trajectory $u^r(x, t)$ and to force the output $u(0, t)$ to track the output trajectory $u^r(0, t)$.

“Trajectory Generation Examples (Part I)” and “Trajectory Generation Examples (Part II)” focus on the generation of the state trajectory $u^r(x, t)$ given the output trajectory $u^r(0, t)$, which is often referred to as motion planning. Our coverage is driven not by a desire to achieve generality but by a desire to achieve clarity. Instead of trying to find a state trajectory for an arbitrary output trajectory, we go through a series of examples with reference outputs common in practice—polynomial and sinusoidal signals (the latter being crucial for the generation of the additive dither signals $S_i(t)$ in NES and ES control problems)—for a series of different PDEs, including the heat equation, the wave equation (including Kelvin–Voigt damping), and a first-order hyperbolic PDE. We also consider different types of outputs, including the “Dirichlet” outputs $u(0, t)$ and “Neumann” outputs $u_x(0, t)$.

Trajectory Generation Examples (Part I)

ES is an effective real-time optimization method for PDE systems in cascade with nonlinear quadratic maps. To address PDEs in the feedback loop, a boundary control law and a redesign of the additive probing signal are mandatory. The latter, commonly called “trajectory generation” or “motion planning,” involves designing perturbation signals that anticipate their propagation through PDEs. Specifically, this requires solving motion-planning problems for systems governed by parabolic and hyperbolic PDEs. In this first part, we address the former cases.

EXAMPLE S-A

As our first example, we consider the plant

$$u_t = u_{xx} \quad (\text{S15})$$

$$u_x(0) = 0 \quad (\text{S16})$$

with reference output

$$u^r(0, t) = 1 + 2t - t^2. \quad (\text{S17})$$

In a physical sense, we want to generate a temperature trajectory at $x = 1$ such that the temperature evolution at $x = 0$ is given by (S17). To find the reference input $u^r(1, t)$, we first need to construct the full-state trajectory $u^r(x, t)$, which simultaneously satisfies (S15), (S16), and (S17). Let us search for the state trajectory in the following form:

$$u^r(x, t) = \sum_{k=0}^{\infty} a_k(t) \frac{x^k}{k!}. \quad (\text{S18})$$

This is a Taylor series in x with time-varying coefficients $a_k(t)$ that need to be determined from (S15)–(S17). From (S17) and (S18), we see that

$$a_0(t) = u^r(0, t) = 1 + 2t - t^2. \quad (\text{S19})$$

The boundary condition (S16) gives

$$a_1(t) = u_x^r(0, t) = 0. \quad (\text{S20})$$

The next step is to substitute (S18) into (S15) as follows:

$$\begin{aligned} \sum_{k=0}^{\infty} \dot{a}_k(t) \frac{x^k}{k!} &= \frac{\partial^2}{\partial x^2} \sum_{k=0}^{\infty} a_k(t) \frac{x^k}{k!} \\ &= \sum_{k=2}^{\infty} a_k(t) \frac{k(k-1)x^{k-2}}{k!} \\ &= \sum_{k=2}^{\infty} a_k(t) \frac{x^{k-2}}{(k-2)!} \\ &= \sum_{k=0}^{\infty} a_{k+2}(t) \frac{x^k}{k!}. \end{aligned} \quad (\text{S21})$$

We get the recursive relationship

$$a_{k+2}(t) = \dot{a}_k(t). \quad (\text{S22})$$

Using (S20) and (S19) with (S22) results in

$$\begin{aligned} a_0 &= 1 + 2t - t^2, & a_1 &= 0 \\ a_2 &= 2 - 2t, & a_3 &= 0 \\ a_4 &= -2, & a_5 &= 0 \\ a_6 &= 0, & a_i &= 0, \quad \text{for } i > 6. \end{aligned}$$

This gives the reference state trajectory

$$u^r(x, t) = 1 + 2t + t^2 + (1-t)x^2 - \frac{1}{12}x^4$$

and the input signal

$$u^r(1, t) = \frac{23}{12} + t - t^2. \quad \diamond$$

Note that the output matches the reference output trajectory only if the initial condition of the plant is satisfied by the state trajectory, that is, if $u(x, 0) = 1 + x^2 - \frac{1}{12}x^4$. To asymptotically track the reference signal as $t \rightarrow \infty$ for an arbitrary initial condition, we need to apply feedback.

The basic idea introduced in Example S-A is to use the “spatial Taylor series” representation (S18) and to find the temporal coefficients $a_k(t)$ using a recursion such as (S22). This idea will permeate our developments in the subsequent examples. However, while in Example S-A the Taylor series was finite, in the subsequent examples, the series will be infinite and will have to be summed.

The subsequent examples, Examples S-B through S-E, will deal with sinusoidal reference functions. The following formulas are useful when calculating the trajectories for sinusoidal reference outputs:

$$\begin{aligned} \cosh(a) &= \sum_{k=0}^{\infty} \frac{a^{2k}}{(2k)!}, & \sinh(a) &= \sum_{k=0}^{\infty} \frac{a^{2k+1}}{(2k+1)!} \\ \cosh(ja) &= \cos(a), & \sinh(ja) &= j \sin(a) \\ \cos(ja) &= \cosh(a), & \sin(ja) &= j \sinh(a). \end{aligned}$$

EXAMPLE S-B

Consider the plant

$$u_t = u_{xx} \quad (\text{S23})$$

$$u_x(0) = 0 \quad (\text{S24})$$

with reference output

$$u^r(0, t) = \sin(\omega t). \quad (\text{S25})$$

Since $\sin(\omega t) = \text{Im}\{e^{j\omega t}\}$, we can get the reference trajectory by setting $\lambda = 0$ and $\alpha = j\omega$ in [47, Example 12.2]. We have

$$u^r(x, t) = \text{Im}\{\cosh(\sqrt{j\omega} x) e^{j\omega t}\}. \quad (\text{S26})$$

Using the identity $\sqrt{j} = (1+j)/2$, we get

$$\begin{aligned} u^r(x, t) &= \text{Im}\{\cosh\left((1+j)\sqrt{\frac{\omega}{2}} x\right) e^{j\omega t}\} \\ &= \text{Im}\left\{\frac{e^{\sqrt{\frac{\omega}{2}} x + j(\omega t + \sqrt{\frac{\omega}{2}} x)} + e^{-\sqrt{\frac{\omega}{2}} x + j(\omega t - \sqrt{\frac{\omega}{2}} x)}}{2}\right\} \end{aligned}$$

or

$$\begin{aligned} u^r(x, t) &= \frac{1}{2} e^{\sqrt{\frac{\omega}{2}} x} \sin\left(\omega t + \sqrt{\frac{\omega}{2}} x\right) \\ &\quad + \frac{1}{2} e^{-\sqrt{\frac{\omega}{2}} x} \sin\left(\omega t - \sqrt{\frac{\omega}{2}} x\right). \end{aligned}$$

Finally, the reference input is

$$u^r(1, t) = \frac{1}{2} e^{\sqrt{\frac{\omega}{2}}} \sin\left(\omega t + \sqrt{\frac{\omega}{2}}\right) + \frac{1}{2} e^{-\sqrt{\frac{\omega}{2}}} \sin\left(\omega t - \sqrt{\frac{\omega}{2}}\right). \quad (\text{S27})$$

◇

Without loss of generality, we assume that the inputs have distinct known (constant) delays that are ordered so that

$$D = \text{diag}\{D_1, D_2, \dots, D_N\}, \quad 0 \leq D_1 \leq \dots \leq D_N. \quad (10)$$

Further, given any \mathbb{R}^N -valued signal f , we introduce

$$f^D(t) := f(t - D) = [f_1(t - D_1) \quad f_2(t - D_2) \quad \dots \quad f_N(t - D_N)]^T. \quad (11)$$

Figure 2 presents a schematic diagram that summarizes the proposed Nash equilibrium policy for each i th player, where its output is given by (2), where the vector $\theta_{-i}(t - D_{-i})$ in Figure 2 represents the delayed actions of all other players.

The additive–multiplicative dithers $S_i(t)$, $M_i(t)$ are

$$S_i(t) = a_i \sin(\omega_i(t + D_i)) \quad (12)$$

$$M_i(t) = \frac{2}{a_i} \sin(\omega_i t) \quad (13)$$

with nonzero constant amplitudes $a_i > 0$ at frequencies $\omega_i \neq \omega_j$. The additive dither signals $S_i(t)$ that anticipate propagation through delays and PDEs are obtained by solving PDE motion-planning problems, as discussed in “Motion Planning for PDEs,” “Trajectory Generation Examples (Part I),” and “Trajectory Generation Examples (Part II).” Such probing frequencies ω_i can be selected as

$$\omega_i = \omega'_i \omega = \mathcal{O}(\omega), \quad i \in 1, 2, \dots, N \quad (14)$$

where ω is a positive constant and ω'_i is a rational number. One possible choice is given in [23] as

$$\omega_i \notin \left\{ \omega'_j, \frac{1}{2}(\omega'_j + \omega'_k), \omega'_j + 2\omega'_k, \omega'_j + \omega'_k \pm \omega'_l \right\} \quad (15)$$

for all distinct i, j, k , and l .

Considering $\hat{\theta}_i(t)$ as an estimate of θ_i^* , one can define the estimation error:

$$\tilde{\theta}_i(t) = \hat{\theta}_i(t) - \theta_i^*. \quad (16)$$

The estimate \hat{G}_i of the unknown gradient of each payoff J_i is given by

$$\hat{G}_i(t) = M_i(t) y_i(t). \quad (17)$$

Computing the average of the resulting signal leads us to

$$\hat{G}_i^{\text{av}}(t) = \sum_{j=1}^N \epsilon_{ij}^j H_{ij}^i \tilde{\theta}_j^{\text{av}}(t - D_j) \quad (18)$$

with Π defined as

$$\Pi := 2\pi \times \text{LCM}\left\{ \frac{1}{\omega_i} \right\} \quad (19)$$

and LCM standing for the least common multiple.

At this point, if we ignore the prediction loop and the low-pass filter (both indicated in red in Figure 2, the control law $U_i(t) = k_i \hat{G}_i(t)$ could be obtained as in the classical

ES approach. In this case, from (16) and (18), we could write the average version of

$$\dot{\hat{\theta}}_i(t) = U_i(t) \quad (20)$$

as

$$\begin{aligned} \dot{\hat{\theta}}_i^{\text{av}}(t) &= k_i \hat{G}_i^{\text{av}}(t) \\ &= k_i \sum_{j=1}^N \epsilon_{ij}^j H_{ij}^i \tilde{\theta}_j^{\text{av}}(t - D_j). \end{aligned} \quad (21)$$

Therefore, by defining

$$\tilde{\theta}^{\text{av}}(t) := [\tilde{\theta}_1^{\text{av}}(t), \tilde{\theta}_2^{\text{av}}(t), \dots, \tilde{\theta}_N^{\text{av}}(t)]^T \in \mathbb{R}^N$$

and taking into account all players, one has

$$\dot{\tilde{\theta}}^{\text{av}}(t) = KH \tilde{\theta}^{\text{av}}(t - D) \quad (22)$$

with $K := \text{diag}\{k_1, \dots, k_N\}$ and H given by (5). Equation (22) means that, even if KH were a Hurwitz matrix, the equilibrium $\tilde{\theta}_e^{\text{av}} = 0$ of the closed-loop average system would not necessarily be stable for arbitrary values of the time delays D_i . This reinforces the demand to employ the prediction feedback $U_i(t) = k_i \hat{G}_i(t + D_i)$ —or even its filtered version—for each player to collectively stabilize the closed-loop system, as illustrated in red in Figure 2.

In such a decentralized scenario, the dither frequencies ω_{-i} , the excitation amplitudes a_{-i} , and consequently the individual control laws $U_{-i}(t)$ are not available to Player i . Recalling that the models of the payoffs (1) and (2) are also assumed to be unknown, it becomes impossible to reconstruct individually or estimate the Hessian matrix H given in (5) by using demodulating signals, such as those in [59].

Following the nonsharing information paradigm, the i th player is only able to estimate the element H_{ii}^i of the H matrix (5) by itself, this being so for all players. Therefore, only the diagonal of H can be properly recovered in the average sense. In this way, the signal $N_i(t)$ is now simply defined as

$$N_i(t) := N_{ii}(t) = \frac{16}{a_i^2} \left(\sin^2(\omega_i t) - \frac{1}{2} \right) \quad (23)$$

according to [59]. Then, the average version of

$$\hat{H}_i(t) = N_i(t) y_i(t) \quad (24)$$

is given by

$$\hat{H}_i^{\text{av}}(t) = [N_i(t) y_i(t)]_{\text{av}} = H_{ii}^i. \quad (25)$$

To compensate for the time delays, the following predictor-based update law was proposed in [60]:

$$\dot{\hat{\theta}}_i(t) = U_i(t) \quad (26)$$

$$\dot{U}_i(t) = -c_i U_i(t) + c_i k_i \left(\hat{G}_i(t) + \hat{H}_i(t) \int_{t-D_i}^t U_i(\tau) d\tau \right) \quad (27)$$

Trajectory Generation Examples (Part II)

All of the examples in “Trajectory Generation Examples (Part I)” have dealt with Dirichlet-type outputs $u(0, t)$. The next example deals with a Neumann-type output, $u_x(0, t)$. In addition, only parabolic PDEs so far have been considered. The remaining examples deal with hyperbolic PDEs.

Example S-C

Consider the wave equation

$$u_{tt} = u_{xx} \quad (S28)$$

$$u(0) = 0 \quad (S29)$$

with the reference output

$$u'_x(0, t) = \sin(\omega t). \quad (S30)$$

Searching for $u^r(x, t)$ in the form of $u^r(x, t) = \sum_{k=0}^{\infty} a_k(t) (x^k/k!)$, we get

$$a_0 = 0, \quad a_1(t) = \sin(\omega t) = \text{Im}\{e^{j\omega t}\}$$

$$a_{i+2} = \ddot{a}_i(t)$$

which gives

$$a_{2k} = 0$$

$$a_{2k+1}(t) = (j\omega)^{2k} \text{Im}\{e^{j\omega t}\}.$$

The state reference becomes

$$u^r(x, t) = \text{Im}\left\{ \frac{e^{j\omega t}}{j\omega} \sum_{k=0}^{\infty} \frac{(j\omega x)^{2k+1}}{(2k+1)!} \right\}$$

$$= \text{Im}\left\{ \frac{e^{j\omega t}}{j\omega} \sinh(j\omega x) \right\}$$

$$= \text{Im}\left\{ \frac{e^{j\omega t}}{\omega} \sin(\omega x) \right\}$$

$$= \frac{1}{\omega} \sin(\omega x) \sin(\omega t)$$

and the reference input is

$$u^r(1, t) = \frac{\sin(\omega)}{\omega} \sin(\omega t). \quad (S31)$$

Note that, for the same desired reference output trajectory, the reference input for the heat equation—Example S-B in “Trajectory Generation Examples (Part I)”—has a much more complicated form. \diamond

Example S-D

Consider the more complex case of a wave equation with Kelvin–Voigt damping:

$$\epsilon u_{tt} = (1 + d\partial_t)u_{xx} \quad (S32)$$

$$u_x(0) = 0 \quad (S33)$$

and with the reference output

$$u^r(0, t) = \sin(\omega t). \quad (S34)$$

Following the procedure in the previous examples, one can obtain the following reference state trajectory:

$$u^r(x, t) = \frac{1}{2} \left[e^{\sqrt{\epsilon} \omega \frac{\sqrt{1+\omega^2 d^2} - 1}{\sqrt{2\sqrt{1+\omega^2 d^2}}}} x \times \sin\left(\omega\left(t + \sqrt{\epsilon} \omega \frac{\sqrt{1+\omega^2 d^2} + 1}{\sqrt{2\sqrt{1+\omega^2 d^2}}} x\right)\right) \right. \\ \left. + e^{-\sqrt{\epsilon} \omega \frac{\sqrt{1+\omega^2 d^2} - 1}{\sqrt{2\sqrt{1+\omega^2 d^2}}}} x \times \sin\left(\omega\left(t - \sqrt{\epsilon} \omega \frac{\sqrt{1+\omega^2 d^2} + 1}{\sqrt{2\sqrt{1+\omega^2 d^2}}} x\right)\right) \right].$$

\diamond

Example S-E

Consider the first-order hyperbolic PDE

$$u_t = u_x + gu(0) \quad (S35)$$

with the reference output

$$u^r(0, t) = \sin(\omega t). \quad (S36)$$

Searching for the reference trajectory in the form of $u^r(x, t) = \sum_{k=0}^{\infty} a_k(t) (x^k/k!)$, we get

$$a_0 = \sin(\omega t) = \text{Im}\{e^{j\omega t}\}$$

$$a_1 = \dot{a}_0 - gu(0) = \text{Im}\{(j\omega - g)e^{j\omega t}\}$$

$$a_{k+1} = \dot{a}_k = \text{Im}\{(j\omega - g)(j\omega)^k e^{j\omega t}\}$$

$$= \text{Im}\left\{ \left(1 - \frac{g}{j\omega}\right) (j\omega)^{k+1} e^{j\omega t} \right\}.$$

The reference trajectory becomes

$$u^r(x, t) = \text{Im}\left\{ e^{j\omega t} + \left(1 - \frac{g}{j\omega}\right) \sum_{k=1}^{\infty} \frac{(j\omega x)^k}{k!} e^{j\omega t} \right\}$$

$$= \text{Im}\left\{ \left[\frac{g}{j\omega} + \left(1 - \frac{g}{j\omega}\right) e^{j\omega x}\right] e^{j\omega t} \right\}$$

$$= -\frac{g}{\omega} [\cos(\omega t) - \cos(\omega(t+x))] + \sin(\omega(t+x))$$

which gives the reference input

$$u^r(1, t) = \frac{g}{\omega} [\cos(\omega(t+1)) - \cos(\omega t)] + \sin(\omega(t+1)). \quad (S37)$$

\diamond

In the particular case of classical transport PDEs (or pure delays), when $g = 0$, the reference input becomes simply a phase-shifted sinusoidal function, that is, $u^r(1, t) = \sin(\omega(t+1))$.

for positive constants k_i and c_i (see also “Basic Idea of Predictor Feedback Design for ODE Systems With Actuator Delay” and “Backstepping Boundary Control of Transport PDEs”).

NONCOOPERATIVE GAMES WITH HEAT PDES

Recapitulating, the optimality of the respective outcomes (outputs) of Players i and j , respectively, $y_i(t) \in \mathbb{R}$ and

$y_j(t) \in \mathbb{R}$, does not depend exclusively on their own strategies (input signals) $\Theta_i(t) \in \mathbb{R}$ and $\Theta_j(t) \in \mathbb{R}$. Moreover, defining $\Theta := [\Theta_1, \dots, \Theta_N]^T$, each player’s payoff function J_i also depends on Θ_j of the other player j , $j \neq i$. An N -tuple of $\Theta^* = [\Theta_1^*, \dots, \Theta_N^*]^T$ is said to be in Nash equilibrium if no Player i can improve its payoff by unilaterally deviating from Θ_i^* , this being so for all $i \in \{1, \dots, N\}$ [9].

Basic Idea of Predictor Feedback Design for ODE Systems With Actuator Delay

Here we consider the linear infinite-dimensional system

$$\dot{X}(t) = AX(t) + BU(t - D) \quad (\text{S38})$$

where $X \in \mathbb{R}^n$ (A, B) is a controllable pair, and the input signal $U(t)$ is delayed by D units of time.

Given a stabilizing gain vector K for the undelayed system, namely, given a vector K such that the matrix $A + BK$ is Hurwitz, our wish is to have a control that achieves

$$U(t - D) = KX(t). \quad (\text{S39})$$

This control can be alternatively written as

$$U(t) = KX(t + D) \quad (\text{S40})$$

and it appears to be nonimplementable since it requires future values of the state. However, with the variation-of-constants formula, treating the current state $X(t)$ as the initial condition, we have

$$X(t + D) = e^{AD}X(t) + \int_{t-D}^t e^{A(t-\theta)}BU(\theta)d\theta, \quad \forall t \geq 0. \quad (\text{S41})$$

This yields a feedback law

$$U(t) = K \left[e^{AD}X(t) + \int_{t-D}^t e^{A(t-\theta)}BU(\theta)d\theta \right], \quad \forall t \geq 0 \quad (\text{S42})$$

which is implementable, but it is infinite dimensional since it contains the distributed delay term involving past controls,

$\int_{t-D}^t e^{A(t-\theta)}BU(\theta)d\theta$. The closed-loop system is delay compensated:

$$\dot{X}(t) = (A + BK)X(t), \quad t \geq D \quad (\text{S43})$$

but this is true only after the control “kicks in” at $t = D$. During the interval $t \in [0, D]$, the system state is governed by

$$X(t) = e^{At}X(0) + \int_0^t e^{A(t-\tau)}BU(\tau - D)d\tau, \quad \forall t \in [0, D]. \quad (\text{S44})$$

The feedback law (S42) was introduced within the framework of “finite spectrum assignment” [S7], [S8] and the “reduction approach” [7]. In the section “Backstepping Boundary Control of Transport PDEs” we derive the same control law, but in a considerably more complicated way, which will pay dividends by providing us with an explicit Lyapunov–Krasovskii function and the ability to conduct stability analysis in the time domain, as discussed in [44]. (See also “Boundary Control and Backstepping for PDEs” and “Backstepping Boundary Control of Transport PDEs.”)

REFERENCES

- [S7] W. H. Kwon and A. E. Pearson, “Feedback stabilization of linear systems with delayed control,” *IEEE Trans. Autom. Control*, vol. 25, no. 2, pp. 266–269, Apr. 1980, doi: [10.1109/TAC.1980.1102288](https://doi.org/10.1109/TAC.1980.1102288).
 [S8] A. Manitius and A. Olbrot, “Finite spectrum assignment problem for systems with delays,” *IEEE Trans. Autom. Control*, vol. 24, no. 4, pp. 541–552, Aug. 1979, doi: [10.1109/TAC.1979.1102124](https://doi.org/10.1109/TAC.1979.1102124).

As shown in Figures 3 and 4, distinct heat equations (with Dirichlet actuation) are assumed in the vector of player actions $\theta(t) \in \mathbb{R}^2$, in the particular duopoly game with $N = 2$.

In the N -player game, the propagated actuator vector $\Theta(t) \in \mathbb{R}^N$ is given by

$$\Theta_i(t) = \alpha_i(0, t), \quad \forall i \in \{1, \dots, N\} \quad (\text{28})$$

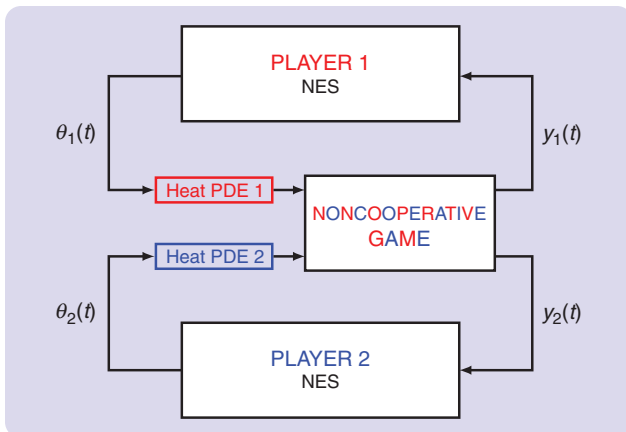


FIGURE 3 NES in a duopoly ($N = 2$) noncooperative game with players acting through heat PDE dynamics.

$$\partial_t \alpha_i(x, t) = \partial_{xx} \alpha_i(x, t), \quad x \in (0, D_i) \quad (\text{29})$$

$$\partial_x \alpha_i(0, t) = 0 \quad (\text{30})$$

$$\alpha_i(D_i, t) = \theta_i(t) \quad (\text{31})$$

where $\alpha_i : [0, D_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}$, and each domain length D_i is known. The solution of (28)–(31) is given by [40, Chapter 15]

$$\alpha_i(x, t) = \mathcal{L}^{-1} \left[\frac{\cosh(x\sqrt{s})}{\cosh(D_i\sqrt{s})} \right] * \theta_i(t) \quad (\text{32})$$

where $\mathcal{L}^{-1}(\cdot)$ denotes the inverse Laplace transformation, and $*$ is the convolution operator.

Given this relation, we define the *diffusion operator* for the PDE (28)–(31) with boundary input and measurement given by $\mathcal{D} = \text{diag}\{\mathcal{D}_1, \dots, \mathcal{D}_N\}$ with

$$\mathcal{D}_i[\varphi(t)] = \mathcal{L}^{-1} \left[\frac{1}{\cosh(D_i\sqrt{s})} \right] * \varphi(t), \quad \text{subject to } \Theta(t) = \mathcal{D}[\theta(t)]. \quad (\text{33})$$

As in (1), we consider games in which the payoff function $y_i(t) = J_i(\mathcal{D}[\theta(t)]) = J_i(\Theta(t))$ of each player is quadratic [9], expressed as a strictly concave combination of their actions propagated through distinct heat PDE dynamics

$$J_i(\Theta(t)) = \frac{1}{2} \sum_{j=1}^N \sum_{k=1}^N \epsilon_{jk}^i H_{jk}^i \Theta_j(t) \Theta_k(t) + \sum_{j=1}^N h_j^i \Theta_j(t) + c_i \quad (34)$$

where $J_i(\Theta): \mathbb{R}^N \rightarrow \mathbb{R}$, obtained by replacing θ for Θ in (1). Equations (2)–(6) can be repeated here simply replacing $\theta^* \in \mathbb{R}^N$ for $\Theta^* \in \mathbb{R}^N$, where Θ^* represents the Nash equilibrium defined in (2).

The *objective* is to design an ES-based strategy to reach the Nash equilibrium in noncooperative games subjected to heat PDEs in the decision variables of the players (input signals).

Since our goal is to find the unknown optimal inputs Θ^* (and θ^*), we define the estimation errors

$$\tilde{\theta}(t) = \hat{\theta}(t) - \theta^*, \quad \vartheta(t) = \hat{\Theta}(t) - \Theta^* \quad (35)$$

where the vectors $\hat{\theta}(t)$ and $\hat{\Theta}(t)$ are estimates of θ^* and Θ^* . To make (35) coherent with the optimizer of the static map Θ^* , we apply the diffusion operator (33) to $\tilde{\theta}_i$ in (35) and get

$$\vartheta_i(t) := \tilde{\alpha}_i(0, t), \quad \forall i \in \{1, \dots, N\} \quad (36)$$

$$\partial_t \tilde{\alpha}_i(x, t) = \partial_{xx} \tilde{\alpha}_i(x, t), \quad x \in (0, D_i) \quad (37)$$

$$\partial_x \tilde{\alpha}_i(0, t) = 0 \quad (38)$$

$$\tilde{\alpha}_i(D_i, t) = \tilde{\theta}_i(t) \quad (39)$$

where $\tilde{\alpha}_i: [0, D_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}$ and $\vartheta(t) := \mathcal{D}[\tilde{\theta}(t)] = \hat{\Theta}(t) - \Theta^*$ is the propagated estimation error $\tilde{\theta}(t)$ through the diffusion domain. For $\lim_{t \rightarrow \infty} \theta(t) = \theta_c$, we have $\lim_{t \rightarrow \infty} \Theta(t) = \Theta_c = \theta_c$, where the index c indicates a constant signal. Indeed, from (32), for a constant input $\theta = \theta_c$, one has $\mathcal{L}\{\theta_c\} = \theta_c/s$, and applying the Laplace limit theorem, we get $\lim_{s \rightarrow 0} \{\theta_c / (\cosh(\sqrt{s} D_i))\} = \theta_c$. In the particular case $\theta = \theta_c = \theta^*$, one has

$$\Theta^* = \theta^*. \quad (40)$$

Figure 4 depicts a schematic diagram that summarizes the proposed Nash equilibrium policy for each player, where their outputs are given by

$$y_i(t) = J_i(\Theta(t)). \quad (41)$$

From the solution of the trajectory generation problem [42, Chapter 12] (see “Trajectory Generation Examples (Part I), Example S-B”), the additive dither signals in the presence of heat PDE dynamics [19] are defined according to

$$S_i(t) = \frac{1}{2} a_i e^{\sqrt{\frac{\omega_i}{2}} D_i} \sin(\omega_i t + \sqrt{\frac{\omega_i}{2}} D_i) + \frac{1}{2} a_i e^{-\sqrt{\frac{\omega_i}{2}} D_i} \sin(\omega_i t - \sqrt{\frac{\omega_i}{2}} D_i) \quad (42)$$

and the multiplicative demodulation signals are given by

$$M_i(t) = \frac{2}{a_i} \sin(\omega_i t) \quad (43)$$

with nonzero constant amplitudes $a_i > 0$ at frequencies $\omega_i \neq \omega_j$. Such probing frequencies ω_i can be selected as

$$\omega_i = \omega'_i \omega = O(\omega), \quad \forall i \in \{1, \dots, N\} \quad (44)$$

where ω is a positive constant and ω'_i is a rational number—one possible choice is given in [23].

Following the nonsharing information paradigm, only the diagonal elements of H can be properly recovered in the average sense by players. In this sense, the signals $N_i(t)$ are simply defined by [23]

$$N_i(t) = \frac{16}{a_i^2} \left(\sin^2(\omega_i t) - \frac{1}{2} \right). \quad (45)$$

Then, the average version of

$$\hat{H}_i(t) = N_i(t) y_i(t) \quad (46)$$

is given by

$$\hat{H}_i^{av}(t) = [N_i(t) y_i(t)]_{av} = H_{ii}^i. \quad (47)$$

Taking into account $\hat{\theta}_i(t)$ as estimates of θ_i^* , one can define from (35) the individual estimation errors

$$\tilde{\theta}_i(t) = \hat{\theta}_i(t) - \theta_i^*, \quad \vartheta_i(t) = \hat{\Theta}_i(t) - \Theta_i^*. \quad (48)$$

The estimates of the unknown gradients of the payoff functions are given by

$$\hat{G}_i(t) = M_i(t) y_i(t) \quad (49)$$

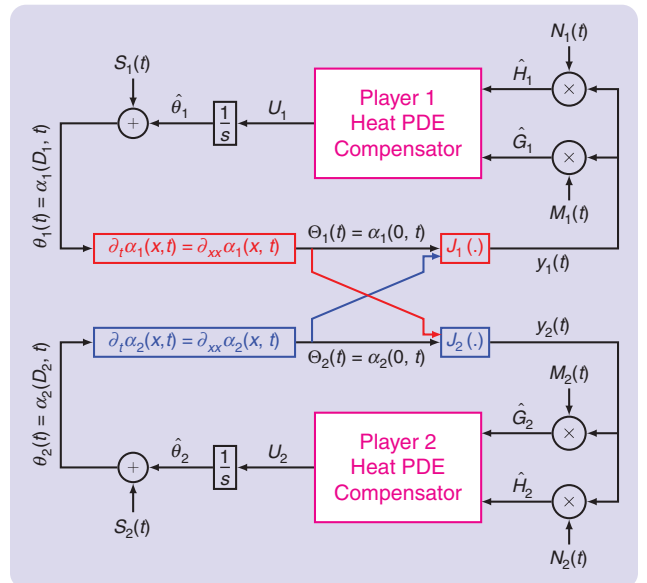


FIGURE 4 Block diagram illustrating the NES strategy performed for each player ($N = 2$). In magenta are shown the boundary controllers used to compensate the individual heat PDEs for the noncooperative game.

Boundary Control and Backstepping for PDEs

Two PDE control settings exist:

- “in domain” control (actuation penetrates inside the domain of the PDE system or is evenly distributed everywhere in the domain, likewise with sensing)
- “boundary” control (actuation and sensing are only through the boundary conditions).

In this sense, boundary control is physically more realistic because actuation and sensing are nonintrusive (think fluid flow where actuation is from the walls). For instance, a “body force” actuation of electromagnetic type is also possible, but it has low control authority, and its spatial distribution typically has a pattern that favors the near-wall region. Moreover, boundary control is a harder problem because the “input operator” (the analog of the B matrix in the linear time-invariant finite-dimensional model $\dot{x} = Ax + Bu$) and the output operator (the analog of the C matrix in $y = Cx$) are unbounded operators.

Many feedback designs for PDEs are inapplicable by boundary control. This article is devoted exclusively to boundary control, using backstepping for PDEs [42], with which destabilizing effects that appear throughout the spatial domain $[0, D]$ can be eliminated. We illustrate the idea of PDE backstepping for the parabolic class. One seeks a change of variables and feedback that transform the PDE system with boundary control (at $x = D$):

$$u_i(x, t) = u_{xx}(x, t) + \int_0^x f(x, y) u(y, t), \quad x \in [0, D] \quad (\text{S45})$$

$$u(D, t) = U(t) \quad (\text{S46})$$

into a target system

$$w_i(x, t) = w_{xx}(x, t), \quad x \in [0, D] \quad (\text{S47})$$

$$w(D, t) = 0 \quad (\text{S48})$$

that has to be exponentially stable. The spatial kernel function $f(x, y)$ in the integral operator in (S45) may contain delta functions like $\delta(x - y)$ and $\delta(y)$. The coordinate transformation

$$w(x, t) = u(x, t) - \int_0^x k(x, y) u(y, t) dy \quad (\text{S49})$$

is proposed in [42] and preceding publications by Krstić and his coauthors, where $k(x, y)$ is the *gain kernel* and the integral in (S49) is a Volterra integral transformation. Since the coordinate transformation (S49) is invertible [42], the stability of the target system translates into stability of the original system. Hence, with the boundary controller chosen as

$$U(t) = u(D, t) = \int_0^D k(D, y) u(y, t) dy \quad (\text{S50})$$

the original system (S45) and (S46) with the controller (S50) is exponentially stable. The gain kernel $k(x, y)$ has to satisfy Goursat PDE conditions that arise from substituting the transformation (S49) into the target system (S47) and (S48) along with the original system equations (S45) and (S46).

Remark S1

There also exists ODE backstepping, especially integrator backstepping [S9], where the feedback law is propagated “backward” through a chain of integrators. The name backstepping for PDEs is caused by the triangular structure of the Volterra integral transformation, similar to ODE backstepping.

Detailed backstepping boundary control methods for transport and heat PDEs are presented in “Backstepping Boundary Control of Transport PDEs” and “Backstepping Boundary Control of Heat PDEs,” respectively.

REFERENCE

[S9] M. Krstić, I. Kanellakopoulos, and P. V. Kokotović, *Nonlinear and Adaptive Control Design*. New York, NY, USA: Wiley, 1995.

and computing the average of the resulting signal leads us to

$$\hat{G}^{\text{av}}(t) = H \hat{\vartheta}^{\text{av}}(t) \quad (50)$$

where the Hessian H is given in (5).

Additionally, from the block diagram in Figure 4, for arbitrary N , one has

$$\hat{\theta}_i(t) = U_i(t), \quad \forall i \in \{1, \dots, N\} \quad (51)$$

and consequently

$$\hat{\theta}_i(t) = U_i(t), \quad \forall i \in \{1, \dots, N\} \quad (52)$$

since $\dot{\theta}(t) = \hat{\theta}(t)$, once θ^* is constant. Taking the time derivative of (36)–(39) and with the help of (35) and (52), the *propagated error dynamics* are written as

$$\dot{\vartheta}_i(t) = u_i(0, t), \quad \forall i \in \{1, \dots, N\} \quad (53)$$

$$\partial_t u_i(x, t) = \partial_{xx} u_i(x, t), \quad x \in (0, D_i) \quad (54)$$

$$\partial_x u_i(0, t) = 0 \quad (55)$$

$$u_i(D_i, t) = U_i(t) \quad (56)$$

where $u_i: [0, D_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}$ and $u_i(x, t) := \partial_t \bar{\alpha}_i(x, t)$.

Hence, from (50) and (53)–(56), it is possible to find a compact form for the overall average estimated gradient according to

$$\dot{\hat{G}}^{\text{av}}(t) = H \hat{\vartheta}^{\text{av}}(t) = H \mathcal{D}[U^{\text{av}}(t)] \quad (57)$$

where $\vartheta^{\text{av}}(t) := [\vartheta_1^{\text{av}}(t), \dots, \vartheta_N^{\text{av}}(t)]^T \in \mathbb{R}^N$, $\hat{G}^{\text{av}}(t) := [\hat{G}_1^{\text{av}}(t), \dots, \hat{G}_N^{\text{av}}(t)]^T \in \mathbb{R}^N$, and $U^{\text{av}}(t) := [U_1^{\text{av}}(t), \dots, U_N^{\text{av}}(t)]^T \in \mathbb{R}^N$ are the average versions of $\vartheta(t) := [\vartheta_1(t), \dots, \vartheta_N(t)]^T$, $\hat{G}(t) := [\hat{G}_1(t), \dots, \hat{G}_N(t)]^T$, and $U(t) := [U_1(t), \dots, U_N(t)]^T$ respectively.

Backstepping Boundary Control of Transport PDEs

The delay in the system (S38) of “Basic Idea of Predictor Feedback Design for ODE Systems With Actuator Delay” can be modeled by the following first-order hyperbolic PDE, also referred to as the “transport PDE”:

$$u_t(x, t) = u_x(x, t) \quad (\text{S51})$$

$$u(D, t) = U(t). \quad (\text{S52})$$

The solution to this equation is

$$u(x, t) = U(t + x - D) \quad (\text{S53})$$

and, therefore, the output

$$u(0, t) = U(t - D) \quad (\text{S54})$$

gives the delayed input. The system (S51) can now be written as

$$\dot{X}(t) = AX(t) + Bu(0, t). \quad (\text{S55})$$

Equations (S51)–(S55) form an ODE–PDE cascade that is driven by the input U from the boundary of the PDE (Figure S2).

Suppose a static state feedback control has been designed for a system with no delay (that is, with $D = 0$) such that

$$U(t) = KX(t) \quad (\text{S56})$$

is a stabilizing controller; that is, the matrix $(A + BK)$ is Hurwitz. Consider the backstepping transformation

$$w(x, t) = u(x, t) - \int_0^x q(x, y)u(y, t)dy - \gamma(x)^T X(t) \quad (\text{S57})$$

with which we want to map the system (S51)–(S55) into the target system

$$\dot{X}(t) = (A + BK)X(t) + Bw(0, t) \quad (\text{S58})$$

$$w_t(x, t) = w_x(x, t) \quad (\text{S59})$$




FIGURE S2 Linear system $\dot{X}(t) = AX(t) + BU(t - D)$ with actuator delay D .

Now, we use an ES strategy based on boundary control to compensate for the diffusion operator $\mathcal{D}[\cdot]$ in (57) due to the multiple and distinct heat PDEs in the players’ actions. Basically, control laws must be able to ensure exponential stabilization of $\hat{G}^{av}(t)$ and

$$w(D, t) = 0 \quad (\text{S60})$$

from the ODE–PDE system

$$\dot{X}(t) = AX(t) + Bu(0, t) \quad (\text{S61})$$

$$u_t(x, t) = u_x(x, t) \quad (\text{S62})$$

$$u(D, t) = U(t). \quad (\text{S63})$$

The transformation (S57) applied to (S58)–(S63) yields

$$q_x(x, y) + q_y(x, y) = 0 \quad (\text{S64})$$

$$q(x, 0) = \gamma(x)^T B \quad (\text{S65})$$

$$\gamma'(x) = A^T \gamma(x) \quad (\text{S66})$$

and

$$\gamma(0) = K^T. \quad (\text{S67})$$

The solution to the ODE (S66) is

$$\gamma(x)^T = Ke^{Ax}. \quad (\text{S68})$$

A general solution to (S64) is

$$q(x, y) = \phi(x - y) \quad (\text{S69})$$

where the function ϕ is determined from (S65). We get

$$q(x, y) = Ke^{A(x-y)}B. \quad (\text{S70})$$

We can now plug the gains $\gamma(x)$ and $q(x, y)$ into the transformation (S57) and set $x = D$ to get the control law

$$u(D, t) = \int_0^D Ke^{A(D-y)}Bu(y, t)dy + Ke^{AD}X(t). \quad (\text{S71})$$

The controller (S71) is expressed in terms of the transport delay state $u(y, t)$. Using (S53), one can also derive the representation in terms of the input signal $U(t)$:

$$U(t) = K \left[e^{AD}X(t) + \int_{t-D}^t e^{A(t-\theta)}BU(\theta)d\theta \right] \quad (\text{S72})$$

which is identical to the controller (S42) in “Basic Idea of Predictor Feedback Design for ODE Systems With Actuator Delay.”

consequently of $\vartheta^{av}(t) = \hat{\Theta}^{av}(t) - \Theta^*$. From (50), it is clear that, if H is invertible, $\vartheta^{av}(t) \rightarrow 0$ as $\hat{G}^{av}(t) \rightarrow 0$. Hence, the convergence of $\vartheta^{av}(t)$ to the origin results in the convergence of $\Theta(t)$ to a small neighborhood of the Nash equilibrium point $\Theta^* = \theta^*$ —see (40)—via averaging theory

[25], without allowing sharing of any information among the players.

Decentralized PDE Boundary Control Using Only the Known Diagonal Terms of the Hessian

In this sense, we can formulate the closed-loop system in a *decentralized* manner, where no knowledge of the payoff or actions of any other player is required.

Inspired by [61], where ES was considered for PDE compensation but not in the context of games, we propose the following boundary-based update laws $\dot{\hat{\theta}}_i(t) = U_i(t)$, $i \in \{1, \dots, N\}$:

$$\begin{aligned} \dot{U}_i(t) = & -c_i U_i(t) \\ & + c_i k_i \left(\hat{G}_i(t) + \hat{H}_i(t) \int_0^{D_i} (D_i - \tau) u_i(\tau, t) d\tau \right) \end{aligned} \quad (58)$$

for positive constants $k_i > 0$ and $c_i > 0$, to compensate for the heat PDEs in (53)–(56). Again, with some abuse of notation, constants c_i were chosen to denote the parameters of the control laws, but they are not related to those that appear in the payoff functions given in (34).

As discussed in [19, Remark 2], the boundary control law (58) could be rewritten as

$$\begin{aligned} \dot{U}_i(t) = & -c_i U_i(t) \\ & + c_i k_i \left[\hat{G}_i(t) + \hat{H}_i(t) (\hat{\theta}_i(t) - \Theta_i(t) + a_i \sin(\omega_i t)) \right] \end{aligned} \quad (59)$$

using the diffusion equations $\partial_t \alpha_i(x, t) = \partial_{xx} \alpha_i(x, t)$, $\forall i \in \{1, \dots, N\}$ and the integration by parts associated with (28)–(31) and (35), and recalling that $\vartheta_i + a_i \sin(\omega_i t) = \Theta_i(t) - \Theta_i^*$, analogously to [19, Equation (25)].

ISS-Like Properties for Parabolic PDE Representation

For the sake of simplicity, let us assume that $c_i \rightarrow +\infty$ in (58), resulting in the following general expression (see “Backstepping Boundary Control of Heat PDEs”):

$$U_i(t) = k_i \left(\hat{G}_i(t) + \hat{H}_i(t) \int_0^{D_i} (D_i - \tau) u_i(\tau, t) d\tau \right). \quad (60)$$

Recalling (53)–(56), the infinite-dimensional closed-loop system (57) and (60) in its average version can be written in the corresponding PDE representation form as

$$\dot{\bar{G}}^{\text{av}}(t) = H u^{\text{av}}(0, t) \quad (61)$$

$$\partial_t u^{\text{av}}(x, t) = D^{-2} \partial_{xx} u^{\text{av}}(x, t), \quad x \in (0, 1) \quad (62)$$

$$\partial_x u^{\text{av}}(0, t) = 0 \quad (63)$$

$$u^{\text{av}}(1, t) = U^{\text{av}}(t) \quad (64)$$

with $D = \text{diag}\{D_1, \dots, D_N\}$.

In the *reduction-like approach* [7] (or finite-spectrum assignment), we use the transformation (for $i \in \{1, \dots, N\}$):

$$\begin{aligned} \bar{G}_i^{\text{av}}(t) &= \hat{G}_i^{\text{av}}(t) + \sum_{j=1}^N \epsilon_{ij}^i H_{ij}^i \int_0^{D_j} (D_j - \tau) u_j^{\text{av}}(\tau, t) d\tau \\ &= \hat{G}_i^{\text{av}}(t) + \sum_{j=1}^N \epsilon_{ij}^i H_{ij}^i \int_0^1 D_j^2 (1 - \xi) u_j^{\text{av}}(\xi, t) d\xi \end{aligned} \quad (65)$$

where $\int_0^{D_j} (D_j - \tau) u_j^{\text{av}}(\tau, t) d\tau = \int_0^1 D_j^2 (1 - \xi) u_j^{\text{av}}(\xi, t) d\xi$. With some mathematical manipulations, it is not difficult to see that \bar{G}^{av} satisfies

$$\dot{\bar{G}}^{\text{av}}(t) = H U^{\text{av}}(t). \quad (66)$$

Now, after adding and subtracting the next terms in blue and red in (60), it is rewritten as

$$\begin{aligned} U_i(t) = & k_i \left(\hat{G}_i(t) + \hat{H}_i(t) \int_0^{D_i} (D_i - \tau) u_i(\tau, t) d\tau \right. \\ & + \sum_{j \neq i} \epsilon_{ij}^i H_{ij}^i \int_0^{D_j} (D_j - \tau) u_j(\tau, t) d\tau \left. \right) \\ & - k_i \sum_{j \neq i} \epsilon_{ij}^i H_{ij}^i \int_0^{D_j} (D_j - \tau) u_j(\tau, t) d\tau \end{aligned} \quad (67)$$

whose average compact form is

$$U^{\text{av}}(t) = K \bar{G}^{\text{av}}(t) + \epsilon K \phi^{\text{av}}(D, t) \quad (68)$$

where the matrix $K := \text{diag}\{k_1, \dots, k_N\}$ with entries $k_i > 0$, and the auxiliary variable $\phi(D, t) := [\phi_1(D, t), \dots, \phi_N(D, t)]^T \in \mathbb{R}^N$ is defined as

$$\begin{aligned} \phi_i(D, t) &:= -\sum_{j \neq i} H_{ij}^i \int_0^{D_j} (D_j - \tau) u_j(\tau) d\tau \\ \phi_i(1, t) &:= -\sum_{j \neq i} H_{ij}^i \int_0^1 D_j^2 (1 - \xi) u_j(\xi, t) d\xi. \end{aligned} \quad (69)$$

Then, it is possible to find a compact form for the overall average game from (66) and (68), such as

$$\dot{\bar{G}}^{\text{av}}(t) = HK \bar{G}^{\text{av}}(t) + \epsilon HK \phi^{\text{av}}(1, t) \quad (70)$$

$$\partial_t u^{\text{av}}(x, t) = D^{-2} \partial_{xx} u^{\text{av}}(x, t), \quad x \in (0, 1) \quad (71)$$

$$\partial_x u^{\text{av}}(0, t) = 0 \quad (72)$$

$$u^{\text{av}}(1, t) = K \bar{G}^{\text{av}}(t) + \epsilon K \phi^{\text{av}}(1, t). \quad (73)$$

From (70), if HK is Hurwitz, it is clear that the dynamics of the ODE state variable $\bar{G}^{\text{av}}(t)$ is exponentially ISS [32] with respect to the PDE state $u(x, t)$ by means of the function $\phi^{\text{av}}(1, t)$. Moreover, the PDE subsystem (71) is ISS [32] with respect to $\bar{G}^{\text{av}}(t)$ in the boundary condition $u^{\text{av}}(1, t)$.

Backstepping Boundary Control of Heat PDEs

Next, we consider the cascade of a heat equation and a linear time-invariant finite-dimensional system given by

$$\dot{X}(t) = AX(t) + Bu(0, t) \quad (S73)$$

$$u_t(x, t) = u_{xx}(x, t) \quad (S74)$$

$$u_x(0, t) = 0 \quad (S75)$$

$$u(D, t) = U(t) \quad (S76)$$

where $X \in \mathbb{R}^p$ is the ODE state, U is the scalar input to the entire system, and $u(x, t)$ is the state of the PDE dynamics of the diffusive actuator. The cascade system is depicted in Figure S3.

The length of the PDE domain D is arbitrary. Thus, we take the diffusion coefficient to be unity without loss of generality. We assume that the pair (A, B) is stabilizable and take K to be a known vector such that $A + BK$ is Hurwitz.

Next, we state a controller that compensates for the diffusive actuator dynamics and guarantees the exponential stability of the resulting closed-loop system (see Figure S4).

We start by formulating an infinite-dimensional transformation of the form

$$w(x, t) = u(x, t) - \int_0^x q(x, y)u(y, t)dy - \gamma(x)X(t) \quad (S77)$$

with kernels $q(x, y)$ and $\gamma(x)$ to be derived, which should transform the plant (S73)–(S76) along with the control law

$$U(t) = K \left[I \ 0 \right] \left\{ e^{\begin{bmatrix} 0 & A \\ I & 0 \end{bmatrix} D} \begin{bmatrix} I \\ 0 \end{bmatrix} X(t) + \int_0^D \left[\int_0^{D-y} e^{\begin{bmatrix} 0 & A \\ I & 0 \end{bmatrix} \xi} d\xi \right] \begin{bmatrix} I \\ 0 \end{bmatrix} Bu(y, t) dy \right\} \quad (S78)$$




FIGURE S3 The cascade of the heat equation PDE dynamics of the actuator with the ODE dynamics of the plant.




FIGURE S4 An arbitrary ODE controlled through a diffusion process.

into the “target system”

$$\dot{X}(t) = (A + BK)X(t) + Bw(0, t) \quad (S79)$$

$$w_t(x, t) = w_{xx}(x, t) \quad (S80)$$

$$w_x(0, t) = 0 \quad (S81)$$

$$w(D, t) = 0. \quad (S82)$$

Combining (S73)–(S82), $\gamma(x)$ and $q(x, y)$ satisfy

$$\gamma''(x) = \gamma(x)A \quad (S83)$$

$$\gamma(0) = K \quad (S84)$$

$$\gamma'(0) = 0 \quad (S85)$$

which happens to represent a second-order ODE in x , and

$$q_{xx}(x, y) = q_{yy}(x, y) \quad (S86)$$

$$q(x, x) = 0 \quad (S87)$$

$$q_y(x, 0) = -\gamma(x)B \quad (S88)$$

which is a second-order hyperbolic PDE of the Goursat type. We then proceed to solve this cascade system explicitly. The explicit solution to the ODE (S83)–(S85) is readily found as

$$\gamma(x) = \begin{bmatrix} K & 0 \end{bmatrix} e^{\begin{bmatrix} 0 & A \\ I & 0 \end{bmatrix} x} \begin{bmatrix} I \\ 0 \end{bmatrix} \quad (S89)$$

and the explicit solution to the PDE (S86)–(S88) is

$$q(x, y) = \int_0^{x-y} \gamma(\sigma)Bd\sigma. \quad (S90)$$

In a manner similar to finding the kernels $q(x, y)$ and $\gamma(x)$ of the direct transformation, the inverse transformation (S77) can be found. To summarize, the direct and inverse backstepping transformations are given by

$$w(x, t) = u(x, t) - \int_0^x m(x-y)u(y, t)dy - KM(x)X(t) \quad (S91)$$

$$u(x, t) = w(x, t) + \int_0^x n(x-y)w(y, t)dy + KN(x)X(t) \quad (S92)$$

where

$$m(s) = \int_0^s KM(\xi)Bd\xi \quad (S93)$$

$$n(s) = \int_0^s KN(\xi)Bd\xi \quad (S94)$$

$$M(\xi) = \begin{bmatrix} I & 0 \end{bmatrix} e^{\begin{bmatrix} 0 & A \\ I & 0 \end{bmatrix} \xi} \begin{bmatrix} I \\ 0 \end{bmatrix} \quad (S95)$$

$$N(\xi) = \begin{bmatrix} I & 0 \end{bmatrix} e^{\begin{bmatrix} 0 & A+BK \\ I & 0 \end{bmatrix} \xi} \begin{bmatrix} I \\ 0 \end{bmatrix}. \quad (S96)$$

STABILITY ANALYSES

We will now show that the ODE–PDE loops [for example, (70)–(73)], for the corresponding delay and heat case, contain small parameters ϵ that can lead to closed-loop stability if

they are chosen sufficiently small. To attain θ^* in real time, without any model information (except for the delays or diffusion domains D_i), each Player i employs the noncooperative ES strategy (27) or (58) via boundary control feedback.

Heat PDEs Case

The next theorem provides the stability/convergence properties of the closed-loop ES feedback for the N -player non-cooperative game with heat PDEs.

Theorem 1

Consider the closed-loop system (61)–(64) under Assumptions 1 and 2, and multiple heat PDEs (28)–(31) with distinct diffusion coefficients D_i for the N -player quadratic non-cooperative game with payoff functions given in (34) and (41) and control laws $U_i(t)$ defined in (58). There exist $c_i > 0$ and $\omega > 0$ sufficiently large as well as $\epsilon > 0$ sufficiently small, such that the closed-loop system with state $\vartheta_i(t)$, $u_i(x, t)$, $\forall i \in \{1, \dots, N\}$, has a unique locally exponentially stable periodic solution in t of period Π in (19), with ω_i in (44) of order $\mathcal{O}(\omega)$ according to (14), denoted by $\vartheta_i^\Pi(t)$, $u_i^\Pi(x, t)$, that satisfies $\forall t \geq 0$:

$$\left(\sum_{i=1}^2 [\vartheta_i^\Pi(t)]^2 + \int_0^{D_i} [u_i^\Pi(x, t)]^2 dx \right)^{1/2} \leq \mathcal{O}\left(\frac{1}{\omega}\right). \quad (74)$$

In particular,

$$\limsup_{t \rightarrow +\infty} |\theta(t) - \theta^*| = \mathcal{O}\left(|a| + \frac{1}{\omega}\right) \quad (75)$$

$$\limsup_{t \rightarrow +\infty} |\theta(t) - \theta^*| = \mathcal{O}\left(|a| e^{\max(D_i)\sqrt{\omega/2}} + \frac{1}{\omega}\right) \quad (76)$$

where $a = [a_1 \ a_2]^T$, and $\theta^* = \Theta^*$ is the unique (unknown) Nash equilibrium given by (6).

Proof

See “Proof of Theorem 1.” ■

Delay Case

On the other hand, the derivative of (18) is

$$\dot{\hat{G}}_i^{\text{av}}(t) = \sum_{j=1}^N \epsilon_{ij}^i H_{ij}^i \dot{\hat{\theta}}_j^{\text{av}}(t - D_j). \quad (77)$$

By delaying by D_i units the time argument of both sides of the average version of (20), we obtain

$$\dot{\hat{\theta}}_i^{\text{av}}(t - D_i) = U_i^{\text{av}}(t - D_i). \quad (78)$$

Thus, (77) can be rewritten as

$$\dot{\hat{G}}_i^{\text{av}}(t) = \sum_{j=1}^N \epsilon_{ij}^i H_{ij}^i U_j^{\text{av}}(t - D_j). \quad (79)$$

Taking into account all players, from (18) and (79), it is possible to find a compact form for the overall average estimated gradient $\hat{G}^{\text{av}}(t) := [\hat{G}_1^{\text{av}}(t), \dots, \hat{G}_N^{\text{av}}(t)]^T \in \mathbb{R}^N$ according to

$$\hat{G}^{\text{av}}(t) = H \hat{\theta}^{\text{av}}(t - D) \quad (80)$$

$$\dot{\hat{G}}^{\text{av}}(t) = H U^{\text{av}}(t - D) \quad (81)$$

where H is given in (5), and

$$U^{\text{av}}(t) := [U_1^{\text{av}}(t), U_2^{\text{av}}(t), \dots, U_N^{\text{av}}(t)]^T \in \mathbb{R}^N.$$

The next theorem provides the stability/convergence properties of the closed-loop ES feedback for the N -player noncooperative game with delays and nonsharing of information.

Theorem 2

Consider the closed-loop system (78) or (79), under Assumptions 1 and 2, and multiple and distinct input delays D_i for the N -player quadratic noncooperative game without information sharing, with payoff functions given in (9) and control laws $U_i(t)$ defined in (27). There exist $c_i > 0$ and $\omega > 0$ sufficiently large as well as $\epsilon > 0$ sufficiently small such that the closed-loop system with state $\hat{\theta}_i(t - D_i)$, $U_i(\tau)$, $\forall \tau \in [t - D_i, t]$, and $\forall i \in \{1, 2, \dots, N\}$ has a unique locally exponentially stable periodic solution in t of period Π in (19), denoted by $\hat{\theta}_i^\Pi(t - D_i)$, $U_i^\Pi(\tau)$, $\forall \tau \in [t - D_i, t]$ satisfying $\forall t \geq 0$:

$$\left(\sum_{i=1}^N [\hat{\theta}_i^\Pi(t - D_i)]^2 + \int_{t-D_i}^t [U_i^\Pi(\tau)]^2 d\tau \right)^{1/2} \leq \mathcal{O}\left(\frac{1}{\omega}\right). \quad (82)$$

In particular,

$$\limsup_{t \rightarrow +\infty} |\theta(t) - \theta^*| = \mathcal{O}\left(|a| + \frac{1}{\omega}\right) \quad (83)$$

where $a = [a_1 \ a_2 \ \dots \ a_N]^T$, and θ^* is the unique Nash equilibrium given by (6).

Proof

See “Proof of Theorem 2.” ■

NONCOOPERATIVE DUOPOLY WITH HETEROGENEOUS TRANSPORT-HEAT PDE DYNAMICS

Here, we propose a nonmodel-based strategy for locally stable convergence to a Nash equilibrium in a quadratic noncooperative game with player actions $\Theta_i(t) \in \mathbb{R}$ subject to heterogeneous PDE dynamics. For simplicity, we consider the duopoly scenario ($N = 2$), where different players use different types of PDEs, one player compensating for a delay (transport PDE) and the other a heat (diffusion) PDE, with each player having access only to its own payoff value, $y_i(t) \in \mathbb{R}$, for $i \in \{1, 2\}$. Heretofore, we solved NES problems with homogeneous games, where the PDE dynamics of distinct nature were not allowed.

Quadratic Payoffs and Heterogeneous PDEs

As shown in Figures 5 and 6, now distinct (transport and heat) PDE equations (with Dirichlet actuation) are assumed in the vector of player actions $\theta(t) \in \mathbb{R}^2$. Thus, the propagated actuator vector $\Theta(t) := [\Theta_1(t), \Theta_2(t)]^T \in \mathbb{R}^2$ is given by the following transport PDE for player P_1 :

$$\Theta_1(t) = \theta_1(t - D_1) = \alpha_1(0, t) \quad (84)$$

$$\partial_t \alpha_1(x, t) = \partial_x \alpha_1(x, t), \quad x \in (0, D_1) \quad (85)$$

$$\partial_x \alpha_1(0, t) = 0 \quad (86)$$

$$\alpha_1(D_1, t) = \theta_1(t) \quad (87)$$

and the next heat PDE for player P_2

$$\Theta_2(t) = \alpha_2(0, t) \quad (88)$$

$$\partial_t \alpha_2(x, t) = \partial_{xx} \alpha_2(x, t), \quad x \in (0, D_2) \quad (89)$$

$$\partial_x \alpha_2(0, t) = 0 \quad (90)$$

$$\alpha_2(D_2, t) = \theta_2(t) \quad (91)$$

Proof of Theorem 1

The proof of Theorem 1 follows steps similar to those used to prove the results about ES under diffusion PDEs in [19, Theorem 1] or even [59, Theorem 1] for pure delays (transport PDEs). In this sense, we will simply point out the main differences for the case of games (not classical ES), instead of giving a full independent proof.

While in [19, Theorem 1] and [59, Theorem 1] it was possible to prove local exponential stability of the average closed-loop system using a Lyapunov functional, a different approach is adopted here for NES in noncooperative games. We will show that it is possible to guarantee the local exponential stability for the average closed-loop system (70)–(73) by means of a small-gain analysis.

First, consider the equivalent parabolic PDE–ODE representation (70)–(73) rewritten for each Player i , $\forall i \in \{1, \dots, N\}$:

$$\dot{\bar{G}}_i^{\text{av}}(t) = H_{ij}^i k_i \bar{G}_i^{\text{av}}(t) + \epsilon H_{ij}^i k_i \phi_i^{\text{av}}(1, t) \quad (S97)$$

$$\partial_t u_i^{\text{av}}(x, t) = D_i^{-2} \partial_{xx} u_i^{\text{av}}(x, t), \quad x \in (0, 1) \quad (S98)$$

$$\partial_x u_i^{\text{av}}(0, t) = 0 \quad (S99)$$

$$u_i^{\text{av}}(1, t) = k_i \bar{G}_i^{\text{av}}(t) + \epsilon k_i \phi_i^{\text{av}}(1, t). \quad (S100)$$

Hence, the average closed-loop system (S97)–(S100) satisfies all of the assumptions (A1) to (A7) of the Small-Gain Theorem [32, Theorem 8.2, p. 205] for the parabolic PDE–ODE loops with $p(z) = 1$, $r(z) = D_i^2$, $q(z) = 0$, $F(\bar{G}_i^{\text{av}}, u_i^{\text{av}}, 0) = H_{ij}^i k_i \bar{G}_i^{\text{av}} + \epsilon H_{ij}^i k_i \phi_i^{\text{av}}(1)$, $g(x, \bar{G}_i^{\text{av}}, u_i^{\text{av}}) = 0$, $f(x, t) = 0$, $\varphi_0(0, u_i^{\text{av}}, \bar{G}_i^{\text{av}}) = b_1 u_i^{\text{av}}(0, t)$, $b_1 < 0$, $b_2 = 1$, $\varphi_1(0, u_i^{\text{av}}, \bar{G}_i^{\text{av}}) = k_i \bar{G}_i^{\text{av}} + \epsilon k_i \phi_i^{\text{av}}(1)$, $a_1 = 1$, $a_2 = 0$, $L = \max(|H_{ij}^i| k_i, 1/\sqrt{3} \epsilon |H_{ij}^i| k_i k_H D_i^2)$, $K_0 = 1$, $B_0 = C_0 = 0$, γ_0 of order $O(\epsilon)$, $K_1 = 1/\sqrt{3} \epsilon k_i k_H D_i^2$, $B_1 = k_i$, $C_1 = 0$, γ_1 of order $O(1)$, $K_2 = B_2 = 0$, and $i \neq j$. The constant $k_H > 0$ is defined just after inequality (S101). Assumption (A6) of [32, Theorem 8.2, p. 205] holds with $M = 1$, $\gamma_3 = 1/\sqrt{3} \epsilon |H_{ij}^i| k_i k_H D_i^2$, and $\sigma = |H_{ij}^i| k_i$, as can be readily verified by means of the variation-of-constants formula

$$\begin{aligned} \bar{G}_i^{\text{av}}(t) &= \exp(-|H_{ij}^i| k_i t) \bar{G}_i^{\text{av}}(0) \\ &+ \int_0^t \exp(-|H_{ij}^i| k_i(t+s)) \epsilon H_{ij}^i k_i \phi_i^{\text{av}}(1, s) ds \end{aligned}$$

and from the application of the Cauchy–Schwarz inequality to the term $\phi_i^{\text{av}}(1, t)$ in (69):

$$\begin{aligned} \phi_i^{\text{av}}(1, t) &\leq \sum_{j \neq i}^N |H_{ij}^i| D_i^2 \left(\int_0^1 (1-\tau)^2 d\tau \right)^{\frac{1}{2}} \times \left(\int_0^1 [u_i^{\text{av}}(\xi, t)]^2 d\xi \right)^{\frac{1}{2}} \\ &\leq \frac{1}{\sqrt{3}} k_H \sum_{j \neq i}^N D_i^2 \left(\int_0^1 [u_i^{\text{av}}(\xi, t)]^2 d\xi \right)^{\frac{1}{2}} \quad (S101) \end{aligned}$$

since $|H_{ij}^i| < k_H < \frac{1}{\epsilon} |H_{ij}^i|$ according to Assumptions 1 and 2, where k_H is a positive constant of order $O(1)$. It follows that the small-gain condition [32, Inequality (8.3.24)]

$$\begin{aligned} \max(\gamma_0 K_0, \gamma_1 K_1) + \sigma^{-1} K_2 &< 1 \\ \gamma_3 \max(\gamma_0 B_0, \gamma_1 B_1) + \gamma_3 \sigma^{-1} B_2 &< 1 \quad (S102) \end{aligned}$$

holds, provided $0 < \epsilon < 1$ is sufficiently small. Therefore, if such a small-gain condition holds, then [32, Theorem 8.2, p. 205] allows us to conclude that there exist constants $\delta, \Delta > 0$ such that for every $u_0^{\text{av}} \in C^0([0, 1])$, $\bar{G}_0^{\text{av}} \in \mathbb{R}^n$, the unique generalized solution of this initial-boundary value problem, with $u^{\text{av}}(x, 0) = u_0^{\text{av}}$ and $\bar{G}^{\text{av}}(0) = \bar{G}_0^{\text{av}}$, satisfies the following estimate:

$$\|\bar{G}^{\text{av}}(t)\| + \|u^{\text{av}}(t)\|_{\infty} \leq \Delta (\|\bar{G}_0^{\text{av}}\| + \|u_0^{\text{av}}\|_{\infty}) \exp(-\delta t). \quad (S103)$$

Therefore, we conclude that the origin of the average closed-loop system (70)–(73) is exponentially stable under the assumption of $0 < \epsilon < 1$ being sufficiently small. Then, from (50) and (65), we conclude the same results in the norm

$$\left(\sum_{i=1}^N [\vartheta_i^{\text{av}}(t)]^2 + \int_0^D [u_i^{\text{av}}(x, t)]^2 dx \right)^{1/2} \quad (S104)$$

since H is nonsingular, that is, $|\vartheta_i^{\text{av}}(t)| \leq |H^{-1}| \|\bar{G}^{\text{av}}(t)\|$.

As developed in [19, Theorem 1], the next steps of the proof would be the application of the local averaging theory for infinite-dimensional systems in [25, Section 2] (see also “[Averaging Theorem for General Infinite-Dimensional Systems](#)”), showing that the periodic solutions satisfy (74) for ω sufficiently large and then the conclusion of the attractiveness of the Nash equilibrium Θ^* according to (75). The final residual set for the error $\theta(t) - \theta^*$ in (76) depends on $\text{lale}^{\max(D_i)\sqrt{\omega/2}}$ due to the amplitude of $S_i(t)$ in (42).

where $\alpha_i: [0, D_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}$, $\forall i \in \{1, 2\}$, and each domain length D_i is known.

The solution of (84)–(87) [40, Chapter 2] is simply

$$\alpha_1(x, t) = \theta_1(t + x - D_1) \quad (92)$$

which represents a delayed action for player P_1 at $x = 0$, whereas the solution of (88)–(91) is given by [40, Chapter 15]

$$\alpha_2(x, t) = \mathcal{L}^{-1} \left[\frac{\cosh(x\sqrt{s})}{\cosh(D_2\sqrt{s})} \right] * \theta_2(t) \quad (93)$$

where $\mathcal{L}^{-1}(\cdot)$ denotes the inverse Laplace transformation, and $*$ is the convolution operator. Given these relations, we define the heterogeneous transport–diffusion operator $\mathcal{H} = \text{diag}\{\mathcal{H}_1, \mathcal{H}_2\}$ for the PDEs (84)–(87) and (88)–(91) with boundary inputs and measurements given by

$$\begin{aligned} \mathcal{H}_1[\varphi(t)] &= \varphi(t + x - D_1), \text{ such that } \Theta_1(t) = \mathcal{H}_1[\theta_1(t)] \\ \mathcal{H}_2[\varphi(t)] &= \mathcal{L}^{-1} \left[\frac{1}{\cosh(D_2\sqrt{s})} \right] * \varphi(t), \\ \text{such that } \Theta_2(t) &= \mathcal{H}_2[\theta_2(t)]. \end{aligned} \quad (94)$$

Averaging Theorem for General Infinite-Dimensional Systems [25]

Consider the infinite-dimensional system, defined in the Banach space \mathcal{X}

$$\dot{z} = \mathcal{A}z + J(\omega t, z) \quad (\text{S105})$$

where $z(0) = z_0 \in \mathcal{X}$, and the operator $\mathcal{A}: \mathcal{D}(\mathcal{A}) \rightarrow \mathcal{X}$ generates an analytic semigroup. Moreover, the nonlinearity $J: \mathbb{R}_+ \times \mathcal{X} \rightarrow \mathcal{X}$ with $t \mapsto J(\omega t, z)$ is Fréchet differentiable in z and strongly continuous and periodic in t uniformly with respect to z in a compact subset of \mathcal{X} . Along with (S105), the average system

$$\dot{z}_{\text{av}} = \mathcal{A}z_{\text{av}} + J_0(z_{\text{av}}) \quad (\text{S106})$$

with $J_0(z_{\text{av}}) = \lim_{T \rightarrow \infty} 1/T \int_0^T J(\tau, z_{\text{av}}) d\tau$ is considered. Suppose that $z_{\text{av}} = 0 \in D \subset \mathcal{X}$ is an exponentially stable equilibrium

point of the average system (S106). Then, for some $\bar{\omega} > 0$ and $\omega > \bar{\omega}$, we have the following:

- 1) There exists a unique exponentially stable periodic solution $t \mapsto \bar{z}(t, 1/\omega)$, continuous in t and $1/\omega$, with $\|\bar{z}(t, 1/\omega)\| \leq \mathcal{O}(1/\omega)$ for $t > 0$.
- 2) With $\|z_0 - z_{\text{av}}(0)\| \leq \mathcal{O}(1/\omega)$, the solution estimate of (S105) is given by

$$\|z(t) - z_{\text{av}}\| \leq \mathcal{O}(1/\omega), \quad t > 0. \quad (\text{S107})$$

- 3) For $\|z_0\| \leq \mathcal{O}(1/\omega)$, and the stable manifold theorem, it holds that

$$\|z(t) - \bar{z}(t, 1/\omega)\| \leq Ce^{-\gamma t}, \quad t > 0 \quad (\text{S108})$$

for some $C, \gamma > 0$.

Proof of Theorem 2

The proof of Theorem 2 follows the same structure of that of Theorem 1, but now considering the equivalent hyperbolic PDE–ODE representation for each Player i , $\forall i \in \{1, \dots, N\}$:

$$\dot{\bar{G}}_i^{\text{av}}(t) = H_i^i k_i \bar{G}_i^{\text{av}}(t) + \epsilon H_i^i k_i \phi_i^{\text{av}}(1, t) \quad (\text{S109})$$

$$\partial_t u_i^{\text{av}}(x, t) = D_i^{-1} \partial_x u_i^{\text{av}}(x, t), \quad x \in (0, 1) \quad (\text{S110})$$

$$u_i^{\text{av}}(1, t) = k_i \bar{G}_i^{\text{av}}(t) + \epsilon k_i \phi_i^{\text{av}}(1, t) \quad (\text{S111})$$

where $H_i^i < 0$, $k_i > 0$, $0 < \epsilon < 1$, $D_i^{-1} > 0$, and

$$\phi_i(1, t) := - \sum_{j \neq i} H_j^i \int_0^1 D_j u_j(\xi, t) d\xi. \quad (\text{S112})$$

In this case, the average closed-loop system (S109)–(S111) satisfies both assumptions (H1) and (H2) of the Small-Gain Theorem [32, Theorem 8.1, p. 198] for the hyperbolic PDE–ODE loops rather than [32, Theorem 8.2, p. 205] employed in the proof of Theorem 1; this latter oriented for parabolic PDE–ODE interconnections—see “Small-Gain Theorem for ODE and Hyperbolic PDE Loops.”

It follows that the small-gain condition in [32, Theorem 8.1, p. 198] holds, provided $0 < \epsilon < 1$ is sufficiently small. Therefore, if such a small-gain condition holds, then [32, Theorem 8.1, p. 198] allows us to conclude that the origin of the average closed-loop system (S109)–(S111) is exponentially stable. Then, from (80) and the following transformation [7]:

$$\begin{aligned} \bar{G}_i^{\text{av}}(t) &= \hat{G}_i^{\text{av}}(t) + \sum_{j=1}^N \epsilon_j^i H_j^i \int_{t-D_j}^t U_j^{\text{av}}(\tau) d\tau \\ &= \hat{G}_i^{\text{av}}(t) + \sum_{j=1}^N \epsilon_j^i H_j^i \int_0^{D_j} u_j^{\text{av}}(\xi, t) d\xi \end{aligned} \quad (\text{S113})$$

we can conclude the same results in the norm

$$\left(\sum_{j=1}^N [\bar{\theta}_j^{\text{av}}(t - D_j)]^2 + \int_0^{D_i} [u_i^{\text{av}}(\tau)]^2 d\tau \right)^{1/2}. \quad (\text{S114})$$

The application of the local averaging theory for functional differential equations in [25] (see also “Averaging Theorem for Functional Differential Equations”) shows that the periodic solutions satisfy inequality (82) and leads to the conclusion of the attractiveness of the Nash equilibrium θ^* according to (83).

We consider games where the payoff function $y_i(t) = J_i(\Theta(t))$, $\forall i \in \{1, 2\}$, of each player is quadratic [9], expressed as a strictly concave combination of their actions propagated through distinct transport–heat PDE dynamics

$$J_1(\Theta(t)) = \frac{H_{11}^1}{2} \Theta_1^2(t) + \frac{H_{22}^1}{2} \Theta_2^2(t) + \epsilon H_{12}^1 \Theta_1(t) \Theta_2(t) + h_1^1 \Theta_1(t) + h_2^1 \Theta_2(t) + c_1 \quad (95)$$

$$J_2(\Theta(t)) = \frac{H_{11}^2}{2} \Theta_1^2(t) + \frac{H_{22}^2}{2} \Theta_2^2(t) + \epsilon H_{21}^2 \Theta_1(t) \Theta_2(t) + h_1^2 \Theta_1(t) + h_2^2 \Theta_2(t) + c_2 \quad (96)$$

where $J_1(\Theta)$, $J_2(\Theta): \mathbb{R}^2 \rightarrow \mathbb{R}$, H_{jk}^i , h_j^i , $c_i \in \mathbb{R}$ are constants, $H_{ii}^i < 0$, $\forall i, j, k \in \{1, 2\}$, and $\epsilon > 0$ without loss of generality.

For the sake of completeness, we provide here, in mathematical terms, the definition of a Nash equilibrium $\Theta^* = [\Theta_1^*, \Theta_2^*]^T$ in a two-player game:

$$J_1(\Theta_1^*, \Theta_2^*) \geq J_1(\Theta_1, \Theta_2^*) \quad \text{and} \quad J_2(\Theta_1^*, \Theta_2^*) \geq J_2(\Theta_1^*, \Theta_2) \quad (97)$$

Hence, no player has any incentive to unilaterally deviate its action from Θ^* .

As done before to determine (6) in the general case, the Nash equilibrium solution $\Theta^* = -H^{-1}h$ for the duopoly is simply

$$\begin{bmatrix} \Theta_1^* \\ \Theta_2^* \end{bmatrix} = - \begin{bmatrix} H_{11}^1 & \epsilon H_{12}^1 \\ \epsilon H_{21}^2 & H_{22}^2 \end{bmatrix}^{-1} \begin{bmatrix} h_1^1 \\ h_2^2 \end{bmatrix}. \quad (98)$$

The *objective* is to design an ES-based strategy to reach the Nash equilibrium in heterogeneous noncooperative games subjected to transport–heat PDEs in the decision variables of the players (input signals).

Since our goal is to find the unknown optimal inputs Θ^* (and θ^*), we define the estimation errors

$$\tilde{\theta}(t) = \hat{\theta}(t) - \theta^*, \quad \vartheta(t) = \hat{\Theta}(t) - \Theta^* \quad (99)$$

where the vectors $\hat{\theta}(t)$ and $\hat{\Theta}(t)$ are estimates of θ^* and Θ^* . To make (99) coherent with the optimizer of the static map Θ^* , we apply the heterogeneous transport–diffusion operator (94) to $\tilde{\theta}_i$ in (99), arriving at

$$\vartheta_1(t) = \tilde{\theta}_1(t - D_1) := \tilde{\alpha}_1(0, t) \quad (100)$$

$$\partial_t \tilde{\alpha}_1(x, t) = \partial_x \tilde{\alpha}_1(x, t), \quad x \in (0, D_1) \quad (101)$$

$$\partial_x \tilde{\alpha}_1(0, t) = 0 \quad (102)$$

$$\tilde{\alpha}_1(D_1, t) = \tilde{\theta}_1(t) \quad (103)$$

and

$$\vartheta_2(t) := \tilde{\alpha}_2(0, t) \quad (104)$$

$$\partial_t \tilde{\alpha}_2(x, t) = \partial_{xx} \tilde{\alpha}_2(x, t), \quad x \in (0, D_2) \quad (105)$$

$$\partial_x \tilde{\alpha}_2(0, t) = 0 \quad (106)$$

$$\tilde{\alpha}_2(D_2, t) = \tilde{\theta}_2(t) \quad (107)$$

where $\tilde{\alpha}_i: [0, D_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}$, $\forall i \in \{1, 2\}$, and $\vartheta(t) := \mathcal{H}[\tilde{\theta}(t)] = \hat{\Theta}(t) - \Theta^*$ is the propagated estimation error $\tilde{\theta}(t)$ through the transport–diffusion domain. For $\lim_{t \rightarrow \infty} \theta(t) = \theta_c$, we have $\lim_{t \rightarrow \infty} \theta(t) = \Theta_c = \theta_c$, where the index c indicates a constant signal. Indeed, from (94), for a constant input $\theta = \theta_c$, we have $\lim_{t \rightarrow \infty} \mathcal{H}_1[\theta_{c1}] = \theta_{c1}$ for Player P_1 . For Player P_2 , we have $\mathcal{L}\{\theta_{c2}\} = \theta_{c2}/s$, and applying the Laplace limit theorem, we get $\lim_{t \rightarrow \infty} \mathcal{H}_2[\theta_{c2}] = \lim_{s \rightarrow 0} \{\theta_{c2}/\cosh(\sqrt{s} D_2)\} = \theta_{c2}$. Thus, in the particular case $\theta = \theta_c = \theta^*$, one has

$$\Theta^* = \theta^*. \quad (108)$$

Figure 6 depicts a schematic diagram that summarizes the proposed Nash equilibrium policy for each player.

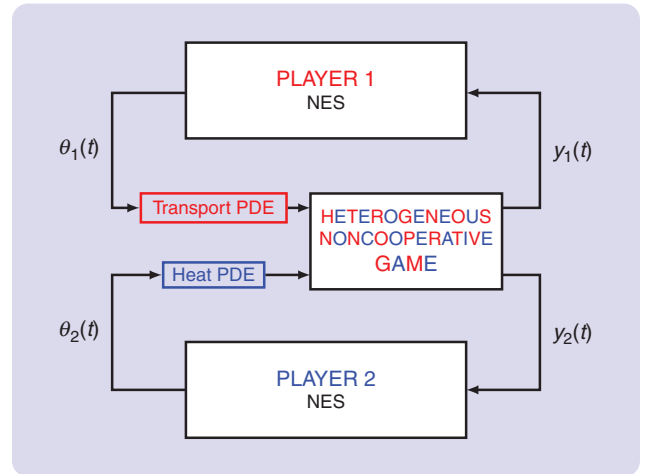


FIGURE 5 NES in a heterogeneous noncooperative game with players acting through transport–heat PDE dynamics.

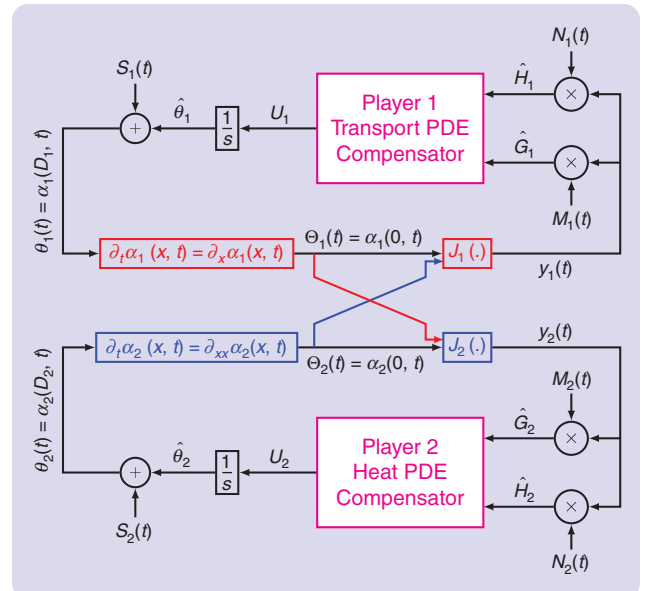


FIGURE 6 Block diagram illustrating the NES strategy performed for each player. In magenta are shown the boundary controllers used to compensate for the individual transport–heat PDEs for the heterogeneous noncooperative game.

Small-Gain Theorem for ODE and Hyperbolic PDE Loops [32, Theorem 8.1, p. 198]

Consider generalized solutions of the following initial-boundary value problem:

$$\dot{x}(t) = F(x(t), u(z, t), v(t)), \quad \forall t \geq 0 \quad (\text{S115})$$

$$u_t(z, t) + cu_z(z, t) = a(z)u(z, t) + g(z, x(t), u(z, t)) + f(z, t), \quad \forall (z, t) \in [0, 1] \times \mathbb{R}_+ \quad (\text{S116})$$

$$u(0, t) = \varphi(d(t), u(z, t), x(t)), \quad \forall t \geq 0, \quad u(z, 0) = u_0, \quad x(0) = x_0. \quad (\text{S117})$$

The state of the system (S115)–(S117) is $(u(z, t), x(t)) \in C^0([0, 1] \times \mathbb{R}_+) \times \mathbb{R}^n$, while the other variables $d \in C^0(\mathbb{R}_+; \mathbb{R}^q)$, $f \in C^0([0, 1] \times \mathbb{R}_+)$ and $v \in C^0(\mathbb{R}_+; \mathbb{R}^m)$ are external inputs. We assume that $(0, 0) \in C^0([0, 1]) \times \mathbb{R}^n$ is an equilibrium point for the input-free system; that is, $F(0, 0, 0) = 0$, $g(z, 0, 0) = 0$, and $\varphi(0, 0, 0) = 0$. Now, we assume that the ODE subsystem satisfies the ISS property:

(H1) There exist constants $M, \sigma > 0$, $b_3, \gamma_3 \geq 0$, such that for every $x_0 \in \mathbb{R}^n$, $u \in C^0([0, 1] \times \mathbb{R}_+)$, and $v \in C^0(\mathbb{R}_+; \mathbb{R}^m)$, the unique solution $x \in C^1(\mathbb{R}_+; \mathbb{R}^n)$ of (S115) with $x(0) = x_0$ satisfies the following estimate, $\forall t \geq 0$:

$$|x(t)| \leq M|x_0| \exp(-\sigma t) + \max_{0 \leq s \leq t} (\gamma_3 \|u(s)\|_\infty + b_3 |v(s)|). \quad (\text{S118})$$

We next need to estimate the static gain of the interconnections. To this purpose, we employ the following further assumption:

(H2) There exist constants $b_2, \gamma_1, \gamma_2, A, B \geq 0$ such that the following growth conditions hold for every $x \in C^1(\mathbb{R}_+; \mathbb{R}^n)$, $u \in C^0([0, 1] \times \mathbb{R}_+)$, and $d \in C^0(\mathbb{R}_+; \mathbb{R}^q)$:

$$|g(z, x, u)| \leq A \|u\|_\infty + \gamma_1 |x|, \quad \forall z \in [0, 1] \quad (\text{S119})$$

$$|\varphi(d, u, x)| \leq B \|u\|_\infty + \gamma_2 |x| + b_2 |d|. \quad (\text{S120})$$

Let $c > 0$ ^a be a given constant and $a \in C^0([0, 1])$ be a given function. Consider the mappings as $F: \mathbb{R}^n \times C^0([0, 1]) \times \mathbb{R}^m \rightarrow \mathbb{R}^n$,

$g: [0, 1] \times \mathbb{R}^n \times C^0([0, 1]) \rightarrow \mathbb{R}$, $\varphi: \mathbb{R}^q \times C^0([0, 1]) \times \mathbb{R}^n \rightarrow \mathbb{R}$ being continuous mappings with $F(0, 0, 0) = 0$ for which there exist constants $L > 0$, $\tilde{N} \in [0, 1[$ such that the inequalities $\max_{0 \leq z \leq 1} (|g(z, x, u) - g(z, y, w)|) + |F(x, u, v) - F(y, w, v)| \leq L \|x - y\| + L \|u - w\|_\infty$ and $|\varphi(d, u, x) - \varphi(d, w, y)| \leq \tilde{N} \|x - y\| + \tilde{N} \|u - w\|_\infty$ hold for all $u, w \in C^0([0, 1])$, $x, y \in \mathbb{R}^n$, $v \in \mathbb{R}^m$, and $d \in \mathbb{R}^q$. Suppose that assumptions (H1) and (H2) hold and that the following small-gain condition is satisfied:

$$(\gamma_1 \gamma_3 + A) c^{-1} \max_{0 \leq z \leq 1} \left(\rho(z) \int_0^z \frac{1}{\rho(l)} dl \right) + (\gamma_2 \gamma_3 + B) \max_{0 \leq z \leq 1} (\rho(z)) + 2 \sqrt{(\gamma_1 \gamma_3 + A) c^{-1} (\gamma_2 \gamma_3 + B) \max_{0 \leq z \leq 1} (\rho(z)) \max_{0 \leq z \leq 1} \left(\rho(z) \int_0^z \frac{1}{\rho(l)} dl \right)} < 1, \quad (\text{S121})$$

with $\rho(z) := \exp\left(c^{-1} \int_0^z a(w) dw\right)$ for $z \in [0, 1]$ [recall (8.2.11) and (8.2.14)] in [32, Section 8.2]. Then, there exist constants $\delta, \Theta, \gamma > 0$ such that, for every $u_0 \in C^0([0, 1])$, $x_0 \in \mathbb{R}^n$, $d \in C^0(\mathbb{R}_+; \mathbb{R}^q)$ with $u_0(0) = \varphi(d(0), u_0, x_0)$, $f \in C^0([0, 1] \times \mathbb{R}_+)$, and $v \in C^0(\mathbb{R}_+; \mathbb{R}^m)$, the unique generalized solution of the initial-boundary value problem (S115), (S116), (S117) satisfies the following estimate:

$$|x(t)| + \|u(t)\|_\infty \leq \Theta (|x_0| + \|u_0\|_\infty) \exp(-\delta t) + \gamma \left[\max_{0 \leq s \leq t} (|v(s)|) + \max_{0 \leq s \leq t} (\|f(s)\|_\infty) + \max_{0 \leq s \leq t} (|d(s)|) \right], \quad \forall t \geq 0. \quad (\text{S122})$$

Analogous small-gain results for parabolic PDE–ODE loops and parabolic–hyperbolic PDE loops can be found in [32, Theorem 8.2, p. 205] and [32, Theorem 11.2 (11.5), p. 269 (277)], respectively. Such theorems are used in “Proof of Theorem 1,” “Proof of Theorem 2,” and “Proof of Theorem 3.”

^aIf the scalar $c < 0$ is considered, the direction of convection must be reversed such that the boundary $u(0, t)$ is replaced by $u(1, t)$ and vice versa.

While the signals $M_i(t)$, $N_i(t)$, $\hat{H}_i(t)$, and $\hat{G}_i(t)$, $\forall i \in \{1, 2\}$, are the same as in the sections “Noncooperative Scenario With Delays” and “Noncooperative Games With Heat PDEs,” the additive dither signals in the presence of transport–heat PDE dynamics [19], [67] are redefined according to

$$\begin{cases} S_1(t) = a_1 \sin(\omega_1(t + D_1)) \\ S_2(t) = \frac{1}{2} a_2 e^{\sqrt{\frac{\omega_2}{2}} D_2} \sin(\omega_2 t + \sqrt{\frac{\omega_2}{2}} D_2) \\ \quad + \frac{1}{2} a_2 e^{-\sqrt{\frac{\omega_2}{2}} D_2} \sin(\omega_2 t - \sqrt{\frac{\omega_2}{2}} D_2) \end{cases}. \quad (\text{109})$$

Hence, computing the average of $\hat{G}_1(t)$ and $\hat{G}_2(t)$ leads us to

$$\begin{cases} \hat{G}_1^{\text{av}}(t) = H_{11}^1 \vartheta_1^{\text{av}}(t) + \epsilon H_{12}^1 \vartheta_2^{\text{av}}(t) \\ \hat{G}_2^{\text{av}}(t) = \epsilon H_{21}^2 \vartheta_1^{\text{av}}(t) + H_{22}^2 \vartheta_2^{\text{av}}(t) \end{cases} \quad (\text{110})$$

Additionally, from the block diagram in Figure 6, one has

$$\dot{\hat{\theta}}_i(t) = U_i(t), \quad \dot{\hat{\theta}}_i(t) = U_i(t), \quad \forall i \in \{1, 2\} \quad (\text{111})$$

since $\ddot{\hat{\theta}}(t) = \dot{\hat{\theta}}(t)$, once θ^* is constant. Taking the time derivative of (100)–(103) and (104)–(107), with the help of (99) and (111), the propagated error dynamics can be written as

$$\dot{\hat{\theta}}_1(t) = u_1(0, t) = U(t - D_1) \quad (\text{112})$$

$$\partial_t u_1(x, t) = \partial_x u_1(x, t), \quad x \in (0, D_1) \quad (\text{113})$$

$$\partial_x u_1(0, t) = 0 \quad (\text{114})$$

$$u_1(D_1, t) = U_1(t) \quad (\text{115})$$

and

Averaging Theorem for Functional Differential Equations [25]

Consider the delay system

$$\dot{x}(t) = f(t/\epsilon, x_t), \quad \forall t \geq 0 \quad (\text{S123})$$

where ϵ is a real parameter, $x_t(\Theta) = x(t + \Theta)$ for $-r \leq \Theta \leq 0$, and $f: \mathbb{R}_+ \times \Omega \rightarrow \mathbb{R}^n$ is a continuous functional from a neighborhood Ω of 0 of the supremum-normed Banach space $X = C([-r, 0]; \mathbb{R}^n)$ of continuous functions from $[-r, 0]$ to \mathbb{R}^n . Assume that $f(t, \varphi)$ is periodic in t uniformly with respect to φ in compact subsets of Ω and that f has a continuous Fréchet derivative $\partial f(t, \varphi)/\partial \varphi$ in φ on $\mathbb{R}_+ \times \Omega$.

If $y = y_0 \in \Omega$ is an exponentially stable equilibrium for the average system

$$\dot{y}(t) = f_0(y_t), \quad \forall t \geq 0 \quad (\text{S124})$$

where $f_0(\varphi) = \lim_{T \rightarrow \infty} 1/T \int_0^T f(s, \varphi) ds$, then, for some $\epsilon_0 > 0$ and $0 < \epsilon \leq \epsilon_0$, there is a unique periodic solution $t \mapsto x^*(t, \epsilon)$ of (S123) with the properties of being continuous in t and ϵ , satisfying $|x^*(t, \epsilon) - y_0| \leq \mathcal{O}(\epsilon)$, for $t \in \mathbb{R}_+$, and such that there is $\rho > 0$ so that if $x(\cdot; \varphi)$ is a solution of (S123) with $x(s) = \varphi$ and $|\varphi - y_0| < \rho$, then $|x(t) - x^*(t, \epsilon)| \leq Ce^{-\gamma(t-s)}$, for $C > 0$ and $\gamma > 0$.

$$\dot{\vartheta}_2(t) = u_2(0, t) \quad (116)$$

$$\partial_x u_2(x, t) = \partial_{xx} u_2(x, t), \quad x \in (0, D_2) \quad (117)$$

$$\partial_x u_2(0, t) = 0 \quad (118)$$

$$u_2(D_2, t) = U_2(t) \quad (119)$$

where $u_i: [0, D_i] \times \mathbb{R}_+ \rightarrow \mathbb{R}$, $u_i(x, t) := \partial_t \bar{\alpha}_i(x, t)$, $\forall i \in \{1, 2\}$, and $\bar{\alpha}_i(x, t) = \alpha_i(x, t) - \beta_i(x, t) - \Theta_i^*$. The term $\beta_i(x, t)$ is the PDE state of the trajectory generation problem [42, Chapter 12] solved to obtain $S_1(t) = \beta_1(D_1, t)$ and $S_2(t) = \beta_2(D_2, t)$ in (109)—for more details, see [53, Equations (19)–(22)].

Hence, from (110), (112)–(115), and (116)–(119), it is possible to find a compact form for the overall average estimated gradient according to

$$\hat{G}^{\text{av}}(t) = H \hat{\vartheta}^{\text{av}}(t) \quad (120)$$

$$\dot{\hat{G}}^{\text{av}}(t) = H \dot{\hat{\vartheta}}^{\text{av}}(t) = H \mathcal{H}[U^{\text{av}}(t)] \quad (121)$$

where the Hessian H is given in (98), and $\vartheta^{\text{av}}(t) := [\vartheta_1^{\text{av}}(t), \vartheta_2^{\text{av}}(t)]^T \in \mathbb{R}^2$, $\hat{G}^{\text{av}}(t) := [\hat{G}_1^{\text{av}}(t), \hat{G}_2^{\text{av}}(t)]^T \in \mathbb{R}^2$, and $U^{\text{av}}(t) := [U_1^{\text{av}}(t), U_2^{\text{av}}(t)]^T \in \mathbb{R}^2$ are the average versions of $\vartheta(t) := [\vartheta_1(t), \vartheta_2(t)]^T$, $\hat{G}(t) := [\hat{G}_1(t), \hat{G}_2(t)]^T$, and $U(t) := [U_1(t), U_2(t)]^T$, respectively.

Decentralized PDE Boundary Control Using Only the Known Diagonal Terms of the Hessian

The control laws must be able to ensure the exponential stabilization of $\hat{G}^{\text{av}}(t)$ by compensating the heterogeneous transport–diffusion operator $\mathcal{H}[\cdot]$ in (121). Following [19] and [59], we propose the following boundary-based update laws $\hat{\vartheta}_i(t) = U_i(t)$, $i \in \{1, 2\}$:

$$\begin{cases} \dot{U}_1(t) = -c_1 U_1(t) + c_1 k_1 \left(\hat{G}_1(t) + \hat{H}_1(t) \int_0^{D_1} u_1(\tau, t) d\tau \right) \\ \dot{U}_2(t) = -c_2 U_2(t) + c_2 k_2 \left(\hat{G}_2(t) + \hat{H}_2(t) \int_0^{D_2} (D_2 - \tau) u_2(\tau, t) d\tau \right) \end{cases} \quad (122)$$

for positive constants $k_1 > 0$, $k_2 > 0$, $c_1 > 0$, and $c_2 > 0$ to compensate for the transport–heat PDEs in (112)–(115) and (116)–(119). As before in the sections “Noncooperative

Scenario with Delays” and “Noncooperative Games with Heat PDEs,” and with some abuse of notation, the constants c_1 and c_2 were chosen to denote the parameters of the control laws, but they have no relation to those that appear in the payoffs given by (95) and (96).

The boundary control law (122) could be rewritten as

$$\begin{cases} \dot{U}_1(t) = -c_1 U_1(t) + c_1 k_1 \left(\hat{G}_1(t) + \hat{H}_1(t) \int_{t-D_1}^t U_1(\tau, t) d\tau \right) \\ \dot{U}_2(t) = -c_2 U_2(t) + c_2 k_2 \left[\hat{G}_2(t) + \hat{H}_2(t) (\hat{\vartheta}_2(t) - \Theta_2(t) + a_2 \sin(\omega_2 t)) \right] \end{cases} \quad (123)$$

using the relation $u_1(x, t) = U_1(t + x - D_1)$ for the transport PDE and the diffusion equation $\partial_x u_2(x, t) = \partial_{xx} u_2(x, t)$ as well as the integration by parts associated with (88)–(91) and (99) and recalling that $\vartheta_2 + a_2 \sin(\omega_2 t) = \Theta_2(t) - \Theta_2^*$, analogously to [19, Equation (25)].

ISS-Like Properties for Hyperbolic–Parabolic PDE Representation

For simplicity, we assume $c_1, c_2 \rightarrow +\infty$ in (122), resulting in the following general expressions:

$$\begin{aligned} U_1(t) &= k_1 \left(\hat{G}_1(t) + \hat{H}_1(t) \int_0^{D_1} u_1(\tau, t) d\tau \right) \\ U_2(t) &= k_2 \left(\hat{G}_2(t) + \hat{H}_2(t) \int_0^{D_2} (D_2 - \tau) u_2(\tau, t) d\tau \right). \end{aligned} \quad (124)$$

Recalling (112)–(115) and (116)–(119), the infinite-dimensional closed-loop system (121) and (124) in its average version can be written in the corresponding PDE representation form, given by

$$\dot{\hat{G}}_1^{\text{av}}(t) = H_{11}^1 u_1^{\text{av}}(0, t) + \epsilon H_{12}^1 u_2^{\text{av}}(0, t) \quad (125)$$

$$\partial_x u_1^{\text{av}}(x, t) = D_1^{-1} \partial_x u_1^{\text{av}}(x, t), \quad x \in (0, 1) \quad (126)$$

$$\partial_x u_1^{\text{av}}(0, t) = 0 \quad (127)$$

$$u_1^{\text{av}}(1, t) = U_1^{\text{av}}(t) \quad (128)$$

and

$$\dot{\hat{G}}_2^{\text{av}}(t) = \epsilon H_{21}^2 u_1^{\text{av}}(0, t) + H_{22}^2 u_2^{\text{av}}(0, t) \quad (129)$$

$$\partial_t u_2^{\text{av}}(x, t) = D_2^{-2} \partial_{xx} u_2^{\text{av}}(x, t), \quad x \in (0, 1) \quad (130)$$

$$\partial_x u_2^{\text{av}}(0, t) = 0 \quad (131)$$

$$u_2^{\text{av}}(1, t) = U_2^{\text{av}}(t). \quad (132)$$

In the *reduction-like approach* [7] (or finite-spectrum assignment), we use the following transformations to write

$$\begin{aligned} \tilde{G}_1^{\text{av}}(t) &= \hat{G}_1^{\text{av}}(t) + \epsilon_{11}^1 H_{11}^1 \int_0^{D_1} u_1^{\text{av}}(\tau, t) d\tau \\ &\quad + \epsilon_{12}^1 H_{12}^1 \int_0^{D_2} (D_2 - \tau) u_2^{\text{av}}(\tau, t) d\tau \end{aligned} \quad (133)$$

and

$$\begin{aligned} \tilde{G}_2^{\text{av}}(t) &= \hat{G}_2^{\text{av}}(t) + \epsilon_{22}^2 H_{22}^2 \int_0^{D_2} (D_2 - \tau) u_2^{\text{av}}(\tau, t) d\tau \\ &\quad + \epsilon_{21}^2 H_{21}^2 \int_0^{D_1} u_1^{\text{av}}(\tau, t) d\tau \end{aligned} \quad (134)$$

where $\epsilon_{11}^1 = \epsilon_{22}^2 = 1$ and $\epsilon_{12}^1 = \epsilon_{21}^2 = \epsilon$.

With some mathematical manipulations, it is not difficult to see that \tilde{G}^{av} satisfies

$$\dot{\tilde{G}}^{\text{av}}(t) = H U^{\text{av}}(t). \quad (135)$$

After adding and subtracting the next terms in blue and red in (124), it can be rewritten as

$$\begin{aligned} U_1(t) &= k_1 \left(\hat{G}_1(t) + \hat{H}_1(t) \int_0^{D_1} u_1(\tau, t) d\tau \right. \\ &\quad \left. + \epsilon H_{12}^1 \int_0^{D_2} (D_2 - \tau) u_2(\tau, t) d\tau \right) \\ &\quad - k_1 \epsilon H_{12}^1 \int_0^{D_2} (D_2 - \tau) u_2(\tau, t) d\tau \end{aligned} \quad (136)$$

$$\begin{aligned} U_2(t) &= k_2 \left(\hat{G}_2(t) + \hat{H}_2(t) \int_0^{D_2} (D_2 - \tau) u_2(\tau, t) d\tau \right. \\ &\quad \left. + \epsilon H_{21}^2 \int_0^{D_1} u_1(\tau, t) d\tau - \epsilon H_{21}^2 \int_0^{D_1} u_1(\tau, t) d\tau \right) \end{aligned} \quad (137)$$

whose average compact form is

$$U^{\text{av}}(t) = K \tilde{G}^{\text{av}}(t) + \epsilon K \phi^{\text{av}}(D, t) \quad (138)$$

where the matrix $K := \text{diag}\{k_1, k_2\}$ with entries $k_1 > 0$, $k_2 > 0$, and the auxiliary variable $\phi(D, t)$ is defined as

$$\begin{aligned} \phi(D, t) &:= - \begin{bmatrix} H_{12}^1 \int_0^{D_2} (D_2 - \tau) u_2(\tau, t) d\tau \\ H_{21}^2 \int_0^{D_1} u_1(\tau, t) d\tau \end{bmatrix} \\ \phi(1, t) &:= - \begin{bmatrix} H_{12}^1 \int_0^1 D_2^2 (1 - \xi) u_2(\xi, t) d\xi \\ H_{21}^2 \int_0^1 D_1 u_1(\xi, t) d\xi \end{bmatrix} \end{aligned} \quad (139)$$

since $\int_0^{D_j} (D_j - \tau) u_j(\tau, t) d\tau = \int_0^1 D_j^2 (1 - \xi) u_j(\xi, t) d\xi$, for $j \in \{1, 2\}$. Then, it is possible to find a compact form for the overall average game from (135) and (138), such as

$$\dot{\tilde{G}}^{\text{av}}(t) = HK \tilde{G}^{\text{av}}(t) + \epsilon HK \phi^{\text{av}}(1, t) \quad (140)$$

$$\partial_t u_1^{\text{av}}(x, t) = D_1^{-1} \partial_x u_1^{\text{av}}(x, t), \quad x \in (0, 1) \quad (141)$$

$$\partial_t u_2^{\text{av}}(x, t) = D_2^{-2} \partial_{xx} u_2^{\text{av}}(x, t), \quad x \in (0, 1) \quad (142)$$

$$\partial_x u^{\text{av}}(0, t) = 0 \quad (143)$$

$$u^{\text{av}}(1, t) = K \tilde{G}^{\text{av}}(t) + \epsilon K \phi^{\text{av}}(1, t). \quad (144)$$

From (140), if HK is Hurwitz, it is clear that the dynamics of the ODE state variable $\tilde{G}^{\text{av}}(t)$ is exponentially ISS [32] with respect to the PDE state $u^{\text{av}}(x, t) = [u_1^{\text{av}}(x, t), u_2^{\text{av}}(x, t)]^T$ using the function $\phi^{\text{av}}(1, t)$. Moreover, the PDE subsystem (141) and (142) is ISS [32] with respect to $\tilde{G}^{\text{av}}(t)$ in the boundary condition $u^{\text{av}}(1, t)$.

Stability Analysis

Next, we will show here that this hyperbolic-parabolic PDE-ODE loop (140)–(144) contains a small parameter ϵ that can lead to closed-loop stability if it is chosen sufficiently small. To this end, we assume the following particular condition for duopoly games [21], which is equivalent to Assumptions 1 and 2 when payoff functions of the form of (95) and (96) are considered.

Assumption 3

The Hessian matrix H given by (98) is strictly diagonal dominant, that is

$$\begin{cases} |\epsilon H_{12}^1| < |H_{11}^1| \\ |\epsilon H_{21}^2| < |H_{22}^2|. \end{cases} \quad (145)$$

The next theorem provides the stability/convergence properties of the closed-loop error system of the proposed ES feedback for the two-player noncooperative game with transport-heat PDEs.

Theorem 3

Consider the closed-loop system (112)–(119) under transport-heat PDEs (84)–(91) of distinct transport-diffusion coefficients D_1 and D_2 for the heterogeneous duopoly quadratic game with payoff functions $y_i(t) = J_i(\Theta(t))$, $\forall i \in \{1, 2\}$, given in (95) and (96), satisfying Assumption 3 and control laws $U_i(t)$ defined in (122) or (123). Thus, there exist $c_1 > 0$, $c_2 > 0$, and $\omega > 0$ sufficiently large as well as $\epsilon > 0$ sufficiently small such that the closed-loop error system with state $\vartheta_i(t)$, $u_i(x, t)$, $\forall i \in \{1, 2\}$, has a unique locally exponentially stable periodic solution in t of period Π in (19), denoted by $\vartheta_i^\Pi(t)$, $u_i^\Pi(x, t)$ and satisfying $\forall t \geq 0$:

$$\left(\sum_{i=1}^2 [\vartheta_i^\Pi(t)]^2 + \int_0^{D_1} [u_1^\Pi(x, t)]^2 dx \right)^{1/2} \leq \mathcal{O}\left(\frac{1}{\omega}\right). \quad (146)$$

In particular,

$$\limsup_{t \rightarrow +\infty} |\Theta(t) - \Theta^*| = \mathcal{O}\left(|a| + \frac{1}{\omega}\right) \quad (147)$$

$$\limsup_{t \rightarrow +\infty} |\theta_1(t) - \theta_1^*| = \mathcal{O}\left(a_1 + \frac{1}{\omega}\right) \quad (148)$$

$$\limsup_{t \rightarrow +\infty} |\theta_2(t) - \theta_2^*| = \mathcal{O}\left(a_2 e^{D_2 \sqrt{\omega/2}} + \frac{1}{\omega}\right) \quad (149)$$

where $a = [a_1 \ a_2]^T$ and $\theta^* = \Theta^*$ is the unique (unknown) Nash equilibrium given by (98).

Proof

See “Proof of Theorem 3.” ■

SIMULATIONS

Because of space limitations, we will restrict ourselves to a unique numerical example considering a heterogeneous noncooperative game with two players that employ the proposed ES strategy for PDE compensation of the section “Noncooperative Duopoly With Heterogeneous Transport–Heat PDE Dynamics.” See also “Derivation of Computational Simulation Models for PDEs.” We revisit the example in [62] and consider the following payoff functions from (95) and (96) subject to transport–heat PDEs (84)–(87) and (88)–(91) with distinct transport–diffusion coefficients $D_1 = 30$ and $D_2 = 3$ in the players’ decisions, $i \in \{1, 2\}$:

$$J_1(\Theta(t)) = -5 \Theta_1^2(t) + 5 \epsilon \Theta_1(t) \Theta_2(t) + 250 \Theta_1(t) - 150 \Theta_2(t) - 3,000 \quad (150)$$

$$J_2(\Theta(t)) = -5 \Theta_2^2(t) + 5 \epsilon \Theta_1(t) \Theta_2(t) - 150 \Theta_1(t) + 150 \Theta_2(t) + 2,500 \quad (151)$$

which, according to (98), yield the unique Nash equilibrium

$$\theta_1^* = \theta_1^* = \frac{100 + 30\epsilon}{4 - \epsilon^2}, \quad \theta_2^* = \theta_2^* = \frac{60 + 50\epsilon}{4 - \epsilon^2}. \quad (152)$$

To attain (152), the players implement the nonmodel-based real-time optimization strategy acting through the transport–heat PDE dynamics (see Figure 6). For comparison purposes, except for transport–heat PDEs in the input signals of the players, the plant and the controller parameters were chosen similarly to those in [21] in the simulation tests: $\epsilon = 1$, $a_1 = 0.075$, $a_2 = 0.05$, $k_1 = 2$, $k_2 = 5$, $\omega_1 = 26.75$ rad/s, $\omega_2 = 22$ rad/s, $\theta_1(0) = \hat{\theta}_1(0) = 50$, and $\theta_2(0) = \hat{\theta}_2(0) = \theta_2^* = 110/3$. In addition, the time constants of the boundary control filters were set to $c_1 = c_2 = 100$.

We can verify that the ES approach proposed in [21] is effective when transport–heat PDEs are not present in the decision variables. However, in the presence of the transport–heat PDEs in the input signals θ_1 and θ_2 , but without considering any kind of PDE compensation, the game collapses with the explosion of its variables (curves not shown). On the other hand, Figure 7(a) and (b) shows that the proposed boundary control-based scheme fixes this with a remarkable evolution in searching the Nash

equilibrium and simultaneously compensating for the effect of the transport–heat PDEs in our heterogeneous noncooperative game.

The evolutions of the infinite-dimensional states $\alpha_1(x, t)$ and $\alpha_2(x, t)$ modeled by the transport–heat PDEs (84)–(87) and (88)–(91) are shown in Figure 8(a)–(d). The values of the boundary inputs $\theta_1(t)$ and $\theta_2(t)$, as well as the boundary outputs $\Theta_1(t)$ and $\Theta_2(t)$, are highlighted in black and red, respectively. The initial condition is in blue.

This first set of simulations indicates that, even under an adversarial scenario of strong coupling between the players with $\epsilon = 1$, the proposed approach has behaved successfully. This suggests that our stability analysis may be conservative, and the theoretical assumption $0 < \epsilon < 1$ may be relaxed, given the performance of the closed-loop control system. In Figure 9(a)–(c), the different values of $\epsilon = 0.75$, $\epsilon = 0.5$, and $\epsilon = 0.25$ are considered to evaluate the robustness of the proposed scheme under different levels of coupling between the two players and the corresponding impact on the transient responses.

FROM NASH EQUILIBRIUM TO EXTREMUM SEEKING

We now present the scalar ES design for distinct families of PDEs. According to Remark 1, scalar ES is a particular case of NES with a single agent. See also “Fundamentals of ES.”

Remark 1

Note that the material in the previous sections addressed the following special cases:

- » multivariable static ES [53], [59], as corollaries of Theorems 1, 2 and 3, when $h_1 = h_N$ and $\epsilon = 1$,
- » scalar static ES, when simply $N = 1$

In this case, the ES goal is to optimize in real time an unknown static map $Q(\cdot)$:

$$y(t) = Q(\Theta(t)) \quad (153)$$

with maximum or minimum unknown output y^* and optimizer Θ^* , by measuring the output $y(t)$ and adjusting the input $\Theta(t)$. As illustrated in the block diagram of Figure 10, the input Θ of the map is governed by a PDE, which must be properly compensated for by means of an appropriate boundary control law.

For maximization problems, the unknown nonlinear map is assumed to be locally quadratic, such that

$$Q(\Theta) = y^* + \frac{H}{2}(\Theta - \Theta^*)^2 \quad (154)$$

where $\Theta^*, y^* \in \mathbb{R}$, and the Hessian $H < 0$ is also unknown (when $H > 0$, we have a minimization problem). Hence, the output of the static map is given by

$$y(t) = y^* + \frac{H}{2}(\Theta(t) - \Theta^*)^2. \quad (155)$$

Proof of Theorem 3

First, consider the equivalent hyperbolic/parabolic PDE–ODE representation (140)–(144) rewritten for each Player P_i , $i \in \{1, 2\}$:

$$\dot{\bar{G}}_i^{\text{av}}(t) = H_{ii}^i k_i \bar{G}_i^{\text{av}}(t) + \epsilon H_{ii}^i k_i \phi_i^{\text{av}}(1, t) \quad (\text{S125})$$

$$\partial_t u_1^{\text{av}}(x, t) = D_1^{-1} \partial_x u_1^{\text{av}}(x, t), \quad x \in (0, 1) \quad (\text{S126})$$

$$\partial_t u_2^{\text{av}}(x, t) = D_2^{-2} \partial_{xx} u_2^{\text{av}}(x, t), \quad x \in (0, 1) \quad (\text{S127})$$

$$\partial_x u_i^{\text{av}}(0, t) = 0 \quad (\text{S128})$$

$$u_i^{\text{av}}(1, t) = k_i \bar{G}_i^{\text{av}}(t) + \epsilon k_i \phi_i^{\text{av}}(1, t). \quad (\text{S129})$$

For Player P_1 , the average closed-loop system, given by (S125) and (S126) and (S128) and (S129), satisfies both assumptions (H1) and (H2) of the Small-Gain Theorem [32, Theorem 8.1, p. 198] for the hyperbolic PDE–ODE loop with $n = 1$, $x(t) = \bar{G}_1^{\text{av}}(t)$, $F(\bar{G}_1^{\text{av}}(t), u_1^{\text{av}}(x, t), v(t)) = H_{11}^1 k_1 \bar{G}_1^{\text{av}}(t) + \epsilon v(t)$ with $v(t) = H_{11}^1 k_1 \phi_1^{\text{av}}(1, t)$, $c = 1/D_1$, $a(x) = 0$, $g(x, \bar{G}_1^{\text{av}}(t), u_1^{\text{av}}(x, t)) = 0$, $f(x, t) = 0$, $\varphi(d(t), u_1^{\text{av}}(x, t), \bar{G}_1^{\text{av}}(t)) = k_1 \bar{G}_1^{\text{av}}(t) + \epsilon d(t)$, where $d(t) = k_1 \phi_1^{\text{av}}(1, t)$, $N = k_1$, $L = |H_{11}^1| k_1$, $B = 0$, $\gamma_2 = k_1$, $b_2 = \epsilon$, and $A = \gamma_1 = 0$. Notice that Assumption (H1) holds with $M = 1$, $\gamma_3 = 0$, $\sigma = |H_{11}^1| k_1$, and $b_3 > 0$ being an appropriate constant of order $O(\epsilon)$, as can be easily verified using the variation-of-constants formula (for $i = 1$)

$$\begin{aligned} \bar{G}_1^{\text{av}}(t) &= \exp(-|H_{ii}^i| k_i t) \bar{G}_1^{\text{av}}(0) \\ &+ \int_0^t \exp(-|H_{ii}^i| k_i(t+s)) \epsilon H_{ii}^i k_i \phi_i^{\text{av}}(1, s) ds \end{aligned} \quad (\text{S130})$$

and from the application of the Cauchy–Schwarz inequality to the term $\phi_1^{\text{av}}(1, t)$ in (139):

$$\begin{aligned} \phi_1^{\text{av}}(1, t) &\leq |H_{12}^1| D_2^2 \left(\int_0^1 (1-\tau)^2 d\tau \right)^{\frac{1}{2}} \times \left(\int_0^1 [u_2^{\text{av}}(\xi, t)]^2 d\xi \right)^{\frac{1}{2}} \\ &\leq \frac{1}{\sqrt{3}} k_H D_2^2 \left(\int_0^1 [u_2^{\text{av}}(\xi, t)]^2 d\xi \right)^{\frac{1}{2}} \end{aligned} \quad (\text{S131})$$

since $|H_{ij}^j| < k_H < (1/\epsilon)|H_{ii}^i|$ according to Assumption 3, where k_H is a positive constant of order $O(1)$. Therefore, if $0 < \epsilon < 1$ is sufficiently small, then [32, Theorem 8.1, p. 198] allows us to conclude that there exist constants $\delta_1, \Delta_1 > 0$ such that, for every $u_{1,0}^{\text{av}} \in C^0([0, 1])$, $\bar{G}_{1,0}^{\text{av}} \in \mathbb{R}$, the unique generalized solution of this initial-boundary value problem, with $u_1^{\text{av}}(x, 0) = u_{1,0}^{\text{av}}$ and $\bar{G}_1^{\text{av}}(0) = \bar{G}_{1,0}^{\text{av}}$, satisfies the following estimate, $\forall t \geq 0$:

$$\begin{aligned} |\bar{G}_1^{\text{av}}(t)| + \|u_1^{\text{av}}(t)\|_{\infty} &\leq \Delta_1 \left(|\bar{G}_{1,0}^{\text{av}}| + \|u_{1,0}^{\text{av}}\|_{\infty} \right) \exp(-\delta_1 t) + \\ &+ \bar{\gamma}_1 \epsilon \max_{0 \leq s \leq t} \left(\|u_2^{\text{av}}(s)\|_{\infty} \right) \end{aligned} \quad (\text{S132})$$

for some adequate constant $\bar{\gamma}_1 > \max(|H_{11}^1| k_1, k_1)$.

For Player P_2 , the average closed-loop system (S125) and (S127)–(S129) satisfies all of the assumptions (A1) to (A7) of the Small-Gain Theorem [32, Theorem 8.2, p. 205] for the parabolic PDE–ODE loop with $n = 1$, $p(x) = 1$,

$r(x) = D_2^2$, $q(x) = 0$, $F(\bar{G}_2^{\text{av}}(t), u_2^{\text{av}}(x, t), v(t)) = H_{22}^2 k_2 \bar{G}_2^{\text{av}}(t) + \epsilon v(t)$, $v(t) = H_{22}^2 k_2 \phi_2^{\text{av}}(1, t)$, $g(x, \bar{G}_2^{\text{av}}(t), u_2^{\text{av}}(x, t)) = 0$, $f(x, t) = 0$, $\varphi_0(d(t), u_2^{\text{av}}(x, t), \bar{G}_2^{\text{av}}(t)) = b_1 u_1^{\text{av}}(0, t)$, $b_1 < 0$, $b_2 = 1$, $\varphi_1(d(t), u_2^{\text{av}}(x, t), \bar{G}_2^{\text{av}}(t)) = k_2 \bar{G}_2^{\text{av}}(t) + \epsilon d(t)$, $d(t) = k_2 \phi_2^{\text{av}}(1, t)$, $a_1 = 1$, $a_2 = 0$, $L = |H_{22}^2| k_2$, $K_0 = |b_1|$, $B_0 = C_0 = 0$, γ_0 of order $O(1)$, $K_1 = 0$, $B_1 = k_2$, $C_1 = \epsilon$, γ_1 of order $O(1)$, and $K_2 = B_2 = 0$. Assumption (A6) of [32, Theorem 8.2, p. 205] holds with $M = 1$, $\gamma_3 = 0$, $\sigma = |H_{22}^2| k_2$, and $b_3 > 0$ being of order $O(\epsilon)$, as can be readily verified by means of the variation-of-constants formula (S130), with $i = 2$, and from the application of the Cauchy–Schwarz inequality to the term $\phi_2^{\text{av}}(1, t)$ in (139):

$$\begin{aligned} \phi_2^{\text{av}}(1, t) &\leq |H_{21}^2| \left(\int_0^1 D_1^2 d\tau \right)^{\frac{1}{2}} \times \left(\int_0^1 [u_1^{\text{av}}(\xi, t)]^2 d\xi \right)^{\frac{1}{2}} \\ &\leq k_H D_1 \left(\int_0^1 [u_1^{\text{av}}(\xi, t)]^2 d\xi \right)^{\frac{1}{2}} \end{aligned} \quad (\text{S133})$$

with the same $k_H > 0$ defined just after (S131). Hence, it follows that the small-gain condition in [32, Inequality (8.3.24)] holds, provided $0 < \epsilon < 1$ is sufficiently small. Thus, [32, Theorem 8.2, p. 205] allows us to conclude that there exist constants $\delta_2, \Delta_2 > 0$ such that

$$\begin{aligned} |\bar{G}_2^{\text{av}}(t)| + \|u_2^{\text{av}}(t)\|_{\infty} &\leq \Delta_2 \left(|\bar{G}_{2,0}^{\text{av}}| + \|u_{2,0}^{\text{av}}\|_{\infty} \right) \exp(-\delta_2 t) + \\ &+ \bar{\gamma}_2 \epsilon \max_{0 \leq s \leq t} \left(\|u_1^{\text{av}}(s)\|_{\infty} \right) \end{aligned} \quad (\text{S134})$$

$\forall t \geq 0$, for some adequate constant $\bar{\gamma}_2 > \max(|H_{22}^2| k_2, k_2)$, $u_2^{\text{av}}(x, 0) = u_{2,0}^{\text{av}}$, and $\bar{G}_2^{\text{av}}(0) = \bar{G}_{2,0}^{\text{av}}$. Since the inequalities (S132) and (S134) are similar to those found in [32, Theorem 11.2, p. 269]—see inequalities (11.2.23) and (11.2.24)—we can finally invoke [32, Theorem 11.5, p. 277], under the condition of $0 < \epsilon < 1$ sufficiently small, to conclude that

$$\begin{aligned} |\bar{G}_1^{\text{av}}(t)| + |\bar{G}_2^{\text{av}}(t)| + \|u_1^{\text{av}}(t)\|_{\infty} + \|u_2^{\text{av}}(t)\|_{\infty} &\leq \\ \Delta \left(|\bar{G}_{1,0}^{\text{av}}| + |\bar{G}_{2,0}^{\text{av}}| + \|u_{1,0}^{\text{av}}\|_{\infty} + \|u_{2,0}^{\text{av}}\|_{\infty} \right) \exp(-\delta t) \end{aligned} \quad (\text{S135})$$

$\forall t \geq 0$, for some $\delta > 0$ and $\Delta > 0$.

Therefore, we conclude that the origin of the average closed-loop system (140)–(144) is exponentially stable under the assumption of $0 < \epsilon < 1$ being sufficiently small. Then, from (120), (133), and (134), we conclude the same results in the norm

$$\left(\sum_{i=1}^2 [\vartheta_i^{\text{av}}(t)]^2 + \int_0^D [u_i^{\text{av}}(x, t)]^2 dx \right)^{1/2}. \quad (\text{S136})$$

After applying the averaging theory for infinite-dimensional systems in [25, Section 2] (see “Averaging Theorem for General Infinite-Dimensional Systems”), we can show that (146) and (147) are indeed satisfied. The final residual sets for the errors $\theta_i(t) - \theta_i^*$ in (148) and (149) depend on a_1 and $a_2 e^{D_2 \sqrt{\omega/2}}$ due to the amplitude of the additive dithers $S_i(t)$ in (109), for $i \in \{1, 2\}$.

For the sake of simplicity, we have considered scalar maps with one single input and one single output, but a more general setup of multiple-input, single-output maps ($\Theta \in \mathbb{R}^N$ and $y \in \mathbb{R}$) could also be addressed, as done in [53] and [59].

While the multiplicative perturbation signals $M(t) = (2/a)\sin(\omega t)$ and $N(t) = -(8/a^2)\cos(\omega t)$ follow classical ES

designs [23], [43], the additive dither $S(t)$ must be redesigned using the trajectory generation paradigm [42, Chapter 12]. We then show the ES design for the simplest cases of hyperbolic and parabolic PDEs: transport (delay) and heat (diffusion) PDEs. While in the case of a delay, it suffices to advance the sinusoidal perturbation in time, in the case of

Derivation of Computational Simulation Models for PDEs

The boundary control laws in this article are implemented by using a modal transformation and PDEs by the finite difference method [S10]. In the following, the simulation model of the diffusion PDE and the corresponding controller are specifically derived.

SIMULATION OF DIFFUSION PDE VIA FINITE DIFFERENCES

Consider the diffusion process (28)–(31), which we note here as a reminder

$$\Theta(t) = \alpha(0, t) \quad (\text{S137})$$

$$\alpha_t(x, t) = \alpha_{xx}(x, t), \quad x \in (0, D) \quad (\text{S138})$$

$$\alpha_x(0, t) = 0 \quad (\text{S139})$$

$$\alpha(D, t) = \theta(t). \quad (\text{S140})$$

First, we normalize the spatial coordinate and the associated PDE state, such that

$$\xi = \frac{x}{D}, \quad \tilde{\alpha}(\xi, t) = \alpha(x, t). \quad (\text{S141})$$

Then, we have the following relation:

$$\alpha_x(x, t) = \frac{1}{D}\tilde{\alpha}_\xi(\xi, t), \quad \alpha_{xx}(x, t) = \frac{1}{D^2}\tilde{\alpha}_{\xi\xi}(\xi, t). \quad (\text{S142})$$

Thus, the scaled diffusion process is

$$\tilde{\alpha}_t(\xi, t) = \frac{1}{D^2}\tilde{\alpha}_{\xi\xi}(\xi, t), \quad \xi \in (0, 1) \quad (\text{S143})$$

$$\tilde{\alpha}_\xi(0, t) = 0 \quad (\text{S144})$$

$$\tilde{\alpha}(1, t) = \theta(t). \quad (\text{S145})$$

To simulate the normalized diffusion process (S143)–(S145), we are using the well-known finite difference method [S10]. We discretize the spatial space in $N + 1$ elements with width $h = 1/N$ such that

$$\tilde{\alpha}_i(t) = \tilde{\alpha}(ih, t), \quad \text{for } i \in [0, N]. \quad (\text{S146})$$

The spatial derivative terms are calculated by the (central) difference quotients

$$\tilde{\alpha}_\xi(ih, t) = \frac{\alpha_{i+1}(t) + \alpha_i(t)}{h} \quad (\text{S147})$$

$$\tilde{\alpha}_{\xi\xi}(ih, t) = \frac{\alpha_{i+1}(t) - 2\alpha_i(t) + \alpha_{i-1}(t)}{h^2}. \quad (\text{S148})$$

Hence, the PDE system (S137)–(S140) can be written as N ODEs and three algebraic equations derived from the boundary conditions

$$\Theta(t) = \tilde{\alpha}_0(t) \quad (\text{S149})$$

$$\tilde{\alpha}_0(t) = \tilde{\alpha}_1(t) \quad (\text{S150})$$

$$\dot{\tilde{\alpha}}_i(t) = \frac{1}{D^2} \frac{\tilde{\alpha}_{i+1}(t) - 2\tilde{\alpha}_i(t) + \tilde{\alpha}_{i-1}(t)}{h^2}, \quad i \in [1, N-1] \quad (\text{S151})$$

$$\dot{\tilde{\alpha}}_N = \dot{\theta}(t) \quad (\text{S152})$$

$$\tilde{\alpha}_N(t) = \theta(t) \quad (\text{S153})$$

where (S150) results from the Neumann boundary condition.

SIMULATION OF THE CONTROLLER VIA MODAL TRANSFORMATION

In the modal transformation, we rewrite the infinite-dimensional closed-loop system as

$$\dot{z}(t) = \Gamma z(t) + J(\omega t, z(t)) \quad (\text{S154})$$

where Γ generates an analytic semigroup, $J(\omega t, z(t))$ satisfies the smoothness properties, and $z(t) = [\vartheta(t) V(t) U(t)]^\top$ is an infinite-dimensional state vector. Now, we only consider the states $V(t)$ and $U(t)$. Since $V(t)$ is the infinite-dimensional part of the state vector, we approximate this with a finite-dimensional state vector $\tilde{V}(t)$, by considering only the first L states/modes $\tilde{V}(t) = [v_1^*, v_2^*, \dots, v_L^*]$. Hence, we arrive at the simulation model for the controller

$$\begin{bmatrix} \dot{\tilde{V}}(t) \\ \dot{U}(t) \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 & \dots & -c\beta_1^* \\ 0 & \ddots & 0 & \vdots \\ \vdots & 0 & \lambda_L & -c\beta_L^* \\ 0 & \dots & 0 & -c \end{bmatrix} \begin{bmatrix} \tilde{V}(t) \\ U(t) \end{bmatrix} + \begin{bmatrix} -c\beta_1 K \left[G(t) + \hat{H}(t) \left(\frac{1}{2} D^2 U(t) + \sum_{k=1}^L \alpha_k v_k^*(t) \right) \right] \\ \vdots \\ -c\beta_L K \left[G(t) + \hat{H}(t) \left(\frac{1}{2} D^2 U(t) + \sum_{k=1}^L \alpha_k v_k^*(t) \right) \right] \\ -cK \left[G(t) + \hat{H}(t) \left(\frac{1}{2} D^2 U(t) + \sum_{k=1}^L \alpha_k v_k^*(t) \right) \right] \end{bmatrix} \quad (\text{S155})$$

with the eigenvalues

$$\lambda_k = -\frac{\pi^2}{4D^2}(2k-1)^2, \quad k = 1, 2, \dots, L \quad (\text{S156})$$

$\alpha_k = \sqrt{2/D}(4D^2/\pi^2(2k-1)^2)$, $\beta_k = -\sqrt{2/D}(2D/\pi(2k-1))(-1)^k$, and $k = 1, 2, \dots$. The ODE system (S155) with $L + 1$ states can be simply implemented in Matlab/Simulink.

REFERENCE

[S10] G. D. Smith, *Numerical Solution of Partial Differential Equations: Finite Difference Methods*. Oxford, U.K.: Oxford Univ. Press, 1985.





FIGURE 7 (a) Action time histories and (b) payoff time histories for P_1 and P_2 , for $\epsilon = 1$. The dashed lines denote the values at the Nash equilibrium, $\Theta_1^* = 43.33$ and $\Theta_2^* = 36.67$ (with $J_1(\Theta^*) = 889$ and $J_2(\Theta^*) = 2,722$).







FIGURE 8 Evolution of the infinite-dimensional states $\alpha_1(x, t)$ and $\alpha_2(x, t)$ of the transport–heat PDEs in a heterogeneous duopoly game with boundary Dirichlet actuation, according to (147)–(149): from $\alpha_1(D_1, t) = \theta_1(t)$ to $\alpha_1(0, t) = \Theta_1(t)$, with $D_1 = 30$ for Player P_1 and from $\alpha_2(D_2, t) = \theta_2(t)$ to $\alpha_2(0, t) = \Theta_2(t)$, with $D_2 = 3$ for Player P_2 . (a) Parameter $\theta_1(t)$ (red) converges to a $O(|a| + 1/\omega)$ -neighborhood of Θ_1^* (dashed green). (b) Parameter $\theta_1(t)$ (black) converges to a $O(a_1 + 1/\omega)$ -neighborhood of θ_1^* (dashed green). (c) Parameter $\theta_2(t)$ (red) converges to a $O(|a| + 1/\omega)$ -neighborhood of Θ_2^* (dashed green). (d) Parameter $\theta_2(t)$ (black) converges to a $O(a_2 e^{D_2 \sqrt{\omega/2}} + 1/\omega)$ -neighborhood of θ_2^* (dashed green).

the heat dynamics, one has to employ a solution to the motion-planning problem, where the output of the heat PDE system is a sinusoid at one of its boundaries, and the input is the signal that must be applied on the other boundary to generate a sinusoid at the output. This input signal happens to consist of sinusoidal and exponential functions.

ES for Transport Hyperbolic PDE

In this case, the following infinite-dimensional and averaging-based predictor feedback is introduced to compensate for the delay [40]:

$$U(t) = \frac{c}{s+c} \left\{ k \left[G(t) + \hat{H}(t) \int_{t-D}^t U(\tau) d\tau \right] \right\} \quad (156)$$




FIGURE 9 Action time histories for P_1 and P_2 for distinct values of the coupling coefficient. (a) $\epsilon = 0.75$. The dashed lines denote $\Theta_1 = 35.64$ and $\Theta_2 = 28.36$. (b) $\epsilon = 0.5$. The dashed lines denote $\Theta_1 = 30.67$ and $\Theta_2 = 22.67$. (c) $\epsilon = 0.25$. The dashed lines denote $\Theta_1 = 27.3$ and $\Theta_2 = 18.41$.

where $k > 0$ and $c > 0$ are sufficiently large; that is, the predictor feedback is of the form of a low-pass filtered nonaveraged version of

$$U_{av}(t) = kG_{av}(t+D) = k \left[G_{av}(t) + H \int_{t-D}^t U_{av}(\sigma) d\sigma \right]. \quad (157)$$

The signal

$$\hat{H}(t) = N(t)y(t) \quad (158)$$

in (156) is used to obtain an estimate of the unknown Hessian H , and the additive dither is given by

$$S(t) = a \sin(\omega(t+D)). \quad (159)$$

ES for Diffusion-Dominated Parabolic PDE

In the case of a diffusion process in the actuation dynamics, the trajectory generation problem to be solved for motion-planning design [47, Chapter 12] is

$$S(t) := \beta(D, t) \quad (160)$$

$$\partial_t \beta(x, t) = \partial_{xx} \beta(x, t), \quad x \in (0, D) \quad (161)$$

$$\partial_x \beta(0, t) = 0 \quad (162)$$

$$\beta(0, t) = a \sin(\omega t) \quad (163)$$

where $\beta: [0, D] \times \mathbb{R}_+ \rightarrow \mathbb{R}$. The explicit solution of (160) is given by

$$S(t) = \frac{1}{2} a e^{\sqrt{\frac{\omega}{2} D}} \sin\left(\omega t + \sqrt{\frac{\omega}{2} D}\right) + \frac{1}{2} a e^{-\sqrt{\frac{\omega}{2} D}} \sin\left(\omega t - \sqrt{\frac{\omega}{2} D}\right). \quad (164)$$

On the other hand, we write the average-based infinite-dimensional control law to compensate for the diffusion process by

$$U(t) = \frac{c}{s+c} \left\{ K \left[G(t) + \hat{H}(t) \int_0^D (D-r) u(r, t) dr \right] \right\} \quad (165)$$

where $c > 0$ is sufficiently large.

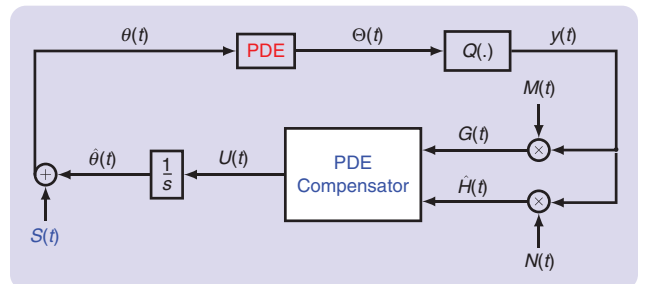


FIGURE 10 A general block diagram for implementation of ES control design for PDE connections (in red) at the input of nonlinear convex maps $Q(\cdot)$. Although the multiplicative perturbation signals $M(t)$ and $N(t)$ are the same as in the classical ES designs [23], [43], the additive dither $S(t)$ (in blue) must be redesigned using the trajectory generation paradigm [42, Chapter 12], and the application of an adequate boundary control law (in blue) for PDE compensation is necessary.

ES for Distinct Families of PDEs

To show that the proposed ES approach for infinite-dimensional systems is general and applicable to a wider class of PDEs, we have formulated in Table 1 the stabilizing boundary control law $U(t)$ and given the explicit solutions to the trajectory generation problem of $S(t)$ for five other classes of

distributed parameter systems [42]: 1) reaction–advection–diffusion (RAD) PDEs [52], [54], 2) wave equations [55], [56], 3) hyperbolic transport PDEs—for constant delays [59], 4) time-varying delays [69], and 5) distributed delays [76], [77].

Regarding the transport and diffusion cases explored in the NES sections, the term $u(x, t)$ that appears in $U(t)$

Fundamentals of ES

ES is a method for real-time nonmodel-based optimization. Though ES was introduced more than a century ago, in 1922, the “turn of the 21st century” has been its golden age, both in terms of the development of theory and in terms of its adoption in industry and in fields outside of control engineering. We provide an overview of the basic gradient-based version of ES with periodic signals for static maps (free of infinite-dimensional dynamics).

Many versions of ES exist, with various approaches to their study of stability [43], [S11], [S6]. The most common version employs perturbation signals for the purpose of estimating the gradient of the unknown map that is being optimized. To understand the basic idea of ES, it is best to first consider the case of a static single-input map of the quadratic form

$$f(\theta) = f^* + \frac{f''}{2}(\theta - \theta^*)^2 \quad (\text{S157})$$

where f^* , f'' and θ^* are all unknown, as shown in Figure S5.

Three different θ 's appear in Figure S5: θ^* is the unknown optimizer of the map, $\hat{\theta}(t)$ is the real-time estimate of θ^* , and $\theta(t)$ is the actual input into the map. The actual input $\theta(t)$ is based on the estimate $\hat{\theta}(t)$ but is perturbed by the signal $a \sin(\omega t)$ for the purpose of estimating the unknown gradient $f'' \cdot (\theta - \theta^*)$ of the map $f(\theta)$ in (S157). The sinusoid is only one choice for a perturbation signal—many other perturbations, from square waves to stochastic noise, can be used in lieu of sinusoids, provided that they are of zero mean. The estimate




FIGURE S5 The simplest perturbation-based ES scheme for a quadratic single-input map $f(\theta)$ in (S1). The user only needs to know the sign of f'' , namely, whether the quadratic map has a maximum or a minimum, and has to choose the adaptation gain k such that $\text{sgn}(k) = -\text{sgn}(f'')$. The user has to also choose the frequency ω as relatively large compared to a , k , and f'' .

$\hat{\theta}(t)$ is generated with the integrator k/s with the adaptation gain k controlling the speed of estimation.

The ES algorithm is successful if the error between the estimate $\hat{\theta}(t)$ and the unknown θ^* , namely the signal

$$\tilde{\theta}(t) = \hat{\theta}(t) - \theta^* \quad (\text{S158})$$

converges toward zero or some small neighborhood of zero as $t \rightarrow +\infty$. Based on Figure S5, the estimate $\hat{\theta}(t)$ is governed by the differential equation $\dot{\hat{\theta}} = k \sin(\omega t) f'(\theta)$, which means that the estimation error is governed by

$$\frac{d\tilde{\theta}}{dt} = k a \sin(\omega t) \left[f' + \frac{f''}{2}(\tilde{\theta} + a \sin(\omega t))^2 \right]. \quad (\text{S159})$$

Expanding the right-hand side, one obtains

$$\begin{aligned} \frac{d\tilde{\theta}(t)}{dt} &= k a f' \underbrace{\sin(\omega t)}_{\text{mean} = 0} + k a^3 \frac{f''}{2} \underbrace{\sin^3(\omega t)}_{\text{mean} = 0} \\ &\quad + k a \frac{f''}{2} \underbrace{\sin(\omega t)}_{\text{fast, mean} = 0} \underbrace{\tilde{\theta}(t)^2}_{\text{slow}} \\ &\quad + k a^2 f'' \underbrace{\sin^2(\omega t)}_{\text{fast, mean} = 1/2} \underbrace{\tilde{\theta}(t)}_{\text{slow}}. \end{aligned} \quad (\text{S160})$$

A theoretically rigorous time-averaging procedure [S1, Section 10.4] allows one to replace the previous sinusoidal signals by their means, yielding the “average system” [S1, p. 404]:

$$\frac{d\tilde{\theta}_{\text{av}}}{dt} = \frac{\overline{k f'' a^2}}{2} \tilde{\theta}_{\text{av}} \quad (\text{S161})$$

which is exponentially stable. The averaging theory guarantees that there exists a sufficiently large ω such that, if the initial estimate $\hat{\theta}(0)$ is sufficiently close to the unknown θ^* , one has

$$|\theta(t) - \theta^*| \leq |\theta(0) - \theta^*| e^{-\frac{k f'' a^2}{2} t} + O\left(\frac{1}{\omega}\right) + O(a), \quad \forall t \geq 0. \quad (\text{S162})$$

For the user, the inequality (S162) guarantees that, if a is chosen small and ω is chosen large, the input $\theta(t)$ exponentially converges to a small interval—of order $O((1/\omega) + a)$ —around the unknown θ^* and, consequently, the output $f(\theta(t))$ converges to the vicinity of the optimal output f^* .

REFERENCE

[S11] S.-J. Liu and M. Krstić, *Stochastic Averaging and Stochastic Extremum Seeking*. London, U.K.: Springer-Verlag, 2012.

of Table 1 is the state of the infinite-dimensional system corresponding to a copy of the PDE model of actuator dynamics.

In [57], differently from what has been done that has dealt with PDEs at the input into an unknown map, and in which we have already advanced from transport PDEs to RAD PDEs to wave PDEs according to Table 1, the authors consider one last configuration in which the input pathway to the map contains a cascade of PDEs from distinct classes. There, we deal with PDEs with input delays, such as the notorious problem of a wave PDE with an input delay, where, if the delay is left uncompensated, an

arbitrarily short delay destroys the closed-loop stability. Then, we move forward to an even more challenging class of problems for parabolic–hyperbolic cascades of PDEs, focusing on a heat equation at the input of a wave PDE. The treatment of such systems with PDE–PDE cascades is again performed by means of boundary control. Local exponential stability and convergence to a small neighborhood of the unknown extremum point are guaranteed by using a backstepping transformation and averaging in infinite dimensions.

PDE–PDE cascades have a great deal in common with PDE–ODE cascades. For instance, a cascade of a delay into a

TABLE 1 ES for distinct classes of PDE systems.

RAD equation	<p>PDE: $\partial_t \alpha(x, t) = \epsilon \partial_{xx} \alpha(x, t) + b \partial_x \alpha(x, t) + \lambda \alpha(x, t), \quad x \in [0, 1]$</p> <p>BC (Dirichlet): $U(t) = \frac{c}{s+c} \left\{ k e^{-\frac{b}{2\epsilon} t} \left[\gamma(1) G(t) + \hat{H}(t) \int_0^1 e^{\frac{b}{2\epsilon} \sigma} m(1-\sigma) u(\sigma, t) d\sigma \right] \right\}$</p> <p>$\gamma(x) = \cosh\left(\sqrt{\frac{\xi}{\epsilon}} x\right) + \frac{b}{2\epsilon} \sqrt{\frac{1}{\xi}} \sinh\left(\sqrt{\frac{\xi}{\epsilon}} x\right), \quad \xi := b^2/(4\epsilon) - \lambda \geq 0, k > 0$</p> <p>$m(x-\sigma) = \frac{1}{\epsilon} \sqrt{\frac{\epsilon}{\xi}} \sinh\left(\sqrt{\frac{\xi}{\epsilon}} (x-\sigma)\right), \quad \epsilon > 0, b \geq 0, \lambda \geq 0$</p> <p>Trajectory generation: $S(t) = e^{-\frac{b}{2\epsilon} t} \sum_{k=0}^{\infty} \frac{a_{2k}(t)}{(2k)!} + \frac{b}{2\epsilon} \frac{a_{2k}(t)}{(2k+1)!}$</p> <p>$a_{2k} := \frac{a}{\epsilon^k} \sin(\omega t) \sum_{n=0}^k \binom{k}{2n} \xi^{k-2n} \omega^{2n} + \frac{a}{\epsilon^k} \cos(\omega t) \sum_{n=0}^k \binom{k}{2n+1} \xi^{k-2n-1} \omega^{2n+1}$</p>
Wave dynamics	<p>PDE: $\partial_{tt} \alpha(x, t) = \partial_{xx} \alpha(x, t), \quad x \in [0, D]$</p> <p>BC (Neumann): $U(t) = \frac{c}{s+c} \left\{ c [k \hat{H}(t) u(D, t) - \partial_t u(D, t)] + \rho(D) G(t) + \hat{H}(t) \int_0^D \rho(D-\sigma) \partial_t u(\sigma, t) d\sigma \right\}, \quad \rho(s) = k [0 \quad \Gamma] e^{As} [0 \quad \Gamma]^T, \quad A = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$</p> <p>Trajectory generation: $S(t) = a \cos(\omega D) \sin(\omega t)$</p>
Constant delay	<p>PDE: $\partial_t \alpha(x, t) = \partial_x \alpha(x, t), \quad x \in [0, D]$</p> <p>BC (Dirichlet): $U(t) = \frac{c}{s+c} \left\{ k \left[G(t) + \hat{H}(t) \int_0^D u(\sigma, t) d\sigma \right] \right\}$</p> <p>Trajectory generation: $S(t) = a \sin(\omega(t+D))$</p>
Variable delay	<p>PDE: $\partial_t \alpha(x, t) = \pi(x, t) \partial_x \alpha(x, t), \quad x \in [0, 1], \quad \pi(x, t) = \frac{1+x \left[\frac{d(\phi^{-1}(t))}{dt} - 1 \right]}{\phi^{-1}(t) - t}$</p> <p>BC (Dirichlet): $U(t) = \frac{c}{s+c} \left\{ k \left[G(t) + \hat{H}(t) \int_0^1 u(\sigma, t) (\phi^{-1}(t) - t) d\sigma \right] \right\}$</p> <p>Trajectory generation: $S(t) = a \sin(\omega t), \quad \phi(t) := t - D(t)$</p> <p>Demodulation: $M(t) = \frac{2}{a} \sin(\omega(t - D(t))), \quad N(t) = -\frac{8}{a^2} \cos(2\omega(t - D(t)))$</p>
Distributed delay	<p>PDE: $\partial_t \alpha(x, t) = \partial_x \alpha(x, t), \quad x \in [0, D], \quad y = Q \left(\int_0^D \Theta(t-\sigma) d\beta(\sigma) \right)$</p> <p>BC (Dirichlet): $U(t) = \frac{c}{s+c} \left\{ k \left[G(t) + \hat{H}(t) \int_0^D (1-\beta(\sigma)) u(D-\sigma, t) d\sigma \right] \right\}$</p> <p>Trajectory generation: $S(t) = \frac{a}{\gamma(\omega)} \int_0^D \sin(\omega(t+\xi)) d\beta(\xi)$</p>

BC: boundary control.

We show how to properly sequence the steps of averaging, backstepping, and Lyapunov functional analysis to prove stability.

PDE is a much more generalized version of an integrator with an input delay. However, in a delay–integrator cascade, the design can be pursued within the predictor feedback framework, with backstepping just employed for an interpretation and for analysis; whereas, in delay–PDE cascades, a predictor for a PDE is too complicated a mathematical object to be of value. Instead, the design is pursued entirely by the backstepping approach. Similarly, while the heat–integrator cascade became familiar and was dealt with through the backstepping design, a more general backstepping design is applied to a heat–wave PDE cascade in a part of the reference [57].

Although our goal in this article is to avoid numbing the reader with lengthy proofs, it is worth providing a general picture of how we carry out the steps to prove the stability results of the ES feedback loop in the presence of PDEs in Figure 10. Figure 11 shows the structure of the proof, divided into six main steps. We take advantage of this opportunity to highlight that our analysis presents a carefully constructed sequence of analytical steps, a predictor-based infinite-dimensional backstepping transformation, a synthesis of a Lyapunov functional (rather than small-gain analysis), and the computation of a Lyapunov estimate, for the overall infinite-dimensional system with nonlinearities, stochastic perturbations, and distributed delays. This analysis process, involving so many steps, has a large number of possible permutations—all *but one* would be wrong. We show how to properly sequence the steps of averaging, backstepping, and Lyapunov functional

analysis to prove stability. This “analysis pathway” will serve the needs of future researchers who deal with stochastic ES under delays. The complete details can be found in the authors’ book [66].

APPLICATIONS

For the distinct cases involving ES plus PDEs in Table 1, we dedicated the remainder of the article to some select applications:

- 1) *traffic control* for linearized Lighthill–Whitham–Richards (LWR) macroscopic PDE models transformed into constant delays [59], [84]
- 2) *optimal oil-drilling control* with ES for wave models [1]
- 3) *deep-sea cable-actuated source seeking* modeled by wave PDEs with Kelvin–Voigt damping [72]
- 4) *additive manufacturing* modeled by the *Stefan PDE* [22], [36]
- 5) *bioreactors* considering ES for models described by parabolic PDEs (reaction–diffusion equations), see [14] and [30]
- 6) *light-source seeking* with infinite-dimensional models represented by Euler–Bernoulli beam equations [42, Chapter 8]
- 7) *NMES problem* for ES with time-varying delays [65], [69].

We offer a tutorial-style guide that synthesizes existing results for ES while presenting new insights and practical examples in engineering applications.

TRAFFIC CONGESTION CONTROL WITH A DOWNSTREAM BOTTLENECK

This section develops boundary control for freeway traffic with a downstream bottleneck [84]. Traffic on a freeway segment with a capacity drop at the outlet of the segment is a common phenomenon that leads to a traffic bottleneck problem. The capacity drop can be caused by lane drops, hills, tunnels, bridges, or curvature on the road. If the incoming traffic flow remains unchanged, traffic congestion forms upstream of the bottleneck because the upstream traffic demand exceeds its capacity. Therefore, it is important to regulate the flow of traffic in the segment to avoid

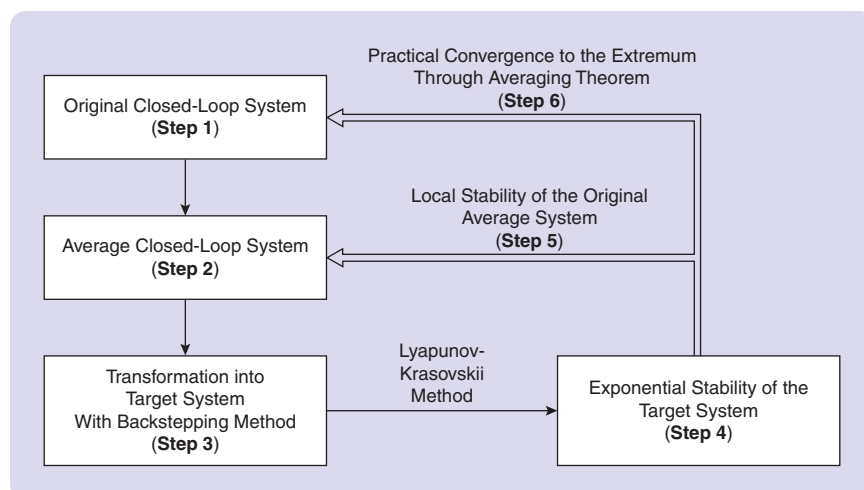


FIGURE 11 Structure of the stability proof for the closed-loop system. We provide theoretical stability guarantees for PDE systems in real-time optimization.

We offer a tutorial-style guide that synthesizes existing results for ES while presenting new insights and practical examples in engineering applications.

overloading the bottleneck area. Traffic densities on the freeway segment are described with the LWR macroscopic PDE model. To mitigate traffic congestion upstream of the bottleneck, the incoming flow at the inlet of the freeway segment is controlled so that an optimal density is reached that maximizes the outgoing flow. The density and traffic flow relation at the bottleneck area, described with a fundamental diagram, is considered to be unknown. We tackle this problem using ES control with delay compensation for LWR PDEs [84]. ES control, a nonmodel-based approach for real-time optimization, is adopted to find the optimal density for the unknown fundamental diagram. A predictor feedback control design is proposed to compensate for the delay effect of traffic dynamics in the freeway segment.

Problem Statement

We consider a traffic congestion problem on a freeway segment with a lane-drop bottleneck downstream of the segment. The freeway segment upstream of the bottleneck and the lane-drop area are shown in Figure 12, which illustrates the clear “Zone C” and the bottleneck “Zone B.” The flow is conserved through the clear Zone C to the bottleneck Zone B. Local road capacity changes are due to the lane drop in Zone B, which could be caused by working zones, accidents, or lane closures. To prevent the traffic in Zone B from overflowing its capacity and then causing congestion in the freeway segment, we aim to find the optimal density ahead of Zone C that maximizes the outgoing flux of Zone B, given the unknown density–flow relation. The traffic dynamics in Zone C is described with the macroscopic LWR traffic model for the aggregated values of traffic density.

Because of the reduction in lanes in Zone B, the fundamental diagram of the flow and density relation usually changes, leading to a capacity drop in Zone B. The control objective is to find the optimal input density at the inlet of Zone C that drives the measurable output flux of Zone B to its unknown optimal value in an unknown fundamental diagram.

LWR Traffic Model

The traffic dynamics in Zone C, upstream of Zone B, is described with a first-order hyperbolic LWR model. The traffic density $\rho(x, t)$ in Zone C is governed by the following nonlinear hyperbolic PDE, where $x \in [0, L]$, $t \in [0, \infty)$:

$$\partial_t \rho + \partial_x(Q_C(\rho)) = 0. \tag{166}$$

The fundamental diagram of the traffic flow and density function $Q_C(\rho)$ is given by $Q_C(\rho) = \rho V(\rho)$, where the traffic

velocity follows an equilibrium velocity–density relation $V(\rho)$. There are different models that describe the flux and density relationship. A basic and popular choice is Greenshield’s model for $V(\rho)$, which is given by $V(\rho) = v_f(1 - (\rho/\rho_m))$, where $v_f \in \mathbb{R}^+$ is defined as the maximum velocity and $\rho_m \in \mathbb{R}^+$ is the maximum density for Zone C [84]. Then the fundamental diagram of the flow and density function $Q_C(\rho)$ is in a quadratic form of density

$$Q_C(\rho) = -\frac{v_f}{\rho_m} \rho^2 + v_f \rho. \tag{167}$$

A critical value of density segregates traffic into a free-flow regime, whose density is smaller than the critical value, and a congested regime, whose density is greater than the critical value. The critical density may be assumed as $\rho_c = \rho_m/2$ for (167) [84]. For the fundamental diagram calibrated with the freeway empirical data, the critical density usually appears at 20% of the maximum value of the density [15], [18].

In practice, the quadratic fundamental diagram sometimes does not fit well with traffic density–flow field data. There are several other equilibrium models, for example, the Greenberg model, the Underwood model, and the diffusion model, for which the fundamental diagrams are nonlinear functions; see [84] and the references therein. However, by Taylor expansion, a second-order differentiable nonlinear function can be approximated by a quadratic function in the neighborhood of its extremum. The following assumption is made for the nonlinear fundamental diagram. The stability results derived here are valid locally for the general form of the fundamental diagram $Q(\rho)$ that satisfies the following assumption. Here we can adopt other density–flow relations for the fundamental diagram $Q(\rho)$, but we require Assumption 4 to be satisfied.

Assumption 4

The fundamental diagram $Q(\rho)$ is a smooth function, and it holds that $Q'(\rho_c) = 0$, $Q''(\rho_c) < 0$.




FIGURE 12 Traffic on a freeway segment with lane drop.

We provide theoretical stability guarantees for PDE systems in real-time optimization.

Under **Assumption 4**, the fundamental diagram can be approximated around the critical density ρ_c as follows: $Q(\rho) = q_c + (Q''(\rho)/2)(\rho(t) - \rho_c)^2$, where $q_c = Q(\rho_c)$ is defined as the road capacity or the maximum flow, with $Q'(\rho) < 0$.

Lane-Drop Bottleneck Control Problem

Because of the reduction in the number of lanes from Zone C to Zone B, we consider the equilibrium density–flow relation of Zone B as shown in **Figure 13**, as pointed out in [75]. There is a capacity drop ΔC of Q_B in Zone B compared to Q_C in Zone C after congestion has formed upstream of the lane-drop area. The capacity drop caused by a sudden lane drop is hard to measure in real time, and the traffic dynamics of Zone B is affected by the lane-changing and merging activities. Therefore, we assume that the fundamental diagram $Q_B(\rho)$ of Zone B is unknown.

In **Figure 13**, the capacity is

$$\Delta C = Q_C(\rho_c) - Q_B((1 + \delta)\rho^*) \quad (168)$$

$$q^* = Q_C(\rho^*) = Q_B((1 + \delta)\rho^*) \quad (169)$$

where ΔC is unknown. $\rho^* \in \mathbb{R}^+$ represents the optimal density that keeps Zone C in the free regime, while $(1 + \delta)\rho^*$ reaches the critical density of Zone B so that the discharging flow rate reaches its maximum value $q^* \in \mathbb{R}^+$. The ratio δ accounts for the density discontinuity before the outlet in Zone C and after the outlet in Zone B. We assume that ΔC and δ are unknown, and therefore the optimal density and flow rate (ρ^*, q^*) are unknown.

When there is a lane-drop bottleneck presenting downstream, the density at the outlet of Zone C is $\rho(L, t)$, governed by the PDE in (166) for $x \in [0, L]$, $t \in [0, \infty)$. The inlet boundary flow is $q_{in}(t) = Q_C(\rho(0, t))$. The output measurement of traffic flow in Zone B, $q_{out}(t)$, is given by $Q(\rho)$ with outlet density $\rho(L, t)$, $q_{out}(t) = Q(\rho(L, t))$, where the

function $Q(\rho)$ of the outlet boundary $x = L$ connecting Zone C and Zone B is defined as follows:

$$Q(\rho(L, t)) = \begin{cases} Q_C(\rho(L, t)), & \rho(L, t) < \rho^* \\ Q_C(\rho^*) = q^* = Q_B((1 + \delta)\rho^*), & \rho(L, t) = \rho^* \\ Q_B((1 + \delta)\rho(L, t)), & \rho(L, t) > \rho^* \end{cases} \quad (170)$$

so that the flow is conserved through the boundary, entering from Zone C to Zone B. Note that, when the optimal density ρ^* is reached, the flow rate at the outlet of Zone C and the input of Zone C reaches equilibrium and its maximum value q^* .

The control objective is to design the traffic flow input $q_{in}(t)$ so that the outgoing flow in the lane-drop area Zone B, $q_{out}(t)$, is maximized. We aim to find the optimal outlet density $\rho(L, t) = \rho^*$ that maximizes $q_{out}(t)$ of Zone B and then use the PDE that describes the dynamics of traffic in Zone C to obtain the desirable flow input $q_{in}(t)$ from the inlet of Zone C. Here we approximate $q_{out}(t)$ with a function that satisfies **Assumption 4**, and $q_{out}(t)$ can be written as

$$q_{out}(t) = q^* + \frac{H}{2}(\rho(L, t) - \rho^*)^2 \quad (171)$$

where $H < 0$ is the unknown Hessian of the approximated static map $q_{out}(t)$.

Note that we use a static fundamental diagram to model traffic in the bottleneck Zone B. Therefore, our model cannot capture the propagating traffic waves upstream from Zone B to Zone C when Zone B is very congested. Since this result is focused on maximizing the discharging flow rate at the bottleneck area, the ES control seeks the optimal traffic density value in its neighborhood. In bottleneck Zone B, the closer the outlet traffic density $\rho(L, t)$ is to the optimal value ρ^* , where $Q'(\rho) = 0$ is satisfied, the smaller is the propagating characteristic speed of the traffic waves $Q'(\rho)$. Therefore, the spillback traffic from Zone B to Zone C is negligible in our model.

To find the unknown optimal density at the bottleneck area, we design ES control for the unknown static map $Q(\rho)$ with actuation dynamics governed by a nonlinear hyperbolic PDE in (166). Next, we linearize the nonlinear PDE, and the traffic dynamics can be represented by the delay effect for the control input design.

Linearized Reference Error System

We linearize the nonlinear LWR model around a constant reference density $\rho_r \in \mathbb{R}^+$, which is assumed to be close to




FIGURE 13 Quadratic fundamental diagram for the clear Zone C and the bottleneck Zone B.

the optimal density ρ^* . Note that the reference density ρ_r is in the free regime of $Q(\rho)$ of Zone C and thus is smaller than the critical density ρ_c ; therefore, the following is satisfied: $\rho_r < \rho_c$. Define the reference error density as

$$\tilde{\rho}(x, t) = \rho(x, t) - \rho_r \quad (172)$$

and the reference flux q_r to be $q_r = Q(\rho_r) > 0$. By the governing equation (166) along with (167), the linearized reference error model is derived as

$$\partial_t \tilde{\rho}(x, t) + u \partial_x \tilde{\rho}(x, t) = 0 \quad (173)$$

$$\tilde{\rho}(0, t) = \rho(0, t) - \rho_r \quad (174)$$

where the constant transport speed u is given by $u = Q'(\rho)|_{\rho=\rho_r} = V(\rho_r) + \rho_r V'(\rho)|_{\rho=\rho_r}$. The equilibrium velocity–density relation $V(\rho)$ is a strictly decreasing function. The reference density ρ_r is in the left half plane of the fundamental diagram $Q_c(\rho)$, which yields the following inequality for the propagation speed: $u > 0$. We define the input density as $\varrho(t) = \rho(0, t)$ and the linearized input at the inlet to be

$$\tilde{\varrho}(t) = \varrho(t) - \rho_r. \quad (175)$$

The linearized error dynamics in (173) and (174) is a transport PDE with an explicit solution for $t > x/u$ and thus is represented with input density $\tilde{\rho}(x, t) = \tilde{\varrho}(t - (x/u))$. The density variation at the outlet is

$$\tilde{\rho}(L, t) = \tilde{\varrho}(t - D) \quad (176)$$

where the time delay is $D = L/u$. Therefore, the density at the outlet is given by a delayed input density variation and the reference

$$\rho(L, t) = \rho_r + \tilde{\rho}(L, t). \quad (177)$$

Finally, substituting (176) and (177) into the static map (171), we arrive at the following:

$$\begin{aligned} q_{\text{out}}(t) &= q^* + \frac{H}{2} (\tilde{\varrho}(t - D) + \rho_r - \rho^*)^2 \\ &= q^* + \frac{H}{2} (\varrho(t - D) - \rho^*)^2. \end{aligned} \quad (178)$$

The control objective is to regulate the input $q_{\text{in}}(t)$ so that $\varrho(t - D)$ reaches an unknown optimal ρ^* , and the maximum of the uncertain quadratic flux–density map $q_{\text{out}}(t)$ can be achieved. We can apply the method of ES for static maps with delays, originally developed in [59]. The ES control is designed to find the extremum of the unknown map.

In practice, control of the density at the inlet can be realized with the coordinated operation of a ramp

metering and a variable speed limit (VSL) at the inlet, which is widely used in freeway traffic management [11], [27], [33], [47], [83], [85], [86]. The controlled density at the inlet is implemented by $\varrho(t) = q_{\text{in}}(t)/v_c$, where v_c is the speed limit implemented by the VSL, and $q_{\text{in}}(t)$ is actuated by an on-ramp metering upstream of the inlet. Note that the linearized model is valid at the optimal density ρ^* since the reference density is assumed to be chosen near the optimal value.

Online Optimization by ES Control

Here, we present the design of ES control with delay by following the procedure in [67]. The block diagram of the delay-compensated ES algorithm applied to the LWR PDE model is depicted in Figure 14.

Let $\hat{\varrho}(t)$ be the estimate of ρ^* , and $e(t)$ be the estimation error defined as

$$e(t) = \hat{\varrho}(t) - \rho^* \quad (179)$$

where $\hat{\varrho}(t)$ is an integrator of the predictor-based feedback signal $U(t)$ as $\hat{\varrho}(t) = U(t)$. From Figure 14, the error dynamics can be written as

$$\dot{e}(t - D) = U(t - D) \quad (180)$$

given the delayed estimation error dynamics modeled by $\epsilon(x, t) = U(t - (x/u))$.

We introduce the additive perturbation $S(t)$

$$S(t) = a \sin(\omega(t + D)) \quad (181)$$

and the multiplicative demodulation signals $(M(t), N(t))$ given by

$$M(t) = \frac{2}{a} \sin(\omega t), \quad N(t) = -\frac{8}{a^2} \cos(2\omega t) \quad (182)$$

where a and ω are, respectively, the amplitude and frequency of a slow periodic perturbation signal $a \sin(\omega t)$ introduced later. Using the demodulation signals, we

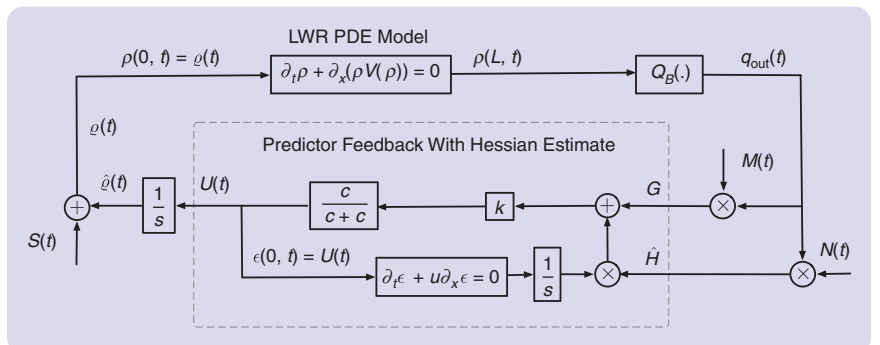


FIGURE 14 Block diagram for implementation of ES control design for the nonlinear LWR PDE model.

calculate estimates of the gradient and Hessian of the cost function, denoted as $(G(t), \hat{H}(t))$:

$$G(t) = M(t)q_{out}(t), \quad \hat{H}(t) = N(t)q_{out}(t) \quad (183)$$

where $\hat{H}(t)$ is the estimate of the unknown Hessian H . The averaging of $G(t)$ and $\hat{H}(t)$ yields

$$G_{av}(t) = He_{av}(t-D), \quad \hat{H}_{av} = (Nq_{out})_{av} = H. \quad (184)$$

Taking the average of (180), we have $\dot{e}_{av}(t-D) = U_{av}(t-D)$, where $U_{av}(t)$ is the averaged value for $U(t)$ to be designed later. Substituting this equation into (184) yields

$$\dot{G}_{av}(t) = HU_{av}(t-D). \quad (185)$$

The motivation for the predictor feedback design is to compensate for the delay by feeding back future states in the equivalent averaged system $G_{av}(t+D)$. Given an arbitrary control gain $k > 0$, we aim to design

$$U_{av}(t) = kG_{av}(t+D), \quad \forall t \geq 0 \quad (186)$$

which requires knowledge of future states. Therefore, we obtain the following by plugging (186) into (180):

$$\dot{e}_{av}(t) = U_{av}(t) = kHe_{av}(t), \quad \forall t \geq D. \quad (187)$$

Recalling that $k > 0, H < 0$, the equilibrium of the average system $e_{av}(t) = 0$ is exponentially stable. Applying the variation-of-constants formula, $G_{av}(t+D) = G_{av}(t) + \hat{H}_{av}(t) \int_{t-D}^t U_{av}(\tau) d\tau$, and from (186), one has

$$U_{av}(t) = k(G_{av}(t) + \hat{H}_{av}(t) \int_{t-D}^t U_{av}(\tau) d\tau) \quad (188)$$

which represents the future state $G_{av}(t+D)$ in (185) in terms of the average control signal $U_{av}(\tau)$ for $\tau \in [t-D, t]$. The control input is infinite dimensional by its use of history over the past D time units.

For the stability analysis in which the averaging theorem for infinite-dimensional systems is used, we employ a low-pass filter for the previous basic predictor feedback controller and then derive an infinite-dimensional and averaging-based predictor feedback given by

$$U(t) = \mathcal{T}\{k(G(t) + \hat{H}(t) \int_{t-D}^t U(\tau) d\tau)\} \quad (189)$$

where $k > 0$ is an arbitrary control gain, and the Hessian estimate $\hat{H}(t)$ is updated according to (183), satisfying the average property in (184). $\mathcal{T}\{\cdot\}$ is the low-pass filter operator defined by

$$\mathcal{T}\{\varphi(t)\} = \mathcal{L}^{-1}\left\{\frac{c}{s+c}\right\} * \varphi(t) \quad (190)$$

where $c \in \mathbb{R}^+$ is the corner frequency, \mathcal{L}^{-1} is the inverse Laplace transformation, and $*$ represents convolution in time.

Hence, according to [49, Theorem 1], we can conclude that

$$\limsup_{t \rightarrow +\infty} |\varrho(t) - \rho^*| = \mathcal{O}(a + 1/\omega) \quad (191)$$

$$\limsup_{t \rightarrow +\infty} |q_{out}(t) - q^*| = \mathcal{O}(a^2 + 1/\omega^2). \quad (192)$$

OPTIMAL OIL-DRILLING CONTROL

A common type of instability in oil drilling is a friction-induced stick-slip oscillation (see [10] and the references therein), which results in torsional vibrations of the drill string and can severely damage the drilling facilities (see Figure 2 from [1]).

Figure 15 shows a modern land-based drilling rig. The tower operates like the derrick of a crane: the traveling block is connected by several steel drill lines with one attached to the deadline anchor and the other being spooled on a drum controlled by ac induction motors. Another electric motor, called the top drive, is connected to the traveling block. The top drive is used to rotate the drill string, a set of hundreds of drill pipes (each about 30 ft long) that conducts the bore hole assembly (BHA). The BHA contains several sensors (for pressure, temperature, and vibration, among others) and the drill bit itself. There are several different types of drill bit designs and materials for drilling different geological formations.

In analogy, the rig operates similarly to a drill press, but with a drill bit that is several inches wide (4–36 in. is a common range) and up to several miles long (an onshore



FIGURE 15 The topside of a drilling rig.

well can be as shallow as 200 yards or as deep as 2 miles). By rotating this drill string and using its weight to generate an axial force, the BHA mills the rocks, drilling the well. Because of the small diameter when compared to its length, the drill string is subject to axial and torsional effects, much like a flexible rod. Because of this elasticity, the force and velocity propagation can be modeled by *wave equations*.

In this particular model, the actuation is the velocity of the traveling block, that is, the axial velocity of the drill string on the surface. Although not considered here, the rotational velocity also influences the rate of penetration (ROP) in a real scenario. The model output is the weight on hook, which somewhat models the weight on bit (WOB). The WOB estimates the contact between the drill bit and the rock formation, and it is the downhole boundary condition to be controlled. In [1], the authors have discussed the feasibility of controlling the hook load to optimize the ROP while drilling.

The key point that enables such an approach is the concept of bit foundering [1], that is, the fact that the ROP tapers off (and sometimes starts decreasing) with increasing weight on the bit past the foundering point. This makes the static mapping between the ROP and the WOB upward convex in an interval around the foundering point. This transfers to an upward convex static mapping between the equilibrium hook load setpoint and feed rate. Consequently, these signals can be used as the plant input and output for the design of a drilling control system. Hence, this physical application motivates our ES scheme for static maps with actuation dynamics described by wave PDEs, as depicted in Figure 16.

DEEP-SEA CABLE-ACTUATED SOURCE SEEKING

This application involving deep-sea cable-actuated source seeking is illustrated in Figure 17. In this scenario, a sensor is suspended on a cable and moved from the sea surface using a surface vessel. The sensor operates without position awareness, primarily because of the challenging undersea environment. The task at hand is to locate the source signal as closely as possible. No external fluid flow (for example, water current) is considered, and the dynamics of the boat is ignored for simplicity [20].

The algorithm proposed here is designed to be applicable to such a source-seeking scenario. The objective of source seeking is to find the source of a signal of an unknown concentration field, which can be chemical, acoustic, electromagnetic, etc. The sensor captures this field, and its strength decays with distance, reaching a

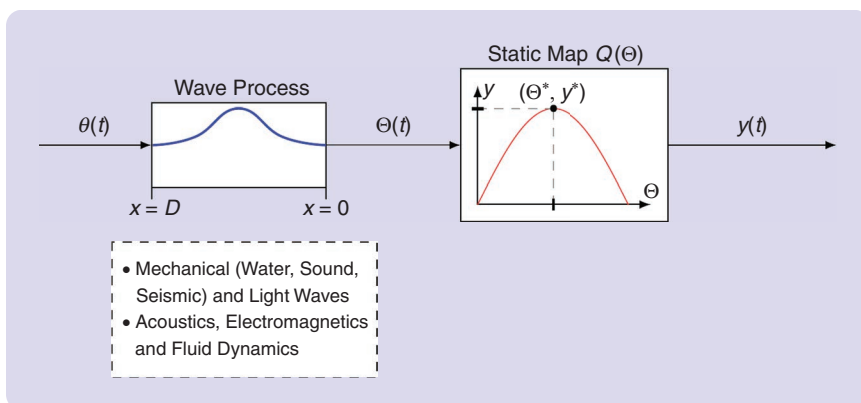


FIGURE 16 Cascade of wave PDE with a static map $y(\Theta(t)) = Q(\Theta(t))$. The extremum $y(t) = y^*$ is achieved for $\Theta(t) = \Theta^*$. Wave PDEs are used generally to model different sorts of processes, such as mechanical, acoustic, electromagnetic, and fluid dynamics processes.

maximum at position x_e (relative to the coordinate system of the surface vessel). Beyond finding the source signal, it also ensures the stabilization of the cable's motion. As the signal source becomes deeper or the cable length increases, the task becomes less demanding for the surface vessel because of the high natural frequency of the longer cable, which reduces the need for rapid vessel motion. Nevertheless, achieving stability in the PDE-compensating ES algorithm may require more extensive memory.




FIGURE 17 Motivating example of underwater search: x_e represents the relative linear position of the source signal with respect to the sensor. The control task aims to drive the sensor to the source signal, meaning that $x_e(t) \rightarrow 0$ (or to a small neighborhood of zero) as $t \rightarrow +\infty$.

The cable of this application is represented by a string described by the following PDE model over an interval $x \in [0, D]$:

$$\varepsilon \alpha_{tt} = (1 + d\partial_t) \alpha_{xx} \quad (193)$$

$$\alpha_x(0, t) = 0 \quad (194)$$

$$\alpha(0, t) = \text{measured} \quad (195)$$

$$\alpha(D, t) = \text{controlled.} \quad (196)$$

Equations (193)–(196) represent the dynamics of a string controlled at the end $x = D$, pinned to the surface vessel, and with a free end at $x = 0$, where the sensor is located. The solution $\alpha(x, t)$ of (193) represents the state variable of the PDE dynamics governing the motion of the cable. Equations (194)–(196) serve as boundary conditions. The constants ε , d , and D are positive. The constant D physically corresponds to the length of the cable. The value $1/\varepsilon$ represents the “stiffness” of the string, which can be expressed as E/ρ , where E denotes Young’s modulus and ρ the density of the material. The term $d\partial_t$ models the “Kelvin–Voigt” damping, representing the internal material damping, not the damping that arises from the viscous interaction of the string with the surrounding medium. We assume that this model takes into account a small amount of damping (d), which is a realistic consideration in any material. We do not rely on the Kelvin–Voigt term as a source of energy dissipation; instead, we use it as a means of enhancing the controllability of the model (193)–(196).

Scalar Maps With Actuation PDE Dynamics

Now, we consider the actuation dynamics described by a wave equation containing Kelvin–Voigt damping with $\varepsilon = 1$, $\theta(t) \in \mathbb{R}$, and the sensor $\Theta(t) \in \mathbb{R}$ given by

$$\Theta(t) = \alpha(0, t) \quad (197)$$

$$\partial_{tt} \alpha(x, t) = \partial_{xx} \alpha(x, t) + d \partial_{xxt} \alpha(x, t) \quad (198)$$

$$\partial_x \alpha(0, t) = 0 \quad (199)$$

$$\alpha(D, t) = \theta(t) \quad (200)$$

where $\alpha: [0, D] \times \mathbb{R}_+ \rightarrow \mathbb{R}$, and D is the known domain length, as mentioned before. The output signal measured with the sensor is represented by the unknown static map

$$y(t) = Q(\Theta(t)) \quad (201)$$

with input $\Theta(t)$ in (197).

We assume that the unknown nonlinear map is locally quadratic, as in (154), resulting in (155). Adopting the proposed scheme in [59] and combining (197)–(200) with the ES approach, the closed-loop ES with actuation dynamics governed by the Kelvin–Voigt PDE is illustrated in Figure 18.

Trajectory Generation for the Probing Signal

The perturbation $S(t)$ is adapted from the basic ES scheme to accommodate actuation dynamics. The trajectory generation problem, as described in [42, Chapter 12], is outlined as follows:

$$S(t) := \beta(D, t) \quad (202)$$

$$\partial_{tt} \beta(x, t) = \partial_{xx} \beta(x, t) + d \partial_{xxt} \beta(x, t) \quad (203)$$

$$\partial_x \beta(0, t) = 0 \quad (204)$$

$$\beta(0, t) = a \sin(\omega t) \quad (205)$$

where $\beta: [0, D] \times \mathbb{R}_+ \rightarrow \mathbb{R}$. The explicit solution of (202) is derived for the reference trajectory

$$\beta(D, t) := \beta^r(D, t) = S(t), \quad \beta(0, t) := \beta^r(0, t) = a \sin(\omega t). \quad (206)$$

This solution is found by postulating the reference solution $\beta^r(x, t)$ as a power series of the spatial variable with time-dependent coefficients: $\beta^r(x, t) = \sum_{i=0}^{\infty} a_i(t) x^i / i!$, as in [44]. The string reference solution is given by [73]

$$\beta^r(x, t) = -a \frac{j}{2} [\cosh(j\sigma x) e^{j\omega t} - \cosh(j\bar{\sigma} x) e^{-j\omega t}] \quad (207)$$

with $\sigma = \omega / \sqrt{1 + j\omega d}$ and $\bar{\sigma}$ being its complex conjugate. Equation (207) can be written as the purely real function

$$\beta^r(x, t) = \frac{a}{2} [e^{\hat{\beta}(\omega)x} \sin(\omega t + \bar{\beta}(\omega)x) + e^{-\hat{\beta}(\omega)x} \sin(\omega t - \bar{\beta}(\omega)x)] \quad (208)$$

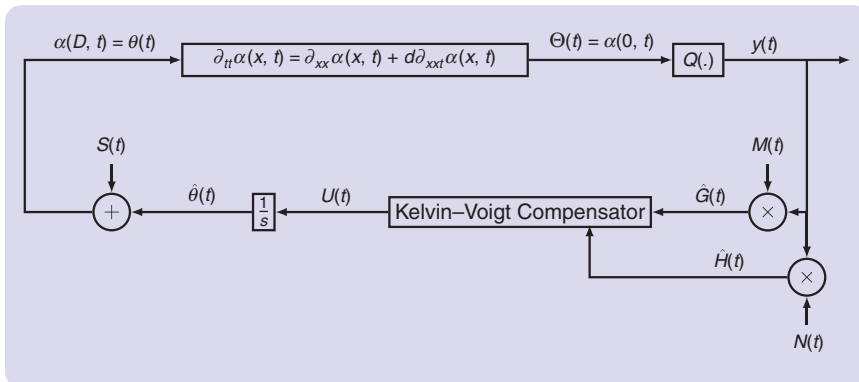


FIGURE 18 Gradient ES control loop.

where the real functions $\bar{\beta}(\omega)$ and $\hat{\beta}(\omega)$ are defined as

$$\bar{\beta}(\omega) = \omega \sqrt{\frac{\sqrt{1 + \omega^2 d^2} + 1}{2(1 + \omega^2 d^2)}} \quad (209)$$

$$\hat{\beta}(\omega) = \omega \sqrt{\frac{\sqrt{1 + \omega^2 d^2} - 1}{2(1 + \omega^2 d^2)}}. \quad (210)$$

On the other hand, the demodulation signals $M(t)$ and $N(t)$, used for estimating the gradient and Hessian,

respectively, of the static map by multiplying them by the output $y(t)$, are defined in [23] as

$$\hat{H}(t) = N(t)y(t) \text{ with } N(t) = -\frac{8}{a^2} \cos(2\omega t) \quad (211)$$

$$G(t) = M(t)y(t) \text{ with } M(t) = \frac{2}{a} \sin(\omega t). \quad (212)$$

Estimation Errors and Error Dynamics

Since our objective is to find Θ^* , which corresponds to the optimal unknown actuator $\theta(t)$, we introduce the following estimates and estimation errors:

$$\hat{\theta}(t) = \theta(t) - S(t), \quad \hat{\Theta}(t) = \Theta(t) - a \sin(\omega t) \quad (213)$$

$$\tilde{\theta}(t) := \hat{\theta}(t) - \Theta^*, \quad \vartheta(t) := \hat{\Theta}(t) - \Theta^*. \quad (214)$$

Let $\tilde{\alpha}: [0, D] \times \mathbb{R}_+ \rightarrow \mathbb{R}$ be defined as $\tilde{\alpha}(x, t) := \alpha(x, t) - \beta(x, t) - \Theta^*$. Manipulating (197)–(200) and (202)–(205) with the help of (213) and (214), we get

$$\vartheta(t) = \tilde{\alpha}(0, t) \quad (215)$$

$$\partial_{tt} \tilde{\alpha}(x, t) = \partial_{xx} \tilde{\alpha}(x, t) + d \partial_{xxt} \tilde{\alpha}(x, t) \quad (216)$$

$$\partial_x \tilde{\alpha}(0, t) = 0 \quad (217)$$

$$\tilde{\alpha}(D, t) = \tilde{\theta}(t). \quad (218)$$

The error dynamics is obtained by taking the time derivative of (215)–(218) and using $\dot{\tilde{\theta}} = U(t)$ and $u(x, t) = \tilde{\alpha}_t(x, t)$:

$$\dot{\vartheta}(t) = u(0, t) \quad (219)$$

$$\partial_{tt} u(x, t) = \partial_{xx} u(x, t) + d \partial_{xxt} u(x, t) \quad (220)$$

$$\partial_x u(0, t) = 0 \quad (221)$$

$$u(D, t) = U(t). \quad (222)$$

Boundary ES Control Law

We consider the PDE–ODE cascade shown in Figure 19 and use the backstepping transformation

$$w(x, t) = u_{av}(x, t) - \bar{K} \int_0^x k(x, \sigma) u_{av}(\sigma, t) d\sigma - \bar{K} \vartheta_{av}(t) \quad (223)$$

to transform the original (219)–(222) into the target system

$$\dot{\vartheta}_{av}(t) = \bar{K} \vartheta_{av}(t) + w(0, t), \quad \bar{K} < 0 \quad (224)$$

$$w_{tt} = (1 + d \partial_t)(w_{xx} - cw), \quad c > 0 \quad (225)$$

$$w_x(0, t) = 0 \quad (226)$$

$$w(D, t) = 0. \quad (227)$$

The gain kernel PDE $k(x, \sigma)$ comes from the solution of (see [42, Section 4.2])

$$k_{xx} = k_{\sigma\sigma} + ck, \quad k_\sigma(x, 0) = 0, \quad k(x, x) = \frac{c}{2}x. \quad (228)$$

The solution to the PDE in (228) is obtained through a summation of successive approximation series [42, Section 4.4]:

$$k(x, \sigma) = -cx \frac{I_1(\sqrt{c(x^2 - \sigma^2)})}{\sqrt{c(x^2 - \sigma^2)}} \quad (229)$$

and, from (223) and (227), the average control law is given by

$$u_{av}(D, t) = \bar{K} \vartheta_{av}(t) - \bar{K} \int_0^D cD \frac{I_1(\sqrt{c(D^2 - \sigma^2)})}{\sqrt{c(D^2 - \sigma^2)}} u_{av}(\sigma, t) d\sigma \quad (230)$$

where $I_1(z) := \sum_{m=0}^{\infty} (z/2)^{1+2m} / m!(m+1)!$ is the modified Bessel function [42, Appendix A.2]. Thus, introducing a result of [23], the averaged versions of the gradient and Hessian estimate are calculated as

$$G_{av}(t) = H \vartheta_{av}(t), \quad \hat{H}_{av}(t) = H. \quad (231)$$

From (222) and (230), choosing $\bar{K} = KH$ with $K > 0$ and plugging in the average gradient and Hessian estimates (231), we obtain

$$U_{av}(t) = KG_{av}(t) - KH \int_0^D cD \frac{I_1(\sqrt{c(D^2 - \sigma^2)})}{\sqrt{c(D^2 - \sigma^2)}} u_{av}(\sigma, t) d\sigma. \quad (232)$$

We introduce a low-pass filter to obtain the nonaverage controller

$$U(t) = \frac{\bar{c}}{s + \bar{c}} \left\{ K \left[G(t) - \hat{H}(t) \int_0^D cD \frac{I_1(\sqrt{c(D^2 - \sigma^2)})}{\sqrt{c(D^2 - \sigma^2)}} u(\sigma, t) d\sigma \right] \right\} \quad (233)$$

with $\bar{c} \rightarrow +\infty$ sufficiently large.

MATERIALS PHASE CHANGE PDE ES CONTROL: FROM FIXED DOMAIN TO MOVING BOUNDARY

We next present the design and analysis of the ES for static maps, with input governed by a PDE of the diffusion type defined on a time-varying spatial domain described by an ODE. We compensate for the average-based actuation dynamics by a controller via backstepping transformation for the moving boundary, which is utilized to transform the original coupled PDE–ODE into a target system whose exponential stability of the average equilibrium of the average system is proved.

One-Phase Stefan Problem

The physical model that describes the 1D Stefan problem in a pure one-component material of length L is




FIGURE 19 The cascade of the PDE dynamics and the ODE integrator.

described in Figure 20. The domain $[0, L]$ is divided into two subdomains $[0, s(t)]$ and $[s(t), L]$, which represent the liquid phase and the solid phase, respectively. The system is controlled by the heat flux $q_c(t)$ at $x=0$ because we are dealing with a Neumann boundary actuation, as shown here:

$$T_t(x, t) = \alpha T_{xx}(x, t), \quad x \in (0, s(t)), \quad \alpha = \frac{k}{\rho C_p} \quad (234)$$

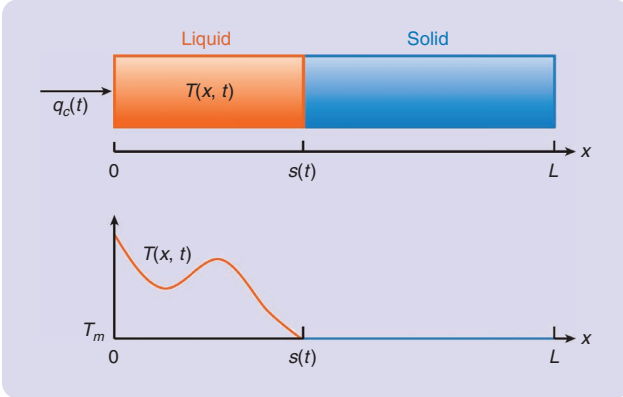


FIGURE 20 Schematic of one-phase Stefan problem [37], [38]. The temperature profile in the solid phase is assumed to be a uniform melting temperature.




FIGURE 21 The cascade of the PDE dynamics and the ODE system.




FIGURE 22 Control objective of the Stefan problem. We aim to design a heat flux input $q_c(t)$ such that the interface position $s(t)$ is driven to the setpoint position s^* .

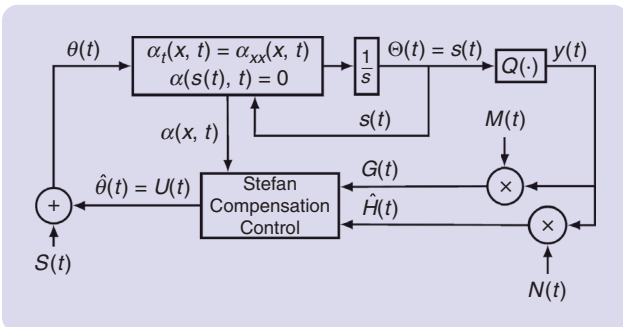


FIGURE 23 ES control loop applied to the one-phase Stefan problem.

$$-kT_x(0, t) = q_c(t) \quad (235)$$

$$T(s(t), t) = T_m \quad (236)$$

$$\dot{s}(t) = -\beta T_x(s(t), t), \quad \beta = \frac{k}{\rho \Delta H^*} \quad (237)$$

where $T(x, t)$, T_m , $q_c(t)$, k , ρ , C_p , and ΔH^* are the distributed temperature of the liquid phase, melting temperature, manipulated heat flux, liquid heat conductivity, liquid density, liquid heat capacity, and latent heat of fusion, respectively. Equations (235) and (236) are the boundary conditions of the system, and (237) is the Stefan condition, which describes the dynamics of the moving boundary. Figure 21 shows the block diagram of the PDE–ODE cascade represented by equations (234)–(237).

Actuation Dynamics and Output Signal

For simplicity, we consider actuation dynamics described by a heat equation with $\alpha, \beta, k = 1$, $\theta(t) \in \mathbb{R}$, and the propagated actuator $\Theta(t) \in \mathbb{R}$ given by

$$\dot{\Theta}(t) = \dot{s}(t) = -\alpha_x(s(t), t), \quad x \in (0, s(t)) \quad (238)$$

$$\partial_t \alpha(x, t) = \partial_{xx} \alpha(x, t) \quad (239)$$

$$\alpha(s(t), t) = 0 \quad (240)$$

$$-\partial_x \alpha(0, t) = \theta(t) \quad (241)$$

where $\alpha: [0, s(t)] \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is $\alpha(x, t) = T(x, t) - T_m$, and $s(t) = \Theta(t)$ is the unknown interface represented as the moving boundary. The output is measured by the unknown static map with input (238), according to (153). The ES goal is to optimize an unknown static map $Q(\cdot)$ using a real-time optimization control with optimal unknown output y^* and optimizer Θ^* as well as measurable output y and input θ . Consequently, the control objectives of the Stefan problem are achieved, that is, $\lim_{t \rightarrow \infty} s(t) = s^*$ and $\lim_{t \rightarrow \infty} T(x, t) = T_m, \forall x \in [0, s^*]$, as illustrated in Figure 22.

On the other hand, the unknown nonlinear map is locally quadratic, given by (154), such that the output of the static map is simply (155). Adapting the proposed scheme in [59] and combining (238)–(241) with the ES approach, the closed-loop ES with actuation dynamics is shown in Figure 23.

The demodulation signal $N(t)$, which is used to estimate the Hessian of the static map by multiplying it with the output $y(t)$ of the static map, is defined in [23] as

$$\hat{H}(t) = N(t)y(t) \quad \text{with} \quad N(t) = -\frac{8}{a^2} \cos(2\omega t) \quad (242)$$

whereas the signal $M(t)$ is used to estimate the gradient of the static map as follows:

$$G(t) = M(t)y(t) \quad \text{with} \quad M(t) = \frac{2}{a} \sin(\omega t). \quad (243)$$

Additive Probing Signal

The perturbation $S(t)$ is adapted from the basic ES to the case of PDE actuation dynamics. The trajectory generation problem, as in [42], is described as follows:

$$S(t) := -\partial_x \beta(0, t), \quad x \in (0, s(t)) \quad (244)$$

$$\partial_t \beta(x, t) = \partial_{xx} \beta(x, t) \quad (245)$$

$$\beta(s(t), t) = 0 \quad (246)$$

$$\beta_x(s(t), t) = -a\omega \cos(\omega t) \quad (247)$$

where $\beta: [0, s(t)] \times \mathbb{R}_+ \rightarrow \mathbb{R}$. The explicit solution of (244) is found, respectively, for the reference trajectory and the reference solution postulated by a power series [16]:

$$s(t) = a \sin(\omega t) \quad (248)$$

$$\beta(x, t) = \sum_{i=0}^{\infty} \frac{a_i(t)}{i!} [x - s(t)]^i. \quad (249)$$

We can calculate the first coefficients of the power series, replacing the boundary conditions (246) and (247) at (249), such that

$$a_0(t) = 0, \quad a_1 = -\dot{s}(t). \quad (250)$$

The general expression $a_i(t) = \dot{a}_{i-2}(t) - a_{i-1}(t)\dot{s}(t)$ is obtained by substituting (249) in (245). We provide here the analytic expression of the first four coefficients of the series (249) so that one can see how the successive derivatives of $s(t)$ appear:

$$a_2(t) = \dot{s}(t)^2 \quad (251)$$

$$a_3(t) = \ddot{s}(t) - \dot{s}(t)^3 \quad (252)$$

$$a_4(t) = \ddot{s}(t)^2 + \ddot{s}(t)\dot{s}(t) + \dot{s}(t)^4. \quad (253)$$

The trajectory generation solution that provides all terms of the power series (249) is given by [28]

$$\beta(x, t) = \sum_{i=0}^{\infty} \frac{1}{(2i)!} \frac{\partial^i}{\partial t^i} [x - s(t)]^{2i}. \quad (254)$$

Although (254) is not an explicit expression, by choosing suitable values for a and ω in (248), the series converges with few iterations of the infinite sum, getting the desirable sinusoidal signal $s(t)$ in the output of the integrator.

According to (244), we take the spatial derivative of (254) and substitute $x = 0$, thus arriving at the final expression of

$$S(t) = -\sum_{i=0}^{\infty} \frac{1}{(2i-1)!} \frac{\partial^i}{\partial t^i} [-a \sin(\omega t)]^{2i-1}. \quad (255)$$

Estimation Errors and PDE Error Dynamics

Since our objective is to find Θ^* , which corresponds to the optimal unknown actuator $\theta(t)$, we introduce the following estimation errors:

$$\hat{\theta}(t) = \theta(t) - S(t), \quad \hat{\Theta}(t) = \Theta(t) - a \sin(\omega t) \quad (256)$$

$$\tilde{\theta}(t) := \hat{\theta}(t) - \Theta^*, \quad \vartheta(t) := \hat{\Theta}(t) - \Theta^* \quad (257)$$

recalling that $\Theta(t) := s(t)$. Combining $\hat{\Theta}(t)$ in (256) and (257), we obtain the relation between the propagated estimation error $\vartheta(t)$, the propagated input $\Theta(t)$, and the optimizer of the static map Θ^* :

$$\Theta(t) - \Theta^* = \vartheta(t) + a \sin(\omega t). \quad (258)$$

Let us define

$$u(x, t) = \alpha(x, t) - \beta(x, t) \quad (259)$$

$$\hat{\theta}(t) = U(t). \quad (260)$$

By (238)–(241) and (244)–(247), with the help of (256) and (257), we have our original system:

$$\dot{\vartheta}(t) = -u_x(s(t), t), \quad x \in (0, s(t)) \quad (261)$$

$$u_t(x, t) = u_{xx}(x, t) \quad (262)$$

$$u(s(t), t) = 0 \quad (263)$$

$$-u_x(0, t) = U(t). \quad (264)$$

Stefan Compensation Control Law

We consider the PDE–ODE cascade (261)–(264) and use the backstepping transformation

$$w(x, t) = u(x, t) - \bar{K} \int_x^{s(t)} (x - \sigma) u(\sigma, t) dy - \bar{K}(x - s(t)) \vartheta(t) \quad (265)$$




FIGURE 24 The initial states $T(x, 0)$ (green curve) and the convergence of $T(s(t), t)$ (blue curve) to $T_m = 100^\circ\text{C}$ in a 3D space for the PDE state $T(x, t)$. The blue curve shows the expansion of the domain of the Stefan PDE.

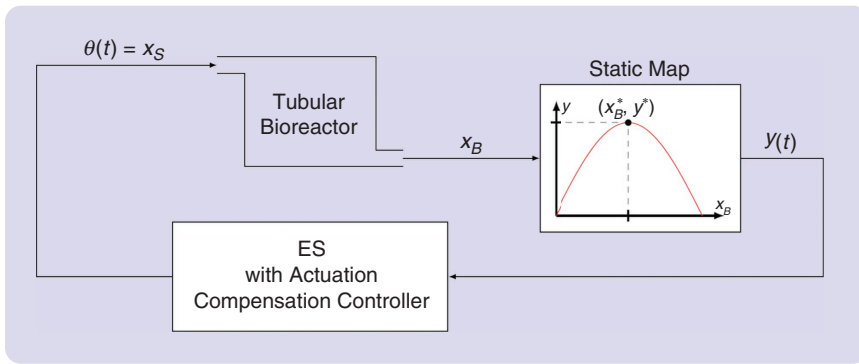


FIGURE 25 ES scheme for tubular bioreactors.

with $\bar{K} > 0$ as an arbitrary controller gain. Equation (265) transforms (261)–(264) into the target system

$$\dot{\vartheta}(t) = -\bar{K}\vartheta(t) - w_x(s(t), t), \quad x \in (0, s(t)) \quad (266)$$

$$w_t(x, t) = w_{xx}(x, t) + \bar{K}\dot{s}(t)\vartheta(t) \quad (267)$$

$$w_x(0, t) = 0 \quad (268)$$

$$w(s(t), t) = 0. \quad (269)$$

The compensation controller can be obtained by taking the derivative of (265) with respect to t and x , respectively, along the solution (261)–(264), and substituting $x = 0$:

$$U(t) = -\bar{K}\left(\vartheta(t) + \int_0^{s(t)} u(x, t) dx\right). \quad (270)$$

Since we have no measurement on $\vartheta(t)$, (270) is not applicable directly. Thus, introducing a result of [23], the average versions of the gradient and Hessian estimates are calculated by

$$G_{av}(t) = H\vartheta_{av}(t), \quad \hat{H}_{av}(t) = H. \quad (271)$$

Averaging (270), choosing $\bar{K} = KH$ with $K < 0$, and plugging in the average gradient and Hessian estimates (271), we obtain

$$U_{av}(t) = -KG_{av}(t) - KH \int_0^{s_{av}(t)} u_{av}(x, t) dx. \quad (272)$$

We introduce a low-pass filter to the controller with the purpose of applying the average theorem for infinite-dimensional systems [25], such that

$$U(t) = \frac{c}{s+c} \left\{ K \left[G(t) + \hat{H}(t) \int_0^{s(t)} u(x, t) dx \right] \right\} \quad (273)$$

for $c > 0$ sufficiently large.

Figure 24 corresponds to the numerical plot of the temperature profile for the closed-loop system, with the control law (273), converging in a 3D space (taking into account the domain $L = 0.9$ and the time t) to a close neighborhood of $T_m = 100$ °C. The sinusoidal movement

of $s(t)$ around $s^* = 0.75$ would violate the usual conditions for the Stefan problem that the temperature remains above or below the melting temperature on the whole interval $(0, s(t))$, forming a periodic chain of liquid and solid. However, it is well known in physics that the phase transition is not sharp but gradual. For a small perturbation amplitude, the system will be operating in the “phase transition range” when it reaches the periodic steady state at the extremum. A more detailed model is needed to

capture the system dynamics in that narrow phase transition range. In addition, we could redesign the algorithm to introduce vanishing probing signals and tapering off the perturbation after the extremum neighborhood is achieved, as studied in [39] to obtain unbiased [81] and safe [78] ES designs for a Stefan PDE with control barrier functions.

BIOREACTORS

Next, we present an application idea for this novel control concept: a tubular bioreactor. We consider the optimization problem in Figure 25 of a tubular bioreactor, where the goal is to operate the bioreactor at an unknown optimal product rate, for example, the growth of biomass, by determining the optimal input and hence the substrate concentration, for example, the glucose concentration. The product concentration x_B is generally not measurable, unlike the product rate.

Since the static map, which determines the product rate depending on the product concentration x_B , is not known or only approximately known, we can apply our introduced control concept. A simple model of a tubular bioreactor is presented by Winkin et al. [79], where the chemical reaction of a reactant R and a product P is given by $R - \lambda P$, where the stoichiometric coefficient λ of the reaction is considered. In its simplest form, the model is linear, since the nonlinear reaction term is simplified to a linear kinetic model depending only on the reactant concentration. The control-loop structure to reach and operate the tubular bioreactor at the optimal (highest product rate) is shown in Figure 26.

The boundary operators of the tubular reactor in Figure 26 are defined as

$$\begin{aligned} \mathcal{R}_D \underline{\alpha} &= \begin{bmatrix} \epsilon \partial_x \alpha_1(D, t) - b \alpha_1(D, t) \\ \epsilon \partial_x \alpha_2(D, t) - b \alpha_2(D, t) \end{bmatrix} = \begin{bmatrix} -b \theta(t) \\ 0 \end{bmatrix} \\ \mathcal{R}_0 \underline{\alpha} &= \begin{bmatrix} \epsilon \partial_x \alpha_1(0, t) \\ \epsilon \partial_x \alpha_2(0, t) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \end{aligned} \quad (274)$$

In a first step, the average-based infinite-dimensional controller used to compensate the tubular reactor process,

which is the actuator dynamics in this case, has to be derived. Note that the dynamics of the reactor has to be known to apply this control concept. Baccoli et al. [8] consider similar dynamics, that is, coupled linear parabolic PDEs, especially reaction–diffusion equations, with boundary input in one variable.

We additionally have an ODE cascade arising from the integrator in the control loop, which has to be stabilized. There seems to be no work that considers a coupled parabolic PDE–ODE cascade with two parabolic PDEs coupled in domain plus boundary input in one variable and boundary measurement in the other variable. The controller derivation and their derivations, like the perturbation signal $S(t)$ for this application, would go beyond the scope of this article. We emphasize that invoking the averaging theorem for infinite-dimensional systems in “Averaging Theorem for General Infinite-Dimensional Systems [25]” would work since the publication [79] showed the analytic semigroup property of the operator, which describes the coupled parabolic PDE system in Figure 26. Furthermore, the exponential stability proof of the average system will follow the same steps but with an extended Lyapunov–Krasovskii functional and more calculation steps. There may be restrictions on the system parameters ϵ, b, k_0 , and λ .

INFINITE-DIMENSIONAL LIGHT-SOURCE SEEKING

While the wave equation is the most appropriate “point of entry” into the realm of hyperbolic PDEs, beam equations are considered a physically relevant benchmark for control of hyperbolic PDEs and structural systems in general.

The simplest beam model is the Euler–Bernoulli model

$$\alpha_{tt} + \alpha_{xxxx} = 0 \quad (275)$$

$$\alpha_x(0, t) = \alpha_{xx}(0, t) = \alpha_{xxx}(0, t) = 0 \quad (\text{free-end condition}) \quad (276)$$

$$\alpha(0, t) = \text{clamped-end condition.} \quad (277)$$

The obvious difference between the PDEs (193)–(196) and (275)–(277) is in the number of spatial derivatives—the wave equation is second order in x , whereas the Euler–Bernoulli beam model is fourth order in x . One consequence of this difference is that a wave equation requires one boundary condition per endpoint [see (195) or (196)], whereas the Euler–Bernoulli beam model requires two boundary conditions per endpoint; see (276) or (277).

A more important difference is in the eigenvalues. Both the beam and the string models have all of their eigenvalues on the imaginary axis. However, while the string eigenvalues are equidistant, the beam eigenvalues get further

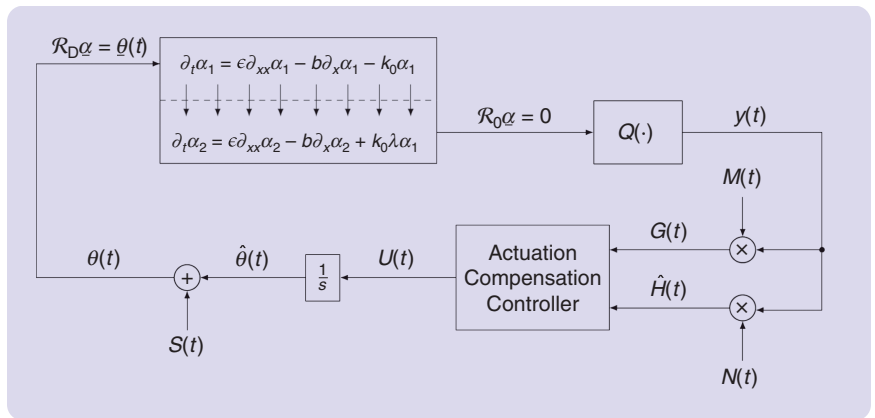


FIGURE 26 Basic gradient ES scheme for tubular bioreactors modeled by RAD PDEs.

and further apart as they go up the imaginary axis. This difference in the eigenvalue pattern is a consequence of the difference in the number of derivatives in x .

The reader might wonder how these differences translate into control. Is it obvious that a beam is more difficult to control than a string? The answer is not clear and is not necessarily “yes.” While the presence of higher derivatives clearly generates some additional issues to deal with in the control design, the wave equation has its own peculiarities that should not be underestimated. For example, controllability results for beams are valid on arbitrarily short time intervals, whereas for strings such results hold only over time intervals that are lower bounded in proportion to the “wave propagation speed” of the string (which physically corresponds to the “string tension”). Also, it is almost intuitively evident that keeping a string from vibrating may not be easier than keeping a beam from vibrating.

Fortunately, both the *backstepping design* and the *trajectory generation problem* for the Euler–Bernoulli model were already studied in [42, Section 8.2 and Example 12.8], respectively. As before, in the source-seeking context, we can assume $\alpha(0, t) = \Theta$ and $\alpha(D, t) = \theta$, where D would be the

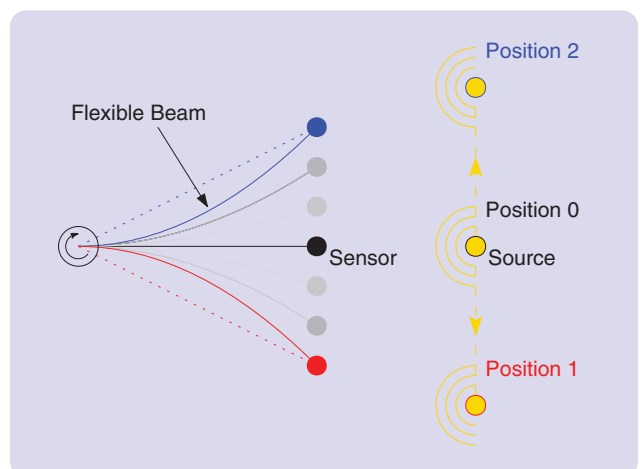


FIGURE 27 Light-source seeking with a coupled flexible beam.

length of the flexible beam. The illustrated system in Figure 27 is composed of a flexible and inextensible Euler–Bernoulli beam clamped to a rotating actuator hub with a light sensor mass at its free end. The control task aims to control the rotating actuator to drive the sensor to the moving light-source signal, while compensating the vibration generated during the beam displacement. Unlike the standard output regulation perspective, the paradigm here

is the maximization of the signal perceived by the light sensor, while position regulation of the end effector is also indirectly achieved.

NMES

The recent article [65] on NMES presents a promising application of ES under nonconstant delays. The authors have proposed a stochastic proportional-integral-derivative (PID) automatic tuner via ES and applied it to precise tracking of a flexion-extension reference for NMES and motor relearning for rehabilitation. Photos of the mechanical apparatus, as well as the electrical stimulation device for NMES experimental tests, are shown in Figure 28. The experimental results are innovative since, unlike for the referred literature, stroke patients rather than only healthy subjects were recruited for the successful tests. Remarkably, the patients are using such a device at the Public Hospital Universitário Clementino Fraga in Rio de Janeiro, Brazil.

In general, NMES devices are applied to clinical work in an open-loop fashion, and their parameters must be set at the beginning of therapy, not to facilitate clinical practice. The levels of electrical stimulation follow precalibrated profiles, requiring the presence of a practitioner to modify the stimulation parameters. This requires protocols aimed at enhancing muscle contraction together with the execution of the intended contractions. The downside of this procedure is that the device always returns the same portion of electrical assistance to the patient if there is no therapist intervention. In addition, the open-loop devices are not prepared to promote a proper association by means of some feedback error between the subject’s intended movement and the artificial activation provided by the NMES system.

In this sense, closed-loop strategies are adequate to generate the NMES electrical current amplitudes based on the




FIGURE 28 Mechanical apparatus for NMES experimental tests. In (b), the point A in the image indicates a goniometer (a simple potentiometer) linked to a steel axis B allowing angular displacement readings. Letter C shows that the wrist has an attachment with linear freedom of movement along the aluminum square rod, while D points out that there is an adjustment for the lateral distance of the elbows. In (a), the controlled joint angle, denoted by y , and the NMES equipment are presented.




FIGURE 29 Open-loop responses for different step-current inputs of a stroke patient with delays of order 300 ms. (Extracted from [65].)

angular displacement (or the measurement of some other variable) related to the upper limbs. Although PID controllers have been explored in different engineering applications, the true limitation is that a PID controller is designed for linear systems, but the neuromuscular plant, which is being controlled, is nonlinear, time-varying, and subject to delays. For instance, it is possible to note a time delay in the subject responses shown in Figure 29. The electromechanical–neuromuscular delay is in general of a time-varying nature and distinct for each subject. Moreover, it is worth mentioning that clinicians’ knowledge of control systems is limited. Therefore, their expertise in tuning controllers is limited as well. Furthermore, in NMES applications, each patient is unique

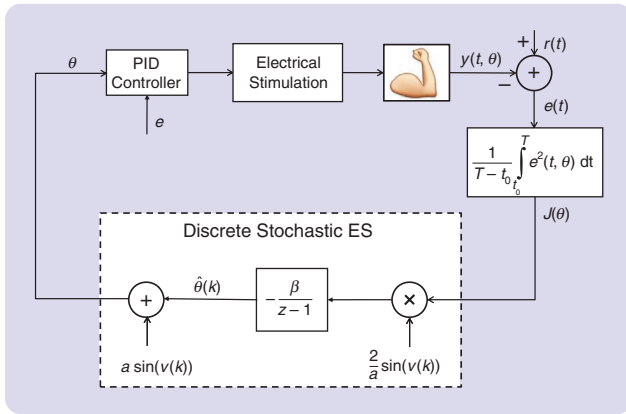


FIGURE 30 Block diagram of the closed-loop system for NMES using discrete-time stochastic ES, where $v(k)$ is the stochastic perturbation vector modeled by Gaussian white-noise signals, and the cost function $J(\theta) := 1/T - t_0 \int_{t_0}^T e^2(t, \theta) dt$ is defined over a period between two time instants t_0 and T . The PID control law is $u(t) = K_p e(t) + K_i \int_{t_0}^t e(\tau) d\tau + K_d (de(t)/dt)$, with $\theta := [K_p, K_i, K_d]^T$ being the PID parameters (gains) to be adapted by means of the ES algorithm.

and requires a particular set of PID parameters. Since it may be difficult to find the proper parameters for each patient, it is well motivated to have better procedures or a more intelligent adaptive controller [51].

On the other hand, manual tuning is a time-consuming task, and analytical methods are based on an exaggerated knowledge of the plant, requiring particular experimental validations to the identification of an acceptable plant model. However, a precise plant model in NMES is not known, and very long identification procedures are not desirable for the patients. This adverse environment of modeling inspires the application of adaptive and robust control methodologies [51] and automatic tuning techniques.

In this context, the model-free PID tuner via multiparameter stochastic ES in [65], inspired by its earlier deterministic version [35], is shown to be importantly fruitful. The proposed PID tuner eliminates the initial offline tests for patients since the control gains are automatically computed to minimize a cost function according to the tracking error $e(t) := y(t) - r(t)$ between the elbow's angle of the patient's arm $y(t)$ and the reference trajectory $r(t)$. This research is highly successful since the stochastic algorithm provides faster (better transient) responses, which is perfect for a self-tuning, fatigue-resistant control method for neuromuscular-based therapies. Moreover, the parameters of the stochastic ES are simpler to tune since the orthogonality assumption on the dither

vector signals of multiparameter deterministic ES imposes additional obstacles in adjusting the frequencies of the sinusoidal perturbations. Finally, deterministic ES may restrict the region of convergence of the algorithm, and adaptation using a periodic deterministic perturbation for learning may be rather poor and unusual in some model-free optimization frameworks. Stochastic perturbations overcome those obstacles as well.

Specifically, ES minimizes a cost function that quantifies the performance of the PID controller and iteratively modifies the arguments of the cost function (the PID parameters) so that its output reaches a local minimum. According to the block diagram in Figure 30, the time-domain implementation of the discrete-time stochastic ES algorithm is given by

$$\hat{\theta}_i(k+1) = \hat{\theta}_i(k) - \beta \frac{2}{a} \sin(v_i(k)) J(\hat{\theta}(k) + a \sin(v(k))) \quad (278)$$

where k is the discrete iteration number, the step size $\beta > 0$ is sufficiently small, the subscript $i = 1, 2, 3$ indicates the i th entry of a vector, $v(k) = [v_1(k) \ v_2(k) \ v_3(k)]^T$, and $\sin(v(k)) = [\sin(v_1(k)) \ \sin(v_2(k)) \ \sin(v_3(k))]^T$. The elements of the stochastic Gaussian perturbation vector $v(k)$ are sequentially and mutually independent such that $E\{v(k)\} = 0$, $E\{v_i^2(k)\} = \sigma_i^2$, and $E\{v_i(k) v_j(k)\} = 0$, $\forall i \neq j$, with $E\{\cdot\}$ denoting the expectation of a signal. In addition, it is assumed that the probability density function of the perturbation vector is symmetric about its mean.

Figures 31–33 show the advantages for the closed-loop responses for a stroke patient, originally presented in [65]. In the clinical scenario, Figure 31 illustrates that, even if the patients know the trajectory of the movement to be performed, they cannot execute it by themselves. Figure 32 also highlights that a fixed-gain PID scheme is not able to adequately control the stroke patient as the number of cycles/iterations increases, because of the time-varying nature of the neuromuscular system under delays. Indeed, PID




FIGURE 31 The graphic portrays the angular elbow joint movement performed by a stroke patient without the help of the proposed NMES controller. It can be seen that the stroke subject is not able to actively contract his arm to the final flexion position.

control with fixed gains is not able to bring satisfactory results in long-running tests. Moreover, a unique fixed-gain tuning is not applicable for different individuals.

On the other hand, the ES adaptive approach for simple adaptation of PID controller parameters is model free, having the interesting ability of controlling multiple subjects without tediously tuning by the designer or practitioner. The response curves in Figure 33 ratify the improved behavior of the adaptive PID control scheme

over a fixed-gain PID controller even in this adversarial scenario for NMES.

Although the delay discussion was not the focus in [65], it was evidenced there and in previous publications [2], [48], [71] that it may represent a significant challenge in NMES. It motivates the application of predictor feedback developed here for delay compensation in ES algorithms plus PID control or even other techniques [2], [3].

CONCLUSIONS

This article has generalized the ES results obtained in [21] to a wider class of infinite-dimensional systems governed by homogeneous and heterogeneous PDEs, rather than being restricted to noncooperative games governed by ODEs. We have introduced a nonmodel-based approach via ES and boundary control to find, in a distributed way, the Nash equilibria of noncooperative games with unknown quadratic payoff functions and the players acting through PDE dynamics. A player could stably attain its Nash equilibrium by measuring only the value of its payoff (no other information about the game is needed).

We did not consider just one kind of PDE dynamics in the players' decision variables. The challenge of PDE dynamics comes from mixed PDE classes since games are decentralized and heterogeneous—each player has its own particular dynamics, possibly of a different nature.

In this scenario, we were able to develop a result for heterogeneous games with distinct transport and diffusion (heat) PDEs to be simultaneously compensated—or homogeneous games with PDEs of different transport and diffusion coefficients—in the action paths of each player, where the players estimate only the diagonal entries of the Hessian matrix due to the players' own payoffs. We were able to dominate sufficiently small off-diagonal terms using a small-gain argument for the average system. Convergence to a small neighborhood of the Nash equilibrium is achieved, even in the presence of transport-heat PDEs.




FIGURE 32 Output responses for a stroke patient: PID with fixed gains ($K_p = 1, K_i = 1, K_d = 1$) not guaranteeing an acceptable trajectory tracking after 120 s.

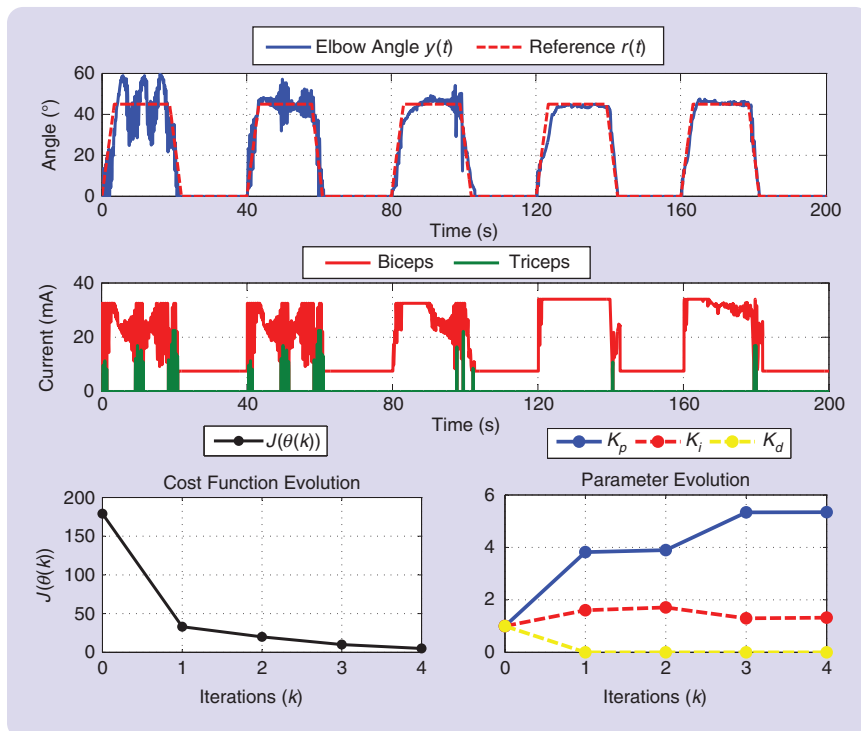


FIGURE 33 Stochastic ES-based PID control for NMES of a stroke patient.

In this article, the introduction to boundary control for noncooperative games was presented for both hyperbolic–transport and parabolic–heat types of PDEs. There is no strong reason why the exposition provided could not have been conducted on some of the other classes of PDEs. However, transport and heat PDEs are particularly convenient because they are sufficiently simple and general to serve as a design template that the reader can use to pursue extensions of the NES design to other classes of PDEs [42], [53] and its integration to other real-world applications described by a game-theoretic framework [31], [46], [87], [88].

We have explored a variety of potential engineering applications in which these advanced strategies could be effectively implemented. These applications included areas such as fluid dynamics, structural mechanics, and advanced energy and transportation systems. The aim was to demonstrate how the theoretical foundations can lead to innovative solutions to real-world engineering challenges, paving the way for new developments.

Future research involves extending these NES results for agents with delays and PDEs to distinct contexts, such as other delay compensation strategies and modern PDE-based frameworks under prescribed-time convergence [82], sampled-data implementation (with uncertainties in sampling or delay) [80], [89], and event-triggered control [67]. Relaxation of the modeling assumptions, particularly regarding the structure of the quadratic payoff functions and the choice of known and constant delays or diffusion coefficients, as well as the use of particular PDE types (for example, parabolic versus hyperbolic) in player dynamics, provides fertile territories for investigation. For instance, noncooperative games, free of PDEs but considering nonquadratic payoff functions and stubborn players, have already been considered in [21]. On the other hand, delay mismatches and uncertain diffusion coefficients have been gracefully addressed in the context of ES [70] and the adaptive control of parabolic or hyperbolic PDEs [6], [74], respectively.

ACKNOWLEDGMENT

The first author acknowledges the support of the Brazilian funding agencies CNPq, CAPES, and FAPERJ.

AUTHOR INFORMATION

Tiago Roux Oliveira (tiagoroux@uerj.br) joined the Department of Electronics and Telecommunication Engineering, State University of Rio de Janeiro, Rio de Janeiro RJ, 20.550-900, Brazil, as an associate professor in 2010. In 2014, he was a visiting scholar at the University of California, San Diego, CA, USA. In 2017, he was nominated as an affiliate member of the Brazilian Academy of Sciences. He has served as an associate editor on the Editorial Board of the *IEEE Transactions on Automatic Control*, *IEEE Open Journal of Control Systems*, *International Journal of Robust and Nonlinear Control*, *Systems & Control Letters*, *IEEE Control Systems Let-*

ters, Automatica, and other journals. In 2020, he was elected and nominated chair of the Technical Committee 1.2 (Adaptive and Learning Systems) of the International Federation of Automatic Control for the triennium 2020–2023 (being reelected for the triennium 2023–2026). In 2021, he received the IEEE Transactions on Control Systems Technology Outstanding Paper Award from the IEEE Control Systems Society. He is coauthor of the book *Extremum Seeking Through Delays and PDEs*, published by the Society for Industrial and Applied Mathematics in 2022. He is a Senior Member of IEEE.

Miroslav Krstić is with the Department of Mechanical and Aerospace Engineering, University of California San Diego, La Jolla, CA 92093-0411 USA. He has received awards that include the IEEE Roger W. Brockett Control Systems Award, the Bellman, Bode Lecture, Society for Industrial and Applied Mathematics (SIAM) Reid, Oldenburger, Ragazzini, Chestnut, Paynter, Nyquist, International Federation of Automatic Control (IFAC) Nonlinear Control, IFAC Ruth Curtain DPS, IFAC Adaptive and Learning Systems, Axelby, and Schuck (1996 and 2019). He has held chief or senior editorial positions of *Systems & Control Letters*, *Automatica*, and *IEEE Transactions on Automatic Control*. He has coauthored 19 books on adaptive, nonlinear, and stochastic control, extremum seeking, control of partial differential equation systems including turbulent flows, and control of delay systems. He is a Fellow of IEEE and of IFAC, SIAM, the American Society of Mechanical Engineers, the American Association for the Advancement of Science, the Institution of Engineering and Technology, the American Institute of Aeronautics and Astronautics, and the Serbian Academy of Sciences and Arts.

Tamer Başar has been with the Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign (UIUC), Urbana, IL 61801 USA, since 1981, where he is the Swanlund Endowed Chair Emeritus, Center for Advanced Study Professor Emeritus of Electrical and Computer Engineering, a research professor with the Coordinated Science Laboratory, and a research professor with the Information Trust Institute. He was with UIUC as director of the Center for Advanced Study from 2014 to 2020. He has more than 1,000 publications in systems, control, communications, and dynamic games, including books on game theory, robust control, network security, and stochastic networked control. He was the recipient of several awards and recognitions over the years, including the highest awards of the IEEE Control Systems Society (CSS), the International Federation of Automatic Control (IFAC), the American Automatic Control Council (AACC), and the International Society of Dynamic Games. He has received the IEEE Control Systems Award and numerous international honorary doctorates and professorships. He is a member of the U.S. National Academy of Engineering, a Fellow of the American Academy of Arts and Sciences,

and a Fellow of the Society for Industrial and Applied Mathematics and IFAC. He was president of the CSS and AACC. He is a Life Fellow of IEEE.

REFERENCES

- [1] U. J. Aarsnes, O. M. Aamo, and M. Krstić, "Extremum seeking for real-time optimal drilling control," in *Proc. IEEE Amer. Control Conf. (ACC)*, Philadelphia, PA, USA, 2019, pp. 5222–5227, doi: [10.23919/ACC.2019.8815162](https://doi.org/10.23919/ACC.2019.8815162).
- [2] N. Alibeji, N. Kirsch, S. Farrokhi, and N. Sharma, "Further results on predictor-based control of neuromuscular electrical stimulation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 23, no. 6, pp. 1095–1105, Nov. 2015, doi: [10.1109/TNSRE.2015.2418735](https://doi.org/10.1109/TNSRE.2015.2418735).
- [3] N. Alibeji and N. Sharma, "A PID-type robust input delay compensation method for uncertain Euler–Lagrange systems," *IEEE Trans. Control Syst. Technol.*, vol. 25, no. 6, pp. 2235–2242, Nov. 2017, doi: [10.1109/TCST.2016.2634503](https://doi.org/10.1109/TCST.2016.2634503).
- [4] T. Alpcan and T. Başar, *Network Security: A Decision and Game Theoretic Approach*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
- [5] S. Amin, G. A. Schwartz, and S. S. Sastry, "Security of interdependent and identical networked control systems," *Automatica*, vol. 49, no. 1, pp. 186–192, 2013, doi: [10.1016/j.automatica.2012.09.007](https://doi.org/10.1016/j.automatica.2012.09.007).
- [6] H. Anfinsen and O. M. Aamo, *Adaptive Control of Hyperbolic PDEs*. Cham, Switzerland: Springer-Verlag, 2019.
- [7] Z. Artstein, "Linear systems with delayed controls: A reduction," *IEEE Trans. Autom. Control*, vol. 27, no. 4, pp. 869–879, Aug. 1982, doi: [10.1109/TAC.1982.1103023](https://doi.org/10.1109/TAC.1982.1103023).
- [8] A. Baccoli, A. Pisano, and Y. Orlov, "Boundary control of coupled reaction–diffusion processes with constant parameters," *Automatica*, vol. 54, pp. 80–90, Apr. 2015, doi: [10.1016/j.automatica.2015.01.032](https://doi.org/10.1016/j.automatica.2015.01.032).
- [9] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. Philadelphia, PA, USA: SIAM, 1999.
- [10] N. Bekiaris-Liberis and M. Krstić, "Compensation of wave actuator dynamics for nonlinear systems," *IEEE Trans. Autom. Control*, vol. 59, no. 6, pp. 1555–1570, Jun. 2014, doi: [10.1109/TAC.2014.2309057](https://doi.org/10.1109/TAC.2014.2309057).
- [11] R. C. Carlson, I. Papamichail, M. Papageorgiou, and A. Messmer, "Optimal motorway traffic flow control involving variable speed limits and ramp metering," *Transp. Sci.*, vol. 44, no. 2, pp. 238–253, 2010, doi: [10.1287/trsc.1090.0314](https://doi.org/10.1287/trsc.1090.0314).
- [12] R. Carmona, J.-P. Fouque, S. M. Mousavi, and L.-H. Sun, "Systemic risk and stochastic games with delay," *J. Optim. Theory Appl.*, vol. 179, no. 2, pp. 366–399, 2018, doi: [10.1007/s10957-018-1267-8](https://doi.org/10.1007/s10957-018-1267-8).
- [13] M. D. Ciletti, "Differential games with information time lag: Norm-invariant systems," *J. Optim. Theory Appl.*, vol. 9, no. 5, pp. 293–301, 1972, doi: [10.1007/BF00932930](https://doi.org/10.1007/BF00932930).
- [14] P. Coughon, D. Dochain, M. Guay, and M. Perrier, "Real-time optimization of a tubular reactor with distributed feed," *AIChE J.*, vol. 52, no. 6, pp. 2120–2128, 2006, doi: [10.1002/aic.10799](https://doi.org/10.1002/aic.10799).
- [15] G. Dervisoglu, G. Gomes, J. Kwon, R. Horowitz, and P. Varaiya, "Automatic calibration of the fundamental diagram and empirical observations on capacity," in *Proc. Transp. Res. Board 88th Annu. Meeting*, 2009, vol. 15, pp. 1–14.
- [16] W. B. Dunbar, N. Petit, P. Rouchon, and P. Martin, "Motion planning for a nonlinear Stefan problem," *ESAIM: Control, Optim. Calculus Variations*, vol. 9, pp. 275–296, Feb. 2003, doi: [10.1051/cocv:2003013](https://doi.org/10.1051/cocv:2003013).
- [17] H. Ehtamo and R. P. Hämmäläinen, "Incentive strategies and equilibria for dynamic games with delayed information," *J. Optim. Theory Appl.*, vol. 63, no. 3, pp. 355–369, Dec. 1989, doi: [10.1007/BF00939802](https://doi.org/10.1007/BF00939802).
- [18] S. Fan and B. Seibold, "Data-fitted first-order traffic models and their second-order generalizations: Comparison by trajectory and sensor data," *Transp. Res. Rec.*, vol. 2391, no. 1, pp. 32–43, 2013, doi: [10.3141/2391-04](https://doi.org/10.3141/2391-04).
- [19] J. Feiling, S. Koga, M. Krstić, and T. R. Oliveira, "Gradient extremum seeking for static maps with actuation dynamics governed by diffusion PDEs," *Automatica*, vol. 95, pp. 197–206, Sep. 2018, doi: [10.1016/j.automatica.2018.05.023](https://doi.org/10.1016/j.automatica.2018.05.023).
- [20] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. Chichester, U.K.: Wiley, 2011.
- [21] P. Frihauf, M. Krstić, and T. Başar, "Nash equilibrium seeking in non-cooperative games," *IEEE Trans. Autom. Control*, vol. 57, no. 5, pp. 1192–1207, May 2012, doi: [10.1109/TAC.2011.2173412](https://doi.org/10.1109/TAC.2011.2173412).
- [22] M. L. Galvao, T. R. Oliveira, and M. Krstić, "Extremum seeking for Stefan PDE with moving boundary and delays," *IFAC-PapersOnLine*, vol. 55, no. 36, pp. 222–227, 2022, doi: [10.1016/j.ifacol.2022.11.361](https://doi.org/10.1016/j.ifacol.2022.11.361).
- [23] A. Ghaffari, M. Krstić, and D. Nešić, "Multivariable Newton-based extremum seeking," *Automatica*, vol. 48, no. 8, pp. 1759–1767, 2012, doi: [10.1016/j.automatica.2012.05.059](https://doi.org/10.1016/j.automatica.2012.05.059).
- [24] V. Y. Glizer and J. Shinar, "Optimal evasion from a pursuer with delayed information," *J. Optim. Theory Appl.*, vol. 111, no. 1, pp. 7–38, 2001, doi: [10.1023/A:1017515129544](https://doi.org/10.1023/A:1017515129544).
- [25] J. K. Hale and S. M. V. Lunel, "Averaging in infinite dimensions," *J. Integral Equ. Appl.*, vol. 2, no. 4, pp. 463–494, 1990, doi: [10.1216/jiea/1181075583](https://doi.org/10.1216/jiea/1181075583).
- [26] Z. Han, D. Niyato, W. Saad, and T. Başar, *Game Theory for Next Generation Wireless and Communication Networks: Modeling, Analysis, and Design*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [27] A. Hegyi, B. De Schutter, and H. Hellendoorn, "Model predictive control for optimal coordination of ramp metering and variable speed limits," *Transp. Res. Part C: Emerg. Technol.*, vol. 13, no. 3, pp. 185–209, 2005, doi: [10.1016/j.trc.2004.08.001](https://doi.org/10.1016/j.trc.2004.08.001).
- [28] C. D. Hill, "Parabolic equations in one space variable and the non-characteristic Cauchy problem," *Commun. Pure Appl. Math.*, vol. 20, no. 3, pp. 619–633, 1967, doi: [10.1002/cpa.3160200309](https://doi.org/10.1002/cpa.3160200309).
- [29] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, U.K.: Cambridge Univ. Press, 1985.
- [30] N. Hudon, M. Guay, M. Perrier, and D. Dochain, "Adaptive extremum-seeking control of convection-reaction distributed reactor with limited actuation," *Comput. Chem. Eng.*, vol. 32, no. 12, pp. 2994–3001, Dec. 2008, doi: [10.1016/j.compchemeng.2008.03.012](https://doi.org/10.1016/j.compchemeng.2008.03.012).
- [31] O. C. Imer, S. Compans, T. Başar, and R. Srikant, "Available bit rate congestion control in ATM networks," *IEEE Control Syst. Mag.*, vol. 21, no. 1, pp. 38–56, Feb. 2001, doi: [10.1109/37.898791](https://doi.org/10.1109/37.898791).
- [32] I. Karafyllis and M. Krstić, *Input-to-State Stability for PDEs*. Cham, Switzerland: Springer-Verlag, 2018.
- [33] I. Karafyllis and M. Papageorgiou, "Feedback control of scalar conservation laws with application to density control in freeways by means of variable speed limits," *Automatica*, vol. 105, pp. 228–236, Jul. 2019, doi: [10.1016/j.automatica.2019.03.021](https://doi.org/10.1016/j.automatica.2019.03.021).
- [34] B. Kaskosz and T. Tadumadze, "A differential game of evasion with delays," *J. Optim. Theory Appl.*, vol. 44, no. 2, pp. 231–268, 1984, doi: [10.1007/BF00935438](https://doi.org/10.1007/BF00935438).
- [35] N. Killingsworth and M. Krstić, "PID tuning using extremum seeking," *IEEE Control Syst. Mag.*, vol. 26, no. 1, pp. 70–79, Feb. 2006, doi: [10.1109/MCS.2006.1580155](https://doi.org/10.1109/MCS.2006.1580155).
- [36] S. Koga, M. Diagne, and M. Krstić, "Control and state estimation of the one-phase Stefan problem via backstepping design," *IEEE Trans. Autom. Control*, vol. 64, no. 2, pp. 510–525, Feb. 2019, doi: [10.1109/TAC.2018.2836018](https://doi.org/10.1109/TAC.2018.2836018).
- [37] S. Koga and M. Krstić, *Materials Phase Change PDE Control & Estimation: From Additive Manufacturing to Polar Ice*. Cham, Switzerland: Birkhäuser, 2020.
- [38] S. Koga and M. Krstić, "Control of Stefan system and applications: A tutorial," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 5, no. 1, pp. 547–577, 2022, doi: [10.1146/annurev-control-042920-014825](https://doi.org/10.1146/annurev-control-042920-014825).
- [39] S. Koga and M. Krstić, "Safe PDE backstepping QP control with high relative degree CBFs: Stefan model with actuator dynamics," *IEEE Trans. Autom. Control*, vol. 68, no. 12, pp. 7195–7208, Dec. 2023, doi: [10.1109/TAC.2023.3250514](https://doi.org/10.1109/TAC.2023.3250514).
- [40] M. Krstić, *Delay Compensation for Nonlinear, Adaptive, and PDE Systems*. Boston, MA, USA: Birkhäuser, 2009.
- [41] M. Krstić, "Compensation of infinite-dimensional actuator and sensor dynamics," *IEEE Control Syst. Mag.*, vol. 30, no. 1, pp. 22–41, Feb. 2010, doi: [10.1109/MCS.2009.934990](https://doi.org/10.1109/MCS.2009.934990).
- [42] M. Krstić and A. Smyshlyayev, *Boundary Control of PDEs: A Course on Backstepping Designs*. Philadelphia, PA, USA: SIAM, 2008.
- [43] M. Krstić and H. H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, no. 4, pp. 595–601, 2000, doi: [10.1016/S0005-1098\(99\)00183-1](https://doi.org/10.1016/S0005-1098(99)00183-1).
- [44] B. Laroche, P. Martin, and P. Rouchon, "Motion planning for a class of partial differential equations with boundary control," in *Proc. 37th IEEE Conf. Decis. Control (Cat. No.98CH36171)*, 1998, vol. 3, pp. 3494–3497, doi: [10.1109/CDC.1998.758247](https://doi.org/10.1109/CDC.1998.758247).
- [45] M. Lewicka and J. J. Manfredi, "Game theoretical methods in PDEs," *Bollettino dell'Unione Matematica Italiana*, vol. 7, no. 3, pp. 211–216, 2014, doi: [10.1007/s40574-014-0011-z](https://doi.org/10.1007/s40574-014-0011-z).

- [46] T. Li, G. Peng, Q. Zhu, and T. Başar, "The confluence of networks, games, and learning a game-theoretic framework for multiagent decision making over networks," *IEEE Control Syst. Mag.*, vol. 42, no. 4, pp. 35–67, Aug. 2022, doi: [10.1109/MCS.2022.3171478](https://doi.org/10.1109/MCS.2022.3171478).
- [47] X. Y. Lu, P. Varaiya, R. Horowitz, D. Su, and S. E. Shladover, "Novel free-way traffic control with variable speed limit and coordinated ramp metering," *Transp. Res. Rec.*, vol. 2229, no. 1, pp. 55–65, 2011, doi: [10.3141/2229-07](https://doi.org/10.3141/2229-07).
- [48] M. Merad, R. J. Downey, S. Obuz, and W. E. Dixon, "Isometric torque control for neuromuscular electrical stimulation with time-varying input delay," *IEEE Trans. Control Syst. Technol.*, vol. 24, no. 3, pp. 971–978, May 2016, doi: [10.1109/TCST.2015.2470637](https://doi.org/10.1109/TCST.2015.2470637).
- [49] K. Mori and E. Shimemura, "Linear differential games with delayed and noisy information," *J. Optim. Theory Appl.*, vol. 13, no. 3, pp. 275–289, 1974, doi: [10.1007/BF00934865](https://doi.org/10.1007/BF00934865).
- [50] J. F. Nash, "Noncooperative games," *Am. Math.*, vol. 54, no. 2, pp. 286–295, 1951.
- [51] T. R. Oliveira, J. P. V. S. Cunha, and L. Hsu, "Adaptive sliding mode control based on the extended equivalent control concept for disturbances with unknown bounds," in *Advances in Variable Structure Systems and Sliding Mode Control—Theory and Applications*, vol. 115, S. Li, X. Yu, L. Fridman, Z. Man, and X. Wang, Eds. Cham, Switzerland: Springer-Verlag, 2018, pp. 149–163, doi: [10.1007/978-3-319-62896-7_6](https://doi.org/10.1007/978-3-319-62896-7_6).
- [52] T. R. Oliveira, J. Feiling, and M. Krstić, "Extremum seeking for maximizing higher derivatives of unknown maps in cascade with reaction-advection-diffusion PDEs," *IFAC-PapersOnLine*, vol. 52, no. 29, pp. 210–215, 2019, doi: [10.1016/j.ifacol.2019.12.651](https://doi.org/10.1016/j.ifacol.2019.12.651).
- [53] T. R. Oliveira, J. Feiling, S. Koga, and M. Krstić, "Multivariable extremum seeking for PDE dynamic systems," *IEEE Trans. Autom. Control*, vol. 65, no. 11, pp. 4949–4956, Nov. 2020, doi: [10.1109/TAC.2020.3005177](https://doi.org/10.1109/TAC.2020.3005177).
- [54] T. R. Oliveira, J. Feiling, S. Koga, and M. Krstić, "Extremum seeking for unknown scalar maps in cascade with a class of parabolic partial differential equations," *Int. J. Adapt. Control Signal Process.*, vol. 35, no. 7, pp. 1162–1187, Jul. 2021, doi: [10.1002/acs.3117](https://doi.org/10.1002/acs.3117).
- [55] T. R. Oliveira and M. Krstić, "Compensation of wave PDEs in actuator dynamics for extremum seeking feedback," *IFAC-PapersOnLine*, vol. 52, no. 29, pp. 134–139, 2019, doi: [10.1016/j.ifacol.2019.12.634](https://doi.org/10.1016/j.ifacol.2019.12.634).
- [56] T. R. Oliveira and M. Krstić, "Extremum seeking feedback with wave partial differential equation compensation," *J. Dyn. Syst. Meas. Contr.*, vol. 143, no. 4, 2021, Art. no. 041002, doi: [10.1115/1.4048586](https://doi.org/10.1115/1.4048586).
- [57] T. R. Oliveira and M. Krstić, "Extremum seeking boundary control for PDE–PDE cascades," *Syst. Control Lett.*, vol. 155, Sep. 2021, Art. no. 105004, doi: [10.1016/j.sysconle.2021.105004](https://doi.org/10.1016/j.sysconle.2021.105004).
- [58] T. R. Oliveira and M. Krstić, *Extremum Seeking through Delays and PDEs*. Philadelphia, PA, USA: SIAM, 2022.
- [59] T. R. Oliveira, M. Krstić, and D. Tsubakino, "Extremum seeking for static maps with delays," *IEEE Trans. Autom. Control*, vol. 62, no. 4, pp. 1911–1926, Apr. 2017, doi: [10.1109/TAC.2016.2564958](https://doi.org/10.1109/TAC.2016.2564958).
- [60] T. R. Oliveira, V. H. P. Rodrigues, M. Krstić, and T. Başar, "Nash equilibrium seeking in quadratic noncooperative games under two delayed information-sharing schemes," *J. Optim. Theory Appl.*, vol. 191, no. 2–3, pp. 700–735, 2021, doi: [10.1007/s10957-020-01757-z](https://doi.org/10.1007/s10957-020-01757-z).
- [61] T. R. Oliveira, V. H. P. Rodrigues, M. Krstić, and T. Başar, "Nash equilibrium seeking with arbitrarily delayed player actions," in *Proc. IEEE Conf. Decis. Control (CDC)*, Jeju, South Korea, 2020, pp. 150–155, doi: [10.1109/CDC42340.2020.9303894](https://doi.org/10.1109/CDC42340.2020.9303894).
- [62] T. R. Oliveira, V. H. P. Rodrigues, M. Krstić, and T. Başar, "Nash equilibrium seeking with players acting through heat PDE dynamics," in *Proc. IEEE Amer. Control Conf. (ACC)*, New Orleans, LA, USA, 2021, pp. 684–689, doi: [10.23919/ACC50511.2021.9483114](https://doi.org/10.23919/ACC50511.2021.9483114).
- [63] T. R. Oliveira, V. H. P. Rodrigues, M. Krstić, and T. Başar, "Nash equilibrium seeking in heterogeneous noncooperative games with players acting through heat PDE dynamics and delays," in *Proc. IEEE Conf. Decis. Control*, Austin, TX, USA, 2021, pp. 1167–1173, doi: [10.1109/CDC45484.2021.9683539](https://doi.org/10.1109/CDC45484.2021.9683539).
- [64] O. M. Pamen, "Optimal control for stochastic delay systems under model uncertainty: A stochastic differential game approach," *J. Optim. Theory Appl.*, vol. 167, no. 3, pp. 998–1031, 2015.
- [65] P. Paz, T. R. Oliveira, A. V. Pino, and A. P. Fontana, "Model-free neuromuscular electrical stimulation by stochastic extremum seeking," *IEEE Trans. Control Syst. Technol.*, vol. 28, no. 1, pp. 238–253, Jan. 2020, doi: [10.1109/TCST.2019.2892924](https://doi.org/10.1109/TCST.2019.2892924).
- [66] J. I. Poveda, M. Krstić, and T. Başar, "Fixed-time Nash equilibrium seeking in time-varying networks," *IEEE Trans. Autom. Control*, vol. 68, no. 4, pp. 1954–1969, Apr. 2023, doi: [10.1109/TAC.2022.3168527](https://doi.org/10.1109/TAC.2022.3168527).
- [67] V. H. P. Rodrigues, T. R. Oliveira, L. Hsu, M. Diagne, and M. Krstić, "Event-triggered and periodic even-triggered extremum seeking control," *Automatica*, vol. 174, Apr. 2025, Art. no. 112161, doi: [10.1016/j.automatica.2025.112161](https://doi.org/10.1016/j.automatica.2025.112161).
- [68] E. O. Roxin, "Differential games with partial differential equations," in *Proc. Theory Appl. Differential Games*, Dordrecht, The Netherlands: Springer-Verlag, 1975, pp. 157–168.
- [69] D. Rusiti, T. R. Oliveira, M. Krstić, and M. Gerdtts, "Newton-based extremum seeking of higher-derivative maps with time-varying delays," *Int. J. Adapt. Control Signal Process.*, vol. 35, no. 7, pp. 1202–1216, 2021.
- [70] D. Rusiti, T. R. Oliveira, M. Krstić, and M. Gerdtts, "Robustness to delay mismatch in extremum seeking," *Eur. J. Control*, vol. 62, pp. 75–83, Nov. 2021, doi: [10.1016/j.ejcon.2021.06.029](https://doi.org/10.1016/j.ejcon.2021.06.029).
- [71] N. Sharma, C. Gregory, and W. E. Dixon, "Predictor-based compensation for electromechanical delay during neuromuscular electrical stimulation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 19, no. 6, pp. 601–611, Dec. 2011, doi: [10.1109/TNSRE.2011.2166405](https://doi.org/10.1109/TNSRE.2011.2166405).
- [72] P. C. S. Silva, P. C. Pellanda, T. R. Oliveira, G. A. de Andrade, and M. Krstić, "Extremum seeking for a class of wave partial differential equations with Kelvin-Voigt damping," *IEEE Contr. Syst. Lett.*, vol. 8, pp. 43–48, 2024, doi: [10.1109/LCSYS.2023.3346314](https://doi.org/10.1109/LCSYS.2023.3346314).
- [73] A. Siranosian, M. Krstić, A. Smyshlyayev, and M. Bement, "Motion planning and tracking for tip displacement and deflection angle for flexible beams," *J. Dyn. Syst. Meas. Contr.*, vol. 131, no. 3, pp. 1–10, 2009, doi: [10.1115/1.13072152](https://doi.org/10.1115/1.13072152).
- [74] A. Smyshlyayev and M. Krstić, *Adaptive Control of Parabolic PDEs*. Philadelphia, PA, USA: SIAM, 2010.
- [75] M. Treiber and A. Kesting, *Traffic Flow Dynamics*. Berlin, Heidelberg, Germany: Springer-Verlag, 2013.
- [76] D. Tsubakino, T. R. Oliveira, and M. Krstić, "Extremum seeking under distributed input delay," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 5423–5428, 2020, doi: [10.1016/j.ifacol.2020.12.1538](https://doi.org/10.1016/j.ifacol.2020.12.1538).
- [77] D. Tsubakino, T. R. Oliveira, and M. Krstić, "Extremum seeking for distributed delays," *Automatica*, vol. 153, Jul. 2023, Art. no. 111044, doi: [10.1016/j.automatica.2023.111044](https://doi.org/10.1016/j.automatica.2023.111044).
- [78] A. Williams, M. Krstić, and A. Scheinker, "Semiglobal safety-filtered extremum seeking with unknown CBFs," *IEEE Trans. Autom. Control*, vol. 70, no. 3, pp. 1698–1713, Mar. 2025, doi: [10.1109/TAC.2024.3469785](https://doi.org/10.1109/TAC.2024.3469785).
- [79] J. J. Winkin, D. Dochain, and P. Ligarius, "Dynamical analysis of distributed parameter tubular reactors," *Automatica*, vol. 36, no. 3, pp. 349–361, 2000, doi: [10.1016/S0005-1098\(99\)00170-3](https://doi.org/10.1016/S0005-1098(99)00170-3).
- [80] X. Yang and E. Fridman, "Extremum seeking in the presence of large delays via time-delay approach to averaging," 2023, *arXiv:2310.09474*.
- [81] C. T. Yilmaz, M. Diagne, and M. Krstić, "Exponential and prescribed-time extremum seeking with unbiased convergence," *Automatica*, vol. 179, Sep. 2025, Art. no. 112392.
- [82] C. T. Yilmaz and M. Krstić, "Prescribed-time extremum seeking for delays and PDEs using chirpy probing," *IEEE Trans. Autom. Control*, vol. 69, no. 11, pp. 7710–7725, Nov. 2024, doi: [10.1109/TAC.2024.3397162](https://doi.org/10.1109/TAC.2024.3397162).
- [83] H. Yu, M. Diagne, L. Zhang, and M. Krstić, "Bilateral boundary control of moving shockwave in LWR model of congested traffic," *IEEE Trans. Autom. Control*, vol. 66, no. 3, pp. 1429–1436, Mar. 2021, doi: [10.1109/TAC.2020.2994031](https://doi.org/10.1109/TAC.2020.2994031).
- [84] H. Yu, S. Koga, T. R. Oliveira, and M. Krstić, "Extremum seeking for traffic congestion control with a downstream bottleneck," *J. Dyn. Syst. Meas. Contr.*, vol. 143, no. 3, pp. 1–10, 2021, doi: [10.1115/1.4048781](https://doi.org/10.1115/1.4048781).
- [85] H. Yu and M. Krstić, "Traffic congestion control for Aw–Rascle–Zhang model," *Automatica*, vol. 100, pp. 38–51, Feb. 2019, doi: [10.1016/j.automatica.2018.10.040](https://doi.org/10.1016/j.automatica.2018.10.040).
- [86] Y. Zhang and P. A. Ioannou, "Coordinated variable speed limit, ramp metering and lane change control of highway traffic," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 5307–5312, 2017, doi: [10.1016/j.ifacol.2017.08.982](https://doi.org/10.1016/j.ifacol.2017.08.982).
- [87] Q. Zhu and T. Başar, "Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: Games-in-games principle for optimal cross-layer resilient control systems," *IEEE Control Syst. Mag.*, vol. 35, no. 1, pp. 46–65, Feb. 2015, doi: [10.1109/MCS.2014.2364710](https://doi.org/10.1109/MCS.2014.2364710).
- [88] Q. Zhu and T. Başar, "Disentangling resilience from robustness: Contextual dualism, interactionism, and game-theoretic paradigms," *IEEE Control Syst. Mag.*, vol. 44, no. 3, pp. 94–103, Jun. 2024, doi: [10.1109/MCS.2024.3382415](https://doi.org/10.1109/MCS.2024.3382415).
- [89] Y. Zhu, E. Fridman, and T. R. Oliveira, "Sampled-data extremum seeking with constant delay: A time-delay approach," *IEEE Trans. Autom. Control*, vol. 68, no. 1, pp. 432–439, 2023, doi: [10.1109/TAC.2022.3140259](https://doi.org/10.1109/TAC.2022.3140259).



SOURCE: GRAPHIC CREATED FROM AUTHOR-GENERATED ARTWORK USING OPENAI CHATGPT. FOR IMAGE CREDITS, PLEASE SEE THE "ACKNOWLEDGMENT" SECTION OF THE ARTICLE.

Hybrid Set-Seeking Systems

MODEL-FREE FEEDBACK OPTIMIZATION VIA HYBRID INCLUSIONS

JORGE I. POVEDA ^{ID} and ANDREW R. TEEL ^{ID}

Hybrid dynamical systems (HDSs), which combine continuous-time and discrete-time dynamics, are pervasive in modern engineering. They naturally arise in cyberphysical systems (CPSs) that integrate physical processes with digital components, spanning domains such as transportation [1], power grids [2], robotics [3], and power electronics [4]. Beyond CPSs, hybrid dynamics also appear in a wide range of applications, including mechanical systems with impacts [5], chemical process control [6], manufacturing systems [7], and synchronization mechanisms in biological systems [8, Ex. 1.2]. In the context of control and decision making, hybrid systems can model algorithms that involve logic modes, timers, resets, impulsive actions, sampled-data structures,

switching dynamics, event-triggered mechanisms, and general “if-then” rules [8], [9], [10]. Hybrid systems are also closely related to frameworks widely studied in computer science and engineering, such as the well-known hybrid automata [11], [12], [13].

Despite the prevalence of hybrid systems in engineering and science, the development of extremum-seeking (ES) methods for systems with hybrid dynamics remained relatively unexplored until recently (see “Summary”). ES systems are feedback control mechanisms designed to achieve real-time steady-state optimization in dynamical systems where both the plant dynamics and the cost function are unknown. This “model-free” characteristic makes ES an adaptable and powerful feedback control scheme for various complex engineering applications where precise mathematical models are unavailable. As a result, ES has become widely adopted across different industries and areas,

Digital Object Identifier 10.1109/MCS.2025.3646104
Date of current version: 26 March 2026

including engine optimization [14], wind turbine control [15], underwater vehicles [16], [17], anti-lock braking systems [18], autotuning systems [19], photolithography in semiconductor manufacturing [20], fuel cells [21], the control of mobile robots [22], heating, ventilation, and air-conditioning (HVAC) systems [23], model predictive control [24], biological systems [25], traffic control [26], [27], mechanical systems [28], [29], [30], sampled-data systems [31], [32], etc. For surveys on ES, we refer the reader to [33], [34], and [35] as well as to [36, Sec. 13.3].

INTRODUCTION

The idea of using feedback control mechanisms for real-time model-free optimization in dynamical systems dates back to the early 1920s with LeBlanc's work [37] and was further advanced in the 1950s [38] and 1960s [39], [40], [41], [42] in the adaptive controls literature under the name of "extremum control" [36, Sec. 13.3]. However, most of the stability results for ES in nonlinear systems emerged in the late 1990s and early 2000s, utilizing multitime-scale methods for ordinary differential equations (ODEs) and identifying the appropriate timescales associated with the components of the algorithms [44]. Since then, different analytical tools for ES have been developed to solve model-free optimization and decision-making problems in single-agent systems [45], [46], [47], [48], [49], [50], multiagent systems [51], [52], [53], partial differential equations [54], stochastic systems [55], and more.

As the foundational theory behind ES relied on perturbation tools for smooth ODEs, particularly averaging [56], [57] and singular perturbations [57], [58], [59], the majority of the ES systems and algorithms were designed for applications with closed-loop smooth dynamics. In particular, the study of ES systems that incorporate continuous-time and discrete-time behaviors—such as "events" triggered by specific conditions, abrupt state "jumps" or "resets" aimed at enhancing performance, or algorithms that "switch" between different dynamics during the seeking process—has received comparatively less attention. Such systems, which are naturally hybrid, also arise across different applications where either the plant is a hybrid system or the controller must be hybrid to address the fundamental limitations of traditional smooth control techniques. These limitations can be related to achieving desired transient or steady-state performance or handling complexities such as dynamic communication networks, resource-constrained actuators, occluded operational spaces, or dynamic adversarial environments. ES systems with discrete states, such as logic modes, also arise when general "if-then" rules are incorporated into the decision-making algorithms, which is now a standard practice in the design and implementation of controllers and optimization algorithms for CPSs [61] and autonomous systems [60]. Moreover, many biological systems also exhibit behaviors that can be modeled as hybrid seeking dynamics, wherein agents switch between distinct exploratory and exploitative

algorithms in response to state-dependent cues, as observed in raptor hunting, bacterial chemotaxis, social insect foraging, and vertebrate motor learning.

Despite the limited available theoretical tools for their analysis and design, ES algorithms with logic states and switching rules have been explored for particular

Summary

This article focuses on the analysis and synthesis of extremum-seeking (ES) systems that incorporate *hybrid dynamics* in the loop, encompassing both continuous-time and discrete-time behaviors. Such ES systems arise from the need to use advanced control and optimization algorithms to overcome the limitations of standard smooth feedback techniques and to meet stringent transient and steady-state demands in high-performance applications. They also emerge in settings where traditional controllers must be implemented on plants that naturally exhibit hybrid behaviors, such as in cyberphysical systems (CPSs) and autonomous systems, which rely on digital devices for computation, actuation, and sensing.

To explore hybrid ES dynamics using control-theoretic tools, we begin by reviewing the key technical concepts that enable the development of perturbation theory for hybrid inclusions. These concepts provide the foundations for both averaging and singular perturbation theory in these systems. We then demonstrate how these tools can be applied to the analysis and synthesis of various hybrid ES algorithms for plants modeled as static maps. In particular, we present several examples of set-valued and switching ES algorithms under various switching scenarios: arbitrarily fast, under dwell-time and average dwell-time constraints, and under average activation-time constraints. Additionally, we explore state-based switching ES in the context of obstacle avoidance problems as well as gradient-Newton switching ES schemes. Other topics include ES with momentum and resets, with intermittent updates, slowly varying parameters, and using safety mechanisms to incorporate constraints. In all these applications, we show how the perturbation-based approach, traditionally used for studying smooth ES control, can naturally be extended to hybrid systems, provided that they satisfy suitable mild regularity conditions and that their solutions are modeled in a hybrid time domain. Finally, in Part 3, we address dynamic plants and discuss how the proposed hybrid ES controllers can be integrated with such plants to solve model-free feedback optimization and decision-making problems. Building on established results for ES and using various illustrative examples, this article aims to provide an introductory and tutorial-style presentation of hybrid ES systems, making it accessible to readers from different backgrounds, including those with no prior experience in hybrid systems.

applications since the early 1970s [62], when a switching hysteresis-based mechanism was designed to stabilize a peak-holding adaptive controller in electrical machines. More modern approaches to ES that incorporate discrete-time behaviors via switches or resets have also been studied for the efficient control of photovoltaic power systems [63] and HVAC systems in industry [23]. However, a comprehensive theory for the study of hybrid ES systems has lagged behind their practical applications despite their potential to enable the design of novel feedback architectures and algorithms capable of addressing a broader class of problems than those handled by traditional ES methods.

As discussed later in the article, ES systems with hybrid dynamics typically exhibit stability properties with respect to general *sets* rather than isolated equilibria. We therefore refer to these controllers as *hybrid set-seeking systems*, which are generally characterized by nonsmooth dynamics modeled by differential and difference inclusions rather than by standard differential or difference equations. Building on this perspective, the primary aim of this article is to provide an accessible introduction to nonsmooth and hybrid set-seeking systems formulated as hybrid inclusions. We begin with well-established smooth ES algorithms extensively studied in the literature [44], [64], [65], [66], [67], which can be viewed as hybrid set-seeking schemes with negligible discrete-time dynamics. We then extend this perspective by incorporating hybrid dynamics through averaging tools, enabling the solution of complex model-free optimization and decision-making problems in static maps. Next, we show how singular perturbation methods for hybrid systems can address seeking problems in plants governed by dynamical systems rather than static maps. Taken together, these developments yield perturbation-based tools for multitime-scale hybrid inclusions, providing a unified framework for analyzing and designing hybrid seeking algorithms that overcome the limitations of traditional smooth ES methods. The effectiveness of this framework is demonstrated through diverse applications and numerical examples, showcasing the flexibility and robustness of the proposed tools.

Perturbation Theory and Seeking Systems

The key mathematical property that enables the design and analysis of most ES algorithms modeled as oscillatory, locally Lipschitz ODEs of the form

$$\dot{x} = f(x, \omega t), \quad x \in \mathbb{R}^n, \quad t \in \mathbb{R}_{\geq 0} \quad (1)$$

is that, for a sufficiently large oscillation frequency ω , the solutions of (1) can be approximated on arbitrarily large finite intervals by those of the so-called *average dynamics*

$$\dot{\bar{x}} = f_{\text{ave}}(\bar{x}) := \frac{1}{T_\omega} \int_0^{T_\omega} f(\bar{x}, \omega t) dt, \quad \bar{x} \in \mathbb{R}^n \quad (2)$$

where $T_\omega = 2\pi/\omega$ is the oscillation period of $f(x, \cdot)$. For example, to minimize an unknown smooth scalar static map $J: \mathbb{R} \rightarrow \mathbb{R}$ representing the plant, the simplest ES continuous-time algorithm based on periodic probing, studied in [44], [65], and [66], can be modeled by the following equations:

$$u = \hat{u} + \varepsilon_a \sin(\omega t) \quad (3a)$$

$$\dot{\hat{u}} = -\frac{2}{\varepsilon_a} J(u) \sin(\omega t) \quad (3b)$$

where $\varepsilon_a > 0$ and $\omega > 0$ are tunable parameters. A Taylor expansion of J around \hat{u} for small values of ε_a and a direct computation of (2) using the identities $\int_0^{T_\omega} \sin(\omega t) dt = 0$ and $(1/T_\omega) \int_0^{T_\omega} \sin(\omega t)^2 dt = 1/2$ show that the average dynamics of (3) correspond to the following ε_a -perturbed system (see Example 8), with state $\bar{u} \in \mathbb{R}$:

$$\dot{\bar{u}} = f_{\text{ave}}(\bar{u}) = -\frac{\partial J(\bar{u})}{\partial \bar{u}} + O(\varepsilon_a) \quad (4)$$

where $O(\varepsilon_a)$ indicates bounded perturbations (on compact sets) of order ε_a . In turn, as $\varepsilon_a \rightarrow 0^+$, the behavior of system (4) can be approximated by a standard gradient flow of J , with state $\bar{u} \in \mathbb{R}$ and dynamics

$$\dot{\bar{u}} = -\frac{\partial J(\bar{u})}{\partial \bar{u}}. \quad (5)$$

Under suitable regularity (for example, locally Lipschitz properties in system (5)), we can use standard robustness results for differential equations to assert that the solutions of (4) and (5) are (T, ϵ) -close when ε_a is sufficiently small [57], [68]. Namely, for each compact set of initial conditions, each $T > 0$, and each $\epsilon > 0$, there exists $\varepsilon_a^* > 0$ such that for all $\varepsilon_a \in (0, \varepsilon_a^*)$, the solutions of (4) and (5) satisfy the bound

$$\sup_{t \in [0, T]} |\bar{u}(t) - \hat{u}(t)| \leq \frac{\epsilon}{2}. \quad (6)$$

Similarly, by leveraging results on averaging theory for differential equations [59] (see “Averaging Theory: From ODEs to Hybrid Systems”), it can be shown that, for a fixed and sufficiently small $\varepsilon_a > 0$, the solutions of (3) and (4) are also (T, ϵ) -close when ω is sufficiently large; that is, for each compact set of initial conditions, each $T > 0$, and each $\epsilon > 0$, there exists a sufficiently large frequency $\omega^* > 0$ such that for all $\omega > \omega^*$, the solutions of (3) and (4) satisfy the bound

$$\sup_{t \in [0, T]} |\hat{u}(t) - \bar{u}(t)| \leq \frac{\epsilon}{2}. \quad (7)$$

By combining the bounds (6) and (7), we can conclude that the trajectories of the ES dynamics (3) and those of the gradient flow (5) are also (T, ϵ) -close in compact sets. This property immediately allows us to obtain qualitative attributes for the behavior of the solutions of the model-free ES dynamics (3) based on the behavior of the solutions of the model-based gradient flow (5), provided the parameters (ω, ε_a) are appropriately tuned. Among such attributes, we

can infer stability properties for system (3) based on the stability properties of (5), provided that such stability properties are “robust” to small perturbations, a property that, under the Lipschitz continuity of (5), can be readily established using, for example, Lyapunov functions. This perturbation-based methodology for the study of smooth ES lies at the heart of the design and stability analysis of most algorithms, including those whose average dynamics (2) are not simple perturbed gradient flows but instead follow alternative well-posed optimization dynamics, such as primal-dual dynamics [69], projected gradient flows [70], Riemannian gradient flows [71], Newton flows [72], pseudo-gradient flows for games [51], decentralized optimization dynamics [73], etc.; see [65] for a unifying framework.

Fortunately, the perturbation methodology described previously extends to systems with hybrid dynamics under the appropriate mild regularity conditions and with an appropriate measure of closeness. For the latter, we note that, unlike smooth ODEs, solutions to hybrid systems are typically discontinuous, making the use of the uniform distance, as in (6) and (7), to measure the closeness between solutions of nominal and perturbed systems problematic. For instance, if \hat{u} is a solution to a hybrid system that has a discontinuity at time t' , and \bar{u} is a solution to an ε_a -perturbed version of the same system that has a discontinuity at time $t' + \varepsilon_a$, then no matter how small ε_a is, \hat{u} and \bar{u} will never be (T, ϵ) -close at time t' , regardless of the choice of ε_a . Thus, a different measure of closeness is required to link the qualitative behaviors of the perturbed and unperturbed systems. Although this issue is well known in the hybrid systems literature [74, Example 3], we will illustrate this phenomenon in a stylized hybrid set-seeking system that exhibits “jumps” in the states, resulting in Zeno-like behavior reminiscent of a bouncing ball without air resistance.

Generalized Solutions for Hybrid Seeking Systems

To extend the perturbation-based methodology to seeking systems with hybrid dynamics, we must generalize both the mathematical models we consider and their corresponding notions of solutions. Specifically, in studying hybrid seeking systems, we will parameterize solutions using a continuous-time index $t \in \mathbb{R}_{\geq 0}$ and a discrete-time index $j \in \mathbb{Z}_{\geq 0}$. The primary benefit of this parameterization is that it allows solutions to be defined in *hybrid time domains*. For such solutions, the concept of (T, ϵ) -closeness can be naturally extended by measuring the distance between the *graphs* of the nominal and perturbed solutions, that is, by using set convergence (between graphs) as opposed to the standard uniform convergence. This can be achieved, provided that the hybrid dynamics exhibit certain mild regularity properties. We will illustrate these regularity properties through various examples of hybrid seeking systems with different types of hybrid time domains. By defining solutions in such domains and employing graphical convergence to analyze

the limiting behavior of perturbed solutions, we can seamlessly extend the perturbation-based methodology traditionally used in ES control [45], [65], [66] to general hybrid set-seeking systems.

Seeking Systems Modeled as Hybrid Inclusions

Most ES dynamics studied via averaging are modeled by ODEs of the form (1). However, it is also useful to generalize these models for the study of hybrid seeking systems. For instance, hybrid systems often exhibit discontinuities, or “jumps,” in the state, governed by a set-valued update rule of the form $x^+ \in G(x)$, where x^+ denotes the state value immediately after the jump. This type of update rule, known as a “difference inclusion,” allows x^+ to take any value within the set $G(x)$, which may contain infinitely many points for a given x . These models are powerful because they can capture more general and flexible behaviors and decision-making rules compared to standard difference equations of the form $x^+ = g(x)$. Similarly, hybrid ES algorithms may also involve continuous-time dynamics described by differential inclusions of the form $\dot{x} \in F(x)$, which can arise from discontinuous vector fields requiring appropriate regularizations (for example, Krasovskii, Filippov, etc.) for analysis or from optimization and decision-making algorithms that are inherently set-valued [75].

By incorporating set-valued updates in the continuous-time and/or discrete-time dynamics, hybrid inclusions can also provide a robust mathematical framework for modeling and analyzing seeking algorithms that are influenced by external discontinuous signals that govern system behavior. Such signals can arise in a variety of contexts, including switching dynamics [76], time-varying optimization problems [77], event-triggered systems [9], and systems with resetting timers [78], to name just a few examples. As such, most of the models in our study of hybrid set-seeking algorithms will be represented using hybrid inclusions.

The remainder of this article is divided into three main sections. In Part 1, we present the preliminaries on hybrid systems and introduce the mathematical language necessary to study hybrid set-seeking systems. Part 2 focuses on hybrid seeking algorithms for plants characterized by static maps, while Part 3 extends these results to systems with dynamic plants. The main developments are complemented by several supporting discussions. “[Continuous-Time Set-Valued Dynamical Systems](#)” and “[Discrete-Time Set-Valued Dynamical Systems](#)” offer brief introductions to dynamical systems modeled as differential and difference inclusions, respectively. “[Averaging Theory: From ODEs to Hybrid Systems](#)” provides a mathematical overview of averaging theory applicable to hybrid set-seeking systems, and “[Singularly Perturbed HDSSs](#)” introduces the mathematical tools needed for studying hybrid systems in singular perturbation form. The notation used throughout this article is defined in [Table 1](#).

PART 1: PRELIMINARIES IN HYBRID SYSTEMS

HDSs, or simply “hybrid systems,” are dynamical systems that combine continuous-time and discrete-time phenomena. Therefore, a natural mathematical representation of such systems would involve both differential equations and discrete-time recursions. In this article, we follow the formalism presented in [8] to study general hybrid systems; see also [74] for a tutorial dedicated exclusively to such systems. In its simplest form, a hybrid system with state $x \in \mathbb{R}^n$, can be represented by the following expressions:

$$x \in C, \quad \dot{x} = f(x) \quad (8a)$$

$$x \in D, \quad x^+ = g(x) \quad (8b)$$

where (8a) describes the continuous-time dynamics, and (8b) describes the discrete-time dynamics. The notation in (8a) indicates that the state x is allowed to evolve according to the continuous-time dynamics $\dot{x} = f(x)$ whenever the

condition $x \in C$ holds, where the set $C \subset \mathbb{R}^n$ is called the *flow set*. Similarly, the notation in (8b) indicates that x is allowed to evolve according to the discrete-time dynamics $x^+ = g(x)$ whenever the condition $x \in D$ holds, where the set $D \subset \mathbb{R}^n$ is called the *jump set*, and where x^+ denotes the value of x after an “instantaneous” change—or jump—in the state. For consistency, we call the functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ the *flow map* and the *jump map*, respectively. In this way, the hybrid system (8) can be represented by the tuple $\mathcal{H} = \{f, C, g, D\}$, which we call the *data* of the system. Note that when $D = \emptyset$, the system (8) recovers a constrained continuous-time dynamical system. Similarly, when $C = \emptyset$, the system (8) recovers a constrained discrete-time dynamical system. Figure 1 shows an abstract representation of a possible “solution” to (8), starting from x_0 and experiencing an initial interval of flow, followed by two consecutive jumps and a final interval of flow that eventually hits the boundary of C , forcing the solution to stop. We will rigorously formalize this notion of “solution” and illustrate its application to study hybrid set-seeking systems.

TABLE 1 Notation used in the article.

Symbol	
\mathbb{R}	Set of real numbers
\mathbb{R}^n	The n -dimensional Euclidean space
$\mathbb{Z}_{\geq 0}$	Set of nonnegative integers
\dot{x}	Time derivative of the state of a system
x^+	State of a system after a jump
\emptyset	The empty set
$\bar{\mathcal{A}}$	The closure of the set \mathcal{A}
$\text{co}(\mathcal{A})$	The convex hull of the set \mathcal{A}
$\text{bd}(\mathcal{A})$	The boundary of the set \mathcal{A}
x^\top	The transpose of the vector x
(x, y)	Equivalent notation for $[x^\top, y^\top]^\top$
$ x $	Euclidean norm of the vector x
\mathbb{B}	The closed unit ball, of appropriate dimension, in the Euclidean norm
\mathbb{S}^1	The set $\{x = (x_1, x_2) \in \mathbb{R}^2 : x_1^2 + x_2^2 = 1\}$
\mathbb{S}^n	The n th-Cartesian product $\mathbb{S}^1 \times \dots \times \mathbb{S}^1$.
$f : \mathbb{R}^n \rightarrow \mathbb{R}^m$	A function from \mathbb{R}^n to \mathbb{R}^m
$F : \mathbb{R}^n \rightrightarrows \mathbb{R}^m$	A set-valued mapping from \mathbb{R}^n to \mathbb{R}^m
\mathbf{R}_o	The matrix $\mathbf{R}_o = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$
\mathcal{K}_∞	Class of functions $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that are zero at zero, continuous, strictly increasing, and unbounded
\mathcal{KL}	Class of functions $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that are nondecreasing in their first argument, nonincreasing in their second argument, and satisfy $\lim_{r \rightarrow 0^+} \beta(r, s) = 0$ for each $s \in \mathbb{R}_{\geq 0}$ and $\lim_{s \rightarrow \infty} \beta(r, s) = 0$ for each $r \in \mathbb{R}_{\geq 0}$

Example 1 (Systems with Periodic Resets)

To illustrate how the data in (8) can be used to model a common class of systems encountered in practice in control, such as systems with periodic resets, let $x = (x_1, x_2) \in \mathbb{R}^2$, and consider the following hybrid dynamics:

$$x \in C := [0, T] \times \mathbb{R}, \quad \dot{x} = \begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = f(x) := \begin{pmatrix} 1 \\ -\gamma x_2 \end{pmatrix}$$

$$x \in D := \{T\} \times \mathbb{R}, \quad x^+ = \begin{pmatrix} x_1^+ \\ x_2^+ \end{pmatrix} = g(x) := \begin{pmatrix} 0 \\ \frac{1}{2} x_2 \end{pmatrix}$$

where $T, \gamma > 0$ are tunable parameters. In this system, the state x_1 can be viewed as a timer that grows linearly during flows and is reset to zero whenever the condition $x_1 = T$ is satisfied. Whenever such resets occur, the state x_2 is also reset to one-half of its previous value. Since resets of x_1 force x to satisfy the conditions $x \in C$ and $x \notin D$, every jump in the system is followed by a period of flow of duration T , during which the continuous-time dynamics make x_2 decrease exponentially at a rate dictated by γ . Figure 2 shows the trajectories of x_1 and x_2 generated by this simple hybrid system, using $T = 10$, $\gamma = 1/10$, and from the initial conditions $(0, 10)$. The dashed lines indicate the jumps in the values of




FIGURE 1 An abstract representation of a solution to a hybrid system (8) from the initial point x_0 .

the states x triggered by the discrete-time dynamics $x^+ = g(x)$. Note that, for this system, the intersection of the flow set and the jump set is not empty. In fact, we have $D \cap C = \{T\}$. However, for initial conditions satisfying $x_1 = T$, the continuous-time dynamics $\dot{x}_1 = 1$ cannot flow and extend x_1 without leaving the set C ; that is, no solutions can evolve via flows from this initialization.

On the other hand, from this same initial condition, the state x_1 can be extended in time via jumps by immediately resetting x_1 and x_2 . This type of behavior will be formalized later when we discuss the notion of solutions to hybrid systems of the form (8). In the meantime, we note that,

while the state x_2 converges to zero in Figure 2, the state x_1 continues to oscillate restricted to the set $[0, T]$. Therefore, for this hybrid system, it is reasonable to study the stability properties with respect to the compact set $\mathcal{A} = [0, T] \times \{0\}$. This simple example serves to illustrate the fact that, unlike differential equations of the form (3), stability properties in hybrid seeking systems will be studied with respect to general sets, as opposed to equilibrium points. \square

Dynamic timers, such as those in Example 1, are common in hybrid systems to model time-triggered jumps, which can arise in sampled-data systems, switching algorithms, event-triggered control, and reset-based controllers, to name just a few examples. Incorporating exogenous timers into hybrid systems is also useful for studying time-varying dynamics, which are particularly relevant in ES systems, such as (3). In this case, the overall dynamics can be written as a hybrid system with states $(x, \tau) \in \mathbb{R}^n \times \mathbb{R}$ and dynamics given by

$$(x, \tau) \in C \times \mathbb{R}_{\geq 0}, \quad \dot{x} = f(x, \tau), \quad \dot{\tau} = \omega \quad (9a)$$

$$(x, \tau) \in D \times \mathbb{R}_{\geq 0}, \quad x^+ = g(x), \quad \tau^+ = \tau \quad (9b)$$

where $\omega > 0$ dictates the rate of evolution of τ . Note that (9) can be written as (8) using the extended state $y = (x, \tau)$.

The following example illustrates a class of time-varying hybrid systems with oscillating flow maps that are of interest in this article, which also exhibit a well-known behavior that can emerge in hybrid systems: Zeno phenomena.

Example 2 (“Bouncing” Seeking Systems)

Bouncing balls are classic examples of hybrid systems. Their dynamics are governed by Newtonian physics while the ball is in the air, and they experience instantaneous jumps upon hitting the ground. These models are most accurate when the ball is rigid, and the deformations caused by impacts are negligible. Inspired by classic bouncing ball

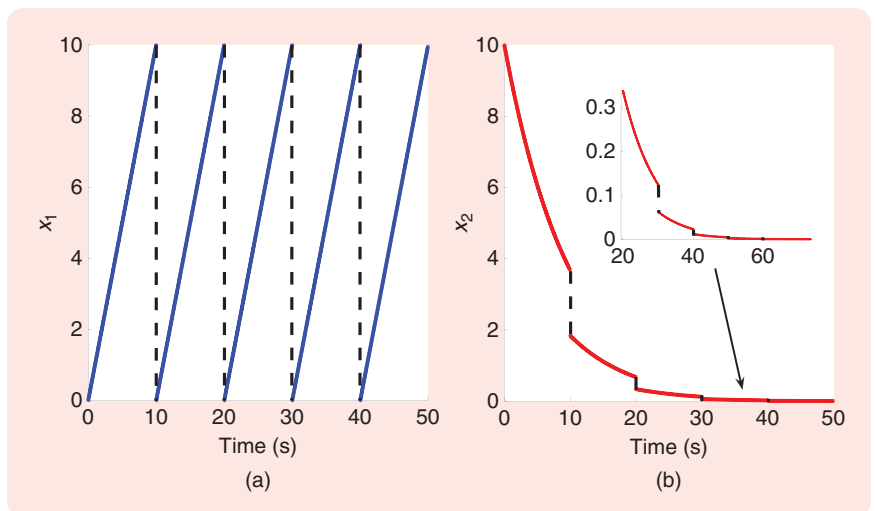


FIGURE 2 The trajectory generated by the hybrid system of Example 1, with the initial condition $x_0 = (0, 10)$. (a) The trajectory of x_1 . (b) The trajectory of x_2 .

systems [8, Ex. 1.1], we can consider a hybrid system of the form (9), with state $x = (x_1, x_2) \in \mathbb{R}^2$, where x_1 models the vertical position of a point-mass ball bouncing vertically on a horizontal surface and x_2 models the vertical velocity of the ball. Following the notation in (9), the data of this system are given by

$$C := \{(x_1, x_2) \in \mathbb{R}^2 : x_1 \geq 0\} \quad (10a)$$

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = f(x, \tau) := \begin{pmatrix} x_2 \\ -\gamma(\tau) \end{pmatrix}, \quad \dot{\tau} = \omega \quad (10b)$$

$$D = \{(x_1, x_2) \in \mathbb{R}^2 : x_1 = 0, x_2 \leq 0\} \quad (10c)$$

$$\begin{pmatrix} x_1^+ \\ x_2^+ \end{pmatrix} = g(x) := \begin{pmatrix} 0 \\ -\lambda x_2 \end{pmatrix}, \quad \tau^+ = \tau \quad (10d)$$

where the time-varying function $\gamma(\cdot)$ is given by

$$\gamma(\tau) = 2\bar{\gamma} \sin(\tau)^2 \quad (11)$$

and where $\bar{\gamma}$ denotes the nominal gravitational acceleration constant, and $\lambda \in (0, 1)$ is a restitution coefficient. According to (10a), whenever the altitude is nonnegative, the ball flows through the air via (10b) under oscillatory acceleration $-\gamma(\tau)$. On the other hand, according to (10c) and (10d), when the ball touches the ground (that is, $x_1 = 0$) and its velocity is nonpositive, the mapping g resets the sign of the velocity and also decreases its magnitude using the restitution coefficient λ . Note that the position of the ball does not change during jumps. Figure 3 illustrates the behavior of the solutions generated by this “bouncing” seeking system when $\omega = 100$ (colored trajectories) and also when $\gamma(\tau) = \bar{\gamma}$ is constant (black trajectories). It can be observed that the highly oscillatory hybrid dynamics that use $\omega = 100$ provide a suitable approximation of the behavior of the hybrid system that uses a constant acceleration γ . This approximation is particularly accurate at the beginning of the simulation, but as shown

on the inset, it eventually deteriorates, as seen by the bouncing times starting to deviate as the number of jumps grows unbounded. Note that the hybrid system with constant acceleration $\dot{x}_2 = -\bar{\gamma}$ exhibits Zeno behavior as time approaches approximately 28 s (as the amount of time between jumps decreases to zero). Similar phenomena typically emerge in certain hybrid control systems, such as in event-triggered controllers [79]. This behavior is also “approximately” inherited by the hybrid system (10) with oscillatory acceleration $\gamma(\tau)$. To study hybrid set-seeking systems, we will formalize this “approximation” property using averaging theory for hybrid systems. \square

Switching systems constitute a class of HDSs that arise frequently in applications. When such systems are modeled within a hybrid systems framework, their qualitative behavior is typically analyzed with respect to prescribed families of switching signals [76], [80]. In this case, set-valued dynamics can provide a convenient mathematical tool to model such signals. We illustrate this idea in the following example.

Example 3 (Source-Seeking and Surveillance)

Consider a point-mass vehicle in the plane, with state $x \in \mathbb{R}^2$ modeling its position and simple dynamics given by a single integrator

$$\dot{x} = \alpha \quad (12)$$

where α is the input to be designed. We study a typical application of ES in the context of source seeking [81], [82], [83], where the vehicle (12) seeks to stabilize its position at the location where a potential field, only available via measurements, attains its maximum intensity. However, unlike standard source-seeking problems with a single static potential field, we consider the presence of multiple intermittent potential fields, which are “active” only during bounded periods of time and never simultaneously. In this case, the goal of the vehicle is to achieve *persistent source*

seeking and surveillance assuming the intermittence of the sources is low; that is, they remain individually active for “sufficiently” long periods of time. This problem can be modeled using the formalism of hybrid set-seeking systems. In particular, let $\mathcal{Q} := \{0, 1, \dots, q_{\max}\}$ denote a finite set of logic modes, where $q_{\max} \in \mathbb{Z}_{\geq 1}$. For each $q \in \mathcal{Q}$, we let $J_q: \mathbb{R}^2 \rightarrow \mathbb{R}$ denote a potential field that attains its maximum at the point x_q^* . For simplicity, we assume that each function $-J_q(\cdot)$ is strongly convex and has a globally Lipschitz gradient ∇J_q . For each fixed potential field $J_q(\cdot)$, the vehicle can achieve source seeking by implementing the following standard seeking feedback law [84]:

$$\alpha = \varepsilon_a \omega \mathbf{R}_o \mu + \frac{2}{\varepsilon_a} J_q(x) \mu \quad (13)$$

where $\varepsilon_a, \omega > 0$ are tunable parameters, and μ is a state that models a dithering signal with dynamics given by the following constrained differential equation:

$$\mu \in \mathbb{S}^1, \quad \dot{\mu} = \omega \mathbf{R}_o \mu. \quad (14)$$

The linear oscillator (14) provides an alternative time-invariant approach to generate sinusoidal functions, such as those in (3), for exploratory purposes. In particular, by direct computation, any solution to (14) has the form

$$\mu(t) = \begin{bmatrix} \cos(\omega t) & \sin(\omega t) \\ -\sin(\omega t) & \cos(\omega t) \end{bmatrix} \begin{bmatrix} \mu_1(0) \\ \mu_2(0) \end{bmatrix} \quad (15)$$

where $\mu_1(0)^2 + \mu_2(0)^2 = 1$. In this way, the feedback law (13) emulates the source-seeking algorithm studied in [84] for single-source-seeking problems, but using a time-invariant model of the dynamics (with μ restricted to evolve in \mathbb{S}^1); see Figure 4 for a graphical representation of the system.

To incorporate the switching behavior of the logic state q that indexes the potential fields J_q , we can consider the

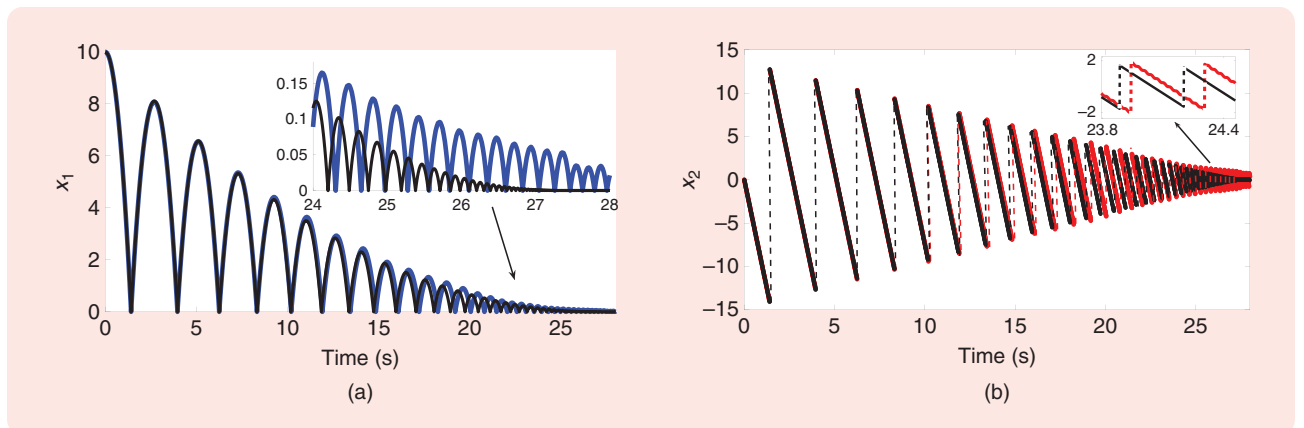


FIGURE 3 Trajectories of the “bouncing” seeking system with time-varying function $\gamma(\tau)$ using $\omega = 100$ (shown in blue and red) and with constant $\gamma = \bar{\gamma} = 10$ (shown in black). (a) Trajectories of the position x_1 . (b) Trajectories of the velocity x_2 , which experience jumps, illustrated with dashed lines.

following set-valued hybrid dynamics with states $(q, \tau) \in \mathcal{Q} \times \mathbb{R}_{\geq 0}$, which are parameterized by the constants $N_0 \geq 1$ and $\eta_d > 0$:

$$(q, \tau) \in \mathcal{Q} \times [0, N_0], \quad \dot{q} = 0, \quad \dot{\tau} \in [0, \eta_d] \quad (16a)$$

$$(q, \tau) \in \mathcal{Q} \times [1, N_0], \quad q^+ \in \mathcal{Q} \setminus \{q\}, \quad \tau^+ = \tau - 1. \quad (16b)$$

Similar to Example 1, in this system, the state τ acts as a timer that regulates how frequently q can jump to other values in the set \mathcal{Q} . However, in this case, the continuous-time dynamics of τ in (16a) are set-valued and allow the derivative




FIGURE 4 A scheme of source-seeking dynamics for point-mass vehicle models under intermittent switching potentials with distinct optimal points x_q^* .




FIGURE 5 The evolution in time of a switching signal q taking values in the set $\{0, 1\}$, indicating which potential field is active in the source-seeking problem.

$\dot{\tau}$ to be any (measurable) selection from the set $[0, \eta_d]$, which includes keeping τ constant via $\dot{\tau} = 0$, increasing τ at the maximum rate $\eta_d > 0$ until $\tau = N_0$, or any other absolutely continuous function τ (not necessarily constant) having a derivative lower bounded by zero and upper-bounded η_d . This flexibility allows the model (16) to capture a large class of switching signals q acting on the feedback law (14), including signals that never switch (corresponding to $\dot{\tau} = 0$ for all time). In fact, for the case where $N_0 = 1$, all switching signals generated by (16) satisfy a *dwell-time* condition [85]. When $N_0 \in \mathbb{Z}_{\geq 2}$, all switching signals generated by (16) satisfy an *average dwell-time* condition [80] (see also the “Seeking With Average Dwell-Time Switching” section). Both classes of switching signals are common in the modeling and stability analysis of switching systems, and by using (16) interconnected with (12)–(14) (with $x^+ = x$), we obtain a hybrid system suitable for the study of the source-seeking feedback law (13) under average dwell-time switches of J_q . Figure 5 shows an example of this switching signal for the case where $\eta_d = 1/100$, $N_0 = 1$, and $q_{\max} = 1$, such that $\mathcal{Q} = \{0, 1\}$. When η_d is sufficiently small, the resulting slowly switching set-seeking system generates the stable behavior shown in Figure 6. In Figure 6(a), we show the trajectories of (12) under the feedback law (13) with $\omega = 1000$, $\epsilon_a = 0.01$, and from different initial conditions shown with red dots. The trajectories converge to a neighborhood of the set $\Omega(K)$, shown in green. This set, called the “Omega-limit set from K ,” described later in (59), turns out to be (semiglobally practically) asymptotically stable for the overall hybrid set-seeking system (12)–(16) as $\eta_d \rightarrow 0^+$, $\epsilon_a \rightarrow 0^+$, and $1/\omega \rightarrow 0^+$. While in certain cases the set $\Omega(K)$ can be computed explicitly [86], its mere existence is typically guaranteed in slowly switching systems with individually asymptotically stable modes, under mild assumptions [8, Sec. 7.4]. \square

The preceding example underscores the advantages of considering hybrid set-seeking systems with *set-valued* dynamics. Consequently, for the remainder of this article,

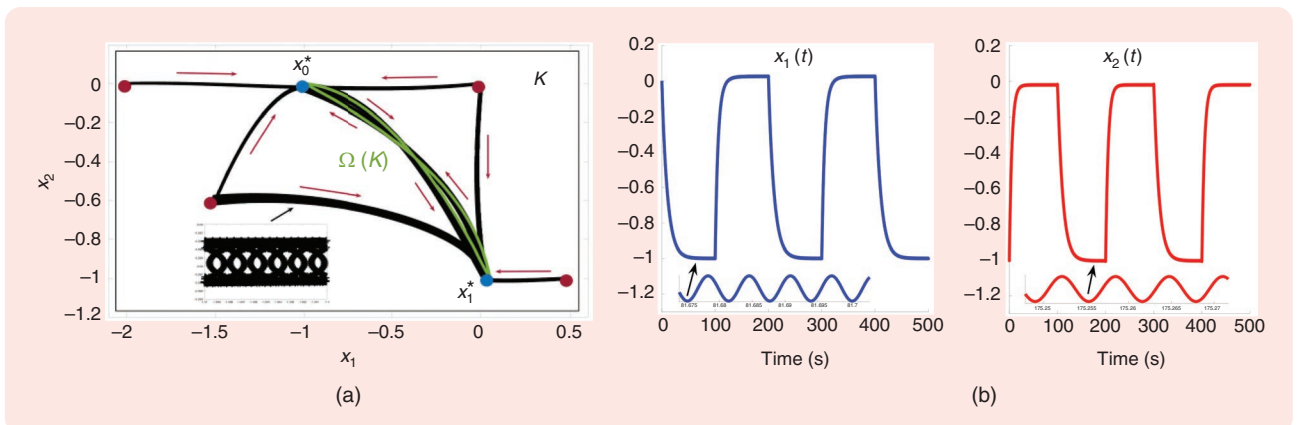


FIGURE 6 (a) The trajectories of a system (12) (shown in black) under the feedback law (13), evolving on the plane, and converging to a neighborhood of the set $\Omega(K)$, shown in green. (b) The evolution in time of the position of the vehicle, which oscillates between the two optimal points x_0^* and x_1^* , which are contained in $\Omega(K)$.

we will focus on hybrid systems that generalize (8) and are represented by the following *hybrid inclusions*:

$$x \in C, \quad \dot{x} \in F(x) \quad (17a)$$

$$x \in D, \quad x^+ \in G(x) \quad (17b)$$

where the flow map $F: \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ and the jump map $G: \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ are in general set-valued. Note that for particular cases in which such maps are single-valued, we can still use the formalism (17) by defining

$$F(x) := \begin{cases} \{f(x)\}, & x \in C \\ \emptyset, & x \notin C \end{cases} \quad (18a)$$

$$G(x) := \begin{cases} \{g(x)\}, & x \in D \\ \emptyset, & x \notin D. \end{cases} \quad (18b)$$

The data of (17) are also denoted $\mathcal{H} := \{C, F, D, G\}$, and similarly to (9), for the case where (17) depends on exogenous time-varying signals, we can consider models of the form

$$(x, \tau) \in C \times \mathbb{R}_{\geq 0}, \quad \dot{x} \in F(x, \tau), \quad \dot{\tau} = \omega \quad (19a)$$

$$(x, \tau) \in D \times \mathbb{R}_{\geq 0}, \quad x^+ \in G(x), \quad \tau^+ = \tau \quad (19b)$$

where $\omega > 0$ dictates the rate of evolution of the auxiliary timer state τ . Note that (19) can be written as (17) using the overall state $y = (x, \tau)$. Also, note that for hybrid systems with periodic continuous-time dynamics, it is also possible to consider models with bounded timers by incorporating periodic resets in τ , similar to Example 1.

The Hybrid Basic Conditions

A crucial factor in establishing the stability properties of hybrid set-seeking systems is ensuring robustness to perturbations in the states and dynamics. To guarantee this property, we will impose the appropriate regularity conditions on the data of the hybrid system (17). For differential equations, such conditions are reduced to the continuity of the vector field. However, to study dynamic inclusions, we will consider set-valued maps that are outer-semicontinuous (OSC) and locally bounded (LB). A set-valued mapping $M: \mathbb{R}^p \rightrightarrows \mathbb{R}^n$ is OSC at $z \in \mathbb{R}^p$ if, for each convergent sequence $\{z_i, s_i\} \rightarrow (z, s) \in \mathbb{R}^p \times \mathbb{R}^n$ satisfying $s_i \in M(z_i)$ for all $i \in \mathbb{Z}_{\geq 0}$, we have $s \in M(z)$. A mapping M is LB at z if there exists an open neighborhood $N_z \subset \mathbb{R}^p$ of z such that $M(N_z)$ is bounded. The mapping M is OSC and LB relative to a set $K \subset \mathbb{R}^p$ if the mapping from \mathbb{R}^p to \mathbb{R}^n defined by $M(z)$ for $z \in K$ and \emptyset for $z \notin K$ is OSC and LB at each $z \in K$. Using OSC and LB as well as standard convexity and closeness notions for sets [87], we consider hybrid systems (17) that satisfy the following *hybrid basic conditions*:

- a) The sets C and D are closed.

- b) The set-valued mapping F is OSC and LB relative to C , and for each $x \in C$, we have that $F(x)$ is nonempty and convex.
- c) The set-valued mapping G is OSC and LB relative to D , and for each $x \in D$, we have that $G(x)$ is nonempty.

Outer semicontinuity of a set-valued mapping $M: \mathbb{R}^p \rightrightarrows \mathbb{R}^n$ relative to a set $K \subset \mathbb{R}^p$ is equivalent to asking for closedness of the graph of M with respect to K ; that is, the set $\text{gph}(M) := \{(z \in \mathbb{R}^p, s \in \mathbb{R}^n) : z \in K, s \in M(z)\}$ is relatively closed in $K \times \mathbb{R}^n$ [8, Lemma 5.10]. Throughout this article, we will work only with hybrid systems that satisfy the hybrid basic conditions. In the context of ES, the controllers will always be designed to satisfy these conditions.

Hybrid Time Domains, Hybrid Arcs, and Solutions

For continuous-time systems of the form (8a) or (17a), solutions $x: \text{dom}(x) \rightarrow \mathbb{R}^n$ are given by absolutely continuous functions with respect to time $t \in \text{dom}(x) \subset \mathbb{R}_{\geq 0}$ (see “Continuous-Time Set-Valued Dynamical Systems”). Similarly, for discrete-time systems of the form (8b) or (17b), solutions $x: \text{dom}(x) \rightarrow \mathbb{R}^n$ are defined as sequences parameterized by a discrete-time index $j \in \text{dom}(x) \subset \mathbb{Z}_{\geq 0}$ (see “Discrete-Time Set-Valued Dynamical Systems”). Therefore, to define solutions for hybrid systems of the form (8) or (17), we consider a “hybrid” parameterization of the solutions that depends on both a continuous-time index $t \in \mathbb{R}_{\geq 0}$, which increases continuously during flows, and a discrete-time index $j \in \mathbb{Z}_{\geq 0}$, which increases by one during jumps. In this way, solutions to (8) or (17) are defined in *hybrid time domains*. A set $E \subset \mathbb{R}_{\geq 0} \times \mathbb{Z}_{\geq 0}$ is called a *compact hybrid time domain* if $E = \cup_{j=0}^{J-1} ([t_j, t_{j+1}], j)$ for some finite sequence of times $0 = t_0 \leq t_1 \leq \dots \leq t_J$. A set $E \subset \mathbb{R}_{\geq 0} \times \mathbb{Z}_{\geq 0}$ is a hybrid time domain if it is the union of a nondecreasing sequence of compact hybrid time domains, namely, E is the union of compact hybrid time domains E_j with the property that $E_0 \subset E_1 \subset E_2 \subset \dots \subset E_j \dots$, etc.

Figure 7 shows five different types of hybrid time domains, including those corresponding to Zeno behavior, average dwell time between jumps, purely discrete solutions, and purely continuous solutions. Using the notion of hybrid time domain, the following definition formalizes what we call *solutions* to hybrid systems of the form (17).

Definition 1 (Hybrid Arcs and Solutions)

A function $x: \text{dom}(x) \rightarrow \mathbb{R}^n$ is a hybrid arc if $\text{dom}(x)$ is a hybrid time domain and $t \mapsto x(t, j)$ is a locally absolutely continuous function for each j such that the interval $I_j := \{t : (t, j) \in \text{dom}(x)\}$ has a nonempty interior. A hybrid arc x is a *solution* to (17) if $x(0, 0) \in C \cup D$ and if the following two conditions hold:

- a) For each $j \in \mathbb{Z}_{\geq 0}$

$$x(t, j) \in C, \text{ for almost all } t \in I_j$$

$$\dot{x}(t, j) \in F(x(t, j)), \text{ for almost all } t \in I_j.$$

Continuous-Time Set-Valued Dynamical Systems

Set-valued continuous-time dynamics modeled as *differential inclusions* typically arise in hybrid seeking systems due to discontinuous feedback-based optimization or decision-making algorithms, plants with uncertain parameters, or auxiliary states used to generate families of switching signals. Here, we review some basic information on differential inclusions [8, Ch. 4], [68], [S1].

DIFFERENTIAL INCLUSIONS

Consider a constrained differential inclusion of the form

$$x \in C, \quad \dot{x} \in F(x) \quad (\text{S1})$$

where $C \subset \mathbb{R}^n$, and $F: \mathbb{R}^n \rightrightarrows \mathbb{R}^n$. System (S1) is said to satisfy the *basic conditions* if C is a closed set, and F is outer semi-continuous (OSC) and locally bounded (LB) relative to C , with $F(x)$ being nonempty and convex for each $x \in C$. Given $T > 0$, an absolutely continuous function $x: [0, T] \rightarrow \mathbb{R}^n$ is said to be a solution to (S1) if: a) $x(0) \in C$, and b) for almost all $t \in [0, T]$

$$\frac{dx(t)}{dt} \in F(x(t)) \quad \text{and} \quad x(t) \in C. \quad (\text{S2})$$

When (S1) satisfies the basic conditions, the existence of solutions can be studied under the so-called viability conditions [8, Lemma 5.26]. To state such conditions, recall that the *tangent cone* to a set $S \subset \mathbb{R}^n$ at a point $x \in \mathbb{R}^n$, denoted $T_S(x)$, is given by all vectors $s \in \mathbb{R}^n$ for which there exist $x_i \in S$, $\tau_i > 0$ with $x_i \rightarrow x$ and $\tau_i \rightarrow 0^+$, such that $s = \lim_{i \rightarrow \infty} \frac{x_i - x}{\tau_i}$. In particular, under the basic conditions, we have that

a) *Necessity*: If $x: [0, T] \rightarrow \mathbb{R}^n$, for some $T > 0$, is a solution to (S1), then

$$F(x(0)) \cap T_C(x(0)) \neq \emptyset. \quad (\text{S3})$$

b) *Sufficiency*: Given $x_0 \in C$, if there exists a neighborhood U of x_0 such that for all $x \in U \cap C$

$$F(x) \cap T_C(x) \neq \emptyset \quad (\text{S4})$$

then there exists $T > 0$ and a solution $x: [0, T] \rightarrow \mathbb{R}^n$ to (S1) with $x(0) = x_0$.

Note that, in general, given an initial condition $x(0) = x_0 \in \mathbb{R}^n$, solutions to (S1) from x_0 might not be unique.

KRASOVSKII SOLUTIONS OF ODES

Differential inclusions (S1) often arise when studying discontinuous differential equations for which solutions might not exist from some initial conditions. In particular, given a set $C \subset \mathbb{R}^n$, a function $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$, and an absolutely continuous function $x: [0, T] \rightarrow \mathbb{R}^n$, the function x is said to be a Krasovskii solution to the constrained system $x \in C$, $\dot{x} = f(x)$ if

$$x(t) \in \overline{C}, \quad \text{and} \quad \dot{x}(t) \in F_K(x(t)) := \bigcap_{\epsilon > 0} \overline{\text{co}} f((x(t) + \epsilon \mathbb{B}) \cap C) \quad (\text{S5})$$

for almost all $t \in [0, T]$. When f is continuous, the mappings F_K and f coincide. The construction (S5) is also valid whenever f is set-valued, even if f does not necessarily satisfy the basic conditions. Indeed, by construction (see [8, Example 6.6]), if f is LB relative to \overline{C} , then the differential inclusion (S5) satisfies the basic conditions.

ROBUSTNESS: CONTINUOUS-TIME HERMES SOLUTIONS

For differential equations with a continuous right-hand side, small additive perturbations to the states do not significantly affect the system's behavior. In contrast, small perturbations can drastically impact the behavior of systems with a discontinuous right-hand side. The effect of arbitrarily small perturbations on such systems can be analyzed using the concept of Hermes solutions [S2]. The function $x: [0, T] \rightarrow \mathbb{R}^n$ is said to be a Hermes solution to the constrained system $x \in C$, $\dot{x} = f(x)$ if there exists a sequence of absolutely continuous functions $x_i: [0, T] \rightarrow \mathbb{R}^n$ and a sequence of measurable functions $e_i: [0, T] \rightarrow \mathbb{R}^n$ such that

$$\begin{aligned} x_i(t) + e_i(t) &\in C, \quad \text{for all } t \in (0, T) \\ \dot{x}_i(t) &= f(x_i(t) + e_i(t)), \quad \text{for almost all } t \in [0, T] \end{aligned}$$

the sequence $\{x_i\}_{i=1}^\infty$ converges uniformly to x on $[0, T]$, and the sequence $\{e_i\}_{i=1}^\infty$ converges uniformly to the zero function on $[0, T]$. In other words, Hermes solutions capture the limiting effect of arbitrarily small additive state perturbations acting on the system $x \in C$, $\dot{x} = f(x)$. The connection between Hermes solutions and Krasovskii solutions is established in the following lemma, corresponding to [8, Thm. 4.3].

Lemma 1 (Continuous-Time Krasovskii Solutions)

Suppose that f is LB, and let $x: [0, T] \rightarrow \mathbb{R}^n$ be an absolutely continuous function. Then, x is a Hermes solution to $\dot{x} = f(x)$ if and only if it is a Krasovskii solution to the same system.

Put simply, Lemma 1 suggests that to accurately characterize the behavior of the solutions of the system $\dot{x} = f(x)$ under arbitrarily small additive perturbations on the state, we should consider its Krasovskii solutions. A different concept of solutions explored in the literature of discontinuous differential equations is that of Filippov solutions [S1]. An absolutely continuous function $x: [0, T] \rightarrow \mathbb{R}^n$ is said to be a Filippov solution to $\dot{x} = f(x)$ if it satisfies

$$\dot{x}(t) \in F_F(x(t)) := \bigcap_{\epsilon > 0} \bigcap_{\mu(S)=0} \overline{\text{co}} f((x(t) + \epsilon \mathbb{B}) \setminus S) \quad (\text{S6})$$

where μ is the Lebesgue measure. While the constructions (S5) and (S6) are similar, the construction in (S6) ignores the behavior of f on sets S with zero measure. Hence, Filippov solutions are Krasovskii solutions, but the converse is generally not true. Thus, in light of Lemma 1, Filippov solu-

(Continued)

Continuous-Time Set-Valued Dynamical Systems (Continued)

tions are problematic as they do not completely capture the behavior of the original dynamics under small additive state perturbations. For further discussions on the difference between Filippov and Krasovskii solutions and for a proof of Lemma 1 for the special case where $C = \mathbb{R}^n$, we refer the reader to [S3].

REFERENCES

- [S1] A. F. Filippov, *Differential Equations With Discontinuous Right-Hand Sides*. Boston, MA, USA: Kluwer, 1988.
 [S2] H. Hermes, "Discontinuous vector fields and feedback control," in *Proc. Int. Symp. Differential Equ. Dynamical Syst.*, 1967.
 [S3] O. Hájek, "Discontinuous differential equations, I," *J. Differential Equ.*, vol. 32, no. 2, pp. 149–170, 1979, doi: 10.1016/0022-0396(79)90056-1.

Discrete-Time Set-Valued Dynamical Systems

Discrete-time dynamical systems modeled as *difference inclusions* are common in optimization, estimation, and control algorithms. They are also relevant for the study of robustness properties in discrete-time systems with a discontinuous right-hand side. Here, we review some fundamentals of difference inclusions.

DIFFERENCE INCLUSIONS

Consider the constrained difference inclusion of the form

$$x \in D, \quad x^+ \in G(x) \quad (S7)$$

where $D \subset \mathbb{R}^n$, and $G: \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ is a set-valued mapping. System (S7) is said to satisfy the *basic conditions* if D is a closed set, G is OSC and LB relative to D , and $G(x)$ is not empty for each $x \in D$. Given $J \in \mathbb{Z}_{\geq 1}$, a sequence $x: \{0, 1, \dots, J\} \rightarrow \mathbb{R}^n$ is said to be a solution to (S7) if it satisfies

$$x(j) \in D, \quad \text{and} \quad x(j+1) \in G(x(j)) \quad (S8)$$

for all $j \in \{0, 1, \dots, J-1\}$. For system (S7), the existence of solutions is always guaranteed provided $D \neq \emptyset$. In general, given a solution $x: \text{dom}(x) \rightarrow \mathbb{R}^n$, it will not be the unique solution from $x(0)$ if $G(x(j))$ is not single-valued for some j in $\text{dom}(x)$ such that $x(j) \in D$.

KRASOVSKII SOLUTIONS TO DISCRETE-TIME SYSTEMS

Difference inclusions of the form (S7) are relevant for the study of robustness properties in discontinuous discrete-time systems. In particular, given a set $D \subset \mathbb{R}^n$ and a function $g: \mathbb{R}^n \rightarrow \mathbb{R}^n$, a sequence $x: \{0, 1, \dots, J\} \rightarrow \mathbb{R}^n$ is said to be a Krasovskii solution to the constrained difference equation

$$x \in D, \quad x^+ = g(x) \quad (S9)$$

if it satisfies

$$x(j) \in \bar{D}, \quad \text{and} \quad x(j+1) \in G_K(x) := \bigcap_{\epsilon > 0} \overline{g((x + \epsilon \mathbb{B}) \cap D)} \quad (S10)$$

for all $j \in \{0, 1, \dots, J\}$ [8, Def. 4.13]. When D is closed, the sets \bar{D} and D coincide. Similarly, when g is continuous, the mappings G_K and g coincide. However, when g is discontinuous, the map G_K is set-valued at the points of discontinuity. In particular, G_K describes the OSC hull of g , namely the unique set-valued mapping G_K satisfying $\text{graph}(G_K) = \overline{\text{graph}(g)}$. The construction in (S10) also holds when g is a set-valued mapping. In this case, by construction (see [8, Ex.6.6]), if g is LB relative to \bar{D} , then the difference inclusion (S10) also satisfies the basic conditions.

Example 4 (Regularization of Sign Function)

Consider the discrete-time system given by

$$x^+ = x - \alpha \cdot \text{sign}(x), \quad x \in \mathbb{R}^n \quad (S11)$$

for $\alpha > 0$, where the sign function is defined as

$$\text{sign}(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases} \quad (S12)$$

To obtain G_K , in this case, it suffices to close the graph of $\text{sign}(x)$. We obtain the set-valued mapping

$$\overline{\text{sign}}(x) = \begin{cases} 1 & \text{if } x > 0 \\ \{-1, 0, 1\} & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases} \quad (S13)$$

and the difference inclusion

$$x^+ \in G_K(x) := x - \alpha \cdot \overline{\text{sign}}(x), \quad x \in \mathbb{R}^n. \quad (S14)$$

ROBUSTNESS: DISCRETE-TIME HERMES SOLUTIONS

As in the continuous-time case, the limiting effect of small additive state disturbances on difference equations can be captured using the notion of Hermes solutions. A function $x: \{0, 1, \dots, J\} \rightarrow \mathbb{R}^n$ is said to be a Hermes solution to the constrained difference equation (S9) if there exists a sequence of functions $x_i: \{0, 1, \dots, J-1\} \rightarrow \mathbb{R}^n$ and $e_i: \{0, 1, \dots, J-1\} \rightarrow \mathbb{R}^n$ such that, for all $j \in \{0, 1, \dots, J-1\}$

$$x_i(j) + e_i(j) \in D, \quad x_i(j+1) = g(x_i(j) + e_i(j))$$

(Continued)

Discrete-Time Set-Valued Dynamical Systems (Continued)

and the sequence $\{x_i\}_{i=1}^{\infty}$ converges uniformly to x on $j \in \{0, 1, \dots, J\}$, and moreover, the sequence $e_i: \{0, 1, \dots, J\} \rightarrow \mathbb{R}^n$ converges uniformly to the zero function on $\{0, 1, \dots, J\}$ [8, Def. 4.12]. As in the continuous-time case, there is a close connection between discrete-time Hermes and Krasovskii solutions, which is stated in the next lemma, corresponding to [8, Thm. 4.17].

Lemma 2 (Discrete-Time Krasovskii Solutions)

Consider the function $x: \{0, 1, \dots, J\} \rightarrow \mathbb{R}^n$, for some $J \in \mathbb{Z}_{\geq 1}$, and suppose that g is LB. Then, x is a Hermes solution to the constrained difference equation (S9) if and only if it is a Krasovskii solution to the same system.

The result of Lemma 2 indicates that to study the robustness properties of difference equations with a discontinuous right-hand side, we should focus on studying their Krasovskii solu-

tions. For further details on the study of difference inclusions, we refer the reader to references [S4] and [S5]. Discrete-time ES algorithms based on difference equations and inclusions have been studied in [88] and [89] using finite-difference approximation methods to estimate the gradients of cost functions. In the context of hybrid set-seeking systems via averaging, difference inclusions can be used to model synchronization mechanisms in multiagent systems [90], switching systems under dwell-time or average dwell-time constraints [91], etc.

REFERENCES

- [S4] C. M. Kellett and A. R. Teel, "On the robustness of \mathcal{KL} -stability for difference inclusions: Smooth discrete-time Lyapunov functions," *SIAM J. Control Optim.*, vol. 44, no. 3, pp. 777–800, 2005, doi: 10.1137/S0363012903435862.
- [S5] C. M. Kellett and A. R. Teel, "Smooth Lyapunov functions and robustness of stability for difference inclusions," *Syst. Control Lett.*, vol. 52, no. 5, pp. 395–405, 2004.

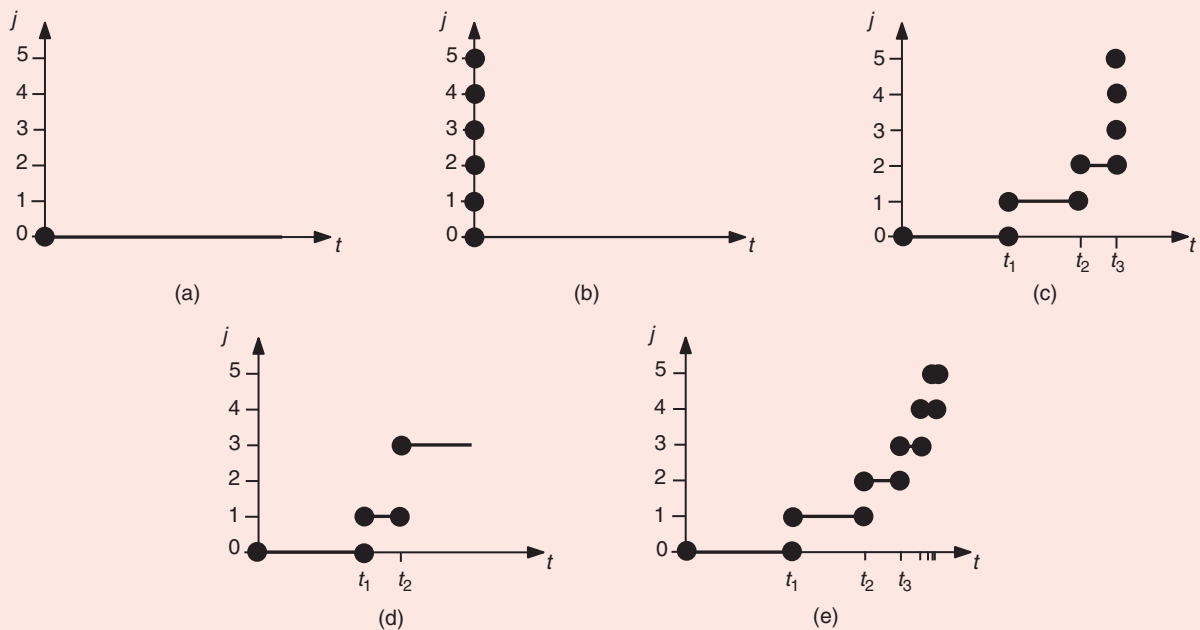


FIGURE 7 Five different hybrid time domains illustrating different types of solutions in a hybrid system. (a) Purely continuous solution. (b) Purely discrete solution. (c) Eventually discrete solution. (d) Solution with average dwell time. (e) Solution with Zeno behavior.

b) For each $(t, j) \in \text{dom}(x)$ such that $(t, j + 1) \in \text{dom}(x)$

$$x(t, j) \in D, \text{ and } x(t, j + 1) \in G(x(t, j)). \quad (20)$$

A solution x to system (17) is said to be *forward precomplete* if its domain is compact or unbounded. It is said to be

forward complete if its domain is unbounded, and it is said to be *maximal* if there does not exist another solution ψ to (17) such that $\text{dom}(x)$ is a proper subset of $\text{dom}(\psi)$, and $x(t, j) = \psi(t, j)$ for all $(t, j) \in \text{dom}(x)$. In general, given an initial condition $x(0, 0) \in C \cup D$, the solutions to system (17) might not be unique. This will be the case if, for example,




FIGURE 8 Hybrid arcs of a solution to the system studied in Example 1, illustrated as a function of time t in Figure 2.

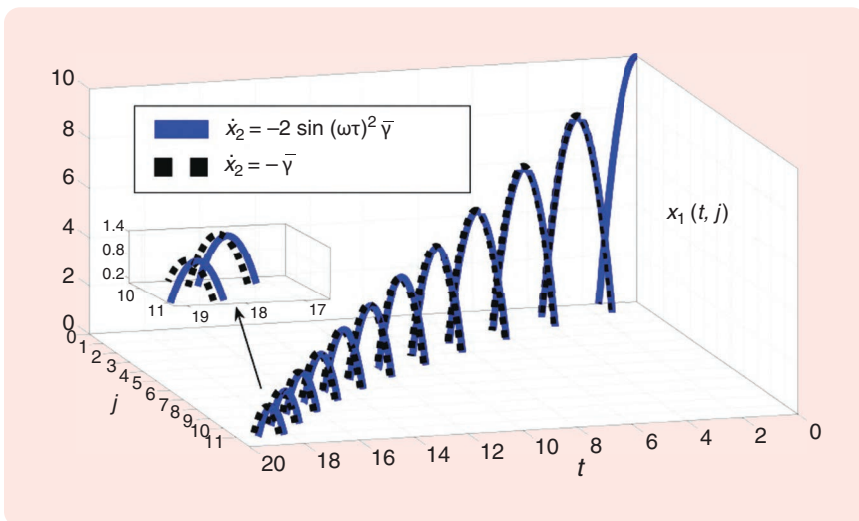


FIGURE 9 Component x_1 of the hybrid arcs associated with the solutions of the hybrid system (10), as studied in Example 2, shown as a function of time t in Figure 3. The blue trajectory is generated using the highly oscillatory acceleration $\gamma(\tau)$, with $\omega = 100$. The dotted black trajectory is generated using the constant acceleration $\bar{\gamma}$, which is just the average of $\gamma(\tau)$.

the intersection $C \cap D$ is not empty and from a point in this intersection, it is possible to flow, or if the flow or jump maps in (17a) and (17b) admit nonunique solutions due to set-valuedness or the lack of Lipschitz continuity in the flow map.

Example 5 [Periodic Hybrid System (Continued)]

Recall the hybrid system discussed in Example 1, which implements periodic jumps in the state x_2 triggered by a periodic resetting timer x_1 with resetting period T . Figure 8 shows the hybrid arcs corresponding to the trajectories x_1 and x_2 shown in Figure 2, from the initial conditions $x_1(0,0) = 0$ and $x_2(0,0) = 10$. As observed, the values of the states x_1 and x_2 can be fully identified via the hybrid time indices for all $(t,j) \in \text{dom}(x)$. \square

Example 6 [Bouncing Seeking System (Continued)]

Recall the “bouncing” seeking system discussed in Example 2. Figure 9 shows the hybrid arcs corresponding to the trajectories x_1 shown in Figure 3(a). The trajectory in black is generated by the hybrid dynamics (10) using constant acceleration; that is, $\dot{x}_2 = -\bar{\gamma}$. On the other hand, the blue trajectory implements highly oscillatory acceleration; that is, $\dot{x}_2 = -2 \sin(\omega\tau)^2 \bar{\gamma}$. Both trajectories were generated from the same initial conditions as in Example 2. As observed, both hybrid arcs remain “close” to each other during the simulation. Specifically, the graphs of the hybrid arcs are similar. This hints to some closeness property, akin to the one discussed when studying system (3) and its average dynamics (4) using inequality (7). In fact, by working with hybrid arcs, hybrid time domains, and a suitable metric for these objects, it is possible to naturally extend the notion of (T, ϵ) -closeness of solutions to hybrid seeking systems, thus opening the door to the development and study of novel ES systems and applications with hybrid dynamics in the loop. \square

Graphical Convergence and (T, ϵ) -Closeness in Hybrid Systems

As discussed in the analysis of the smooth ES dynamics (3), using the average dynamics (4) and the nominal dynamics (5), the closeness-of-solutions property is crucial for the study of ES controllers based on periodic dithering. Specifically, inequalities (6) and (7) showed that the solutions of the seeking dynamics (3) behave similarly to those of a gradient flow, with this “similarity” being quantified, in compact sets and time domains, using the uniform distance between solutions. However, this approach does not extend directly to hybrid systems. The following example, which continues the discussion from Example 2, illustrates some of the challenges that arise when using the standard uniform distance to gauge the closeness between solutions of nominal and perturbed hybrid systems. These challenges can be addressed by working with solutions defined on hybrid time domains, as opposed to solutions defined on $\mathbb{R}_{\geq 0}$, and by studying the closeness of solutions using the distance between the graphs of the hybrid arcs.

Example 7 [Bouncing Seeking System (Continued)]

Consider the “bouncing” seeking system with hybrid dynamics (10), which experiences a high-frequency oscillating acceleration given by $-\gamma(\tau)$. Using the fact that the average of $\sin(\tau)^2$ over one period of oscillations is equal to $1/2$, the average dynamics of (10b) can be explicitly computed (see “Averaging Theory: From ODEs to Hybrid Systems”) to be

$$C := \{\bar{x} \in \mathbb{R}^2 : \bar{x}_1 \geq 0\}, \quad \dot{\bar{x}} = \bar{f}(\bar{x}) := \begin{pmatrix} \bar{x}_2 \\ -\bar{\gamma} \end{pmatrix} \quad (21a)$$

$$D = \{\bar{x} \in \mathbb{R}^2 : \bar{x}_1 = 0, \bar{x}_2 \leq 0\}, \quad \bar{x}^+ = g(\bar{x}) := \begin{pmatrix} 0 \\ -\lambda \bar{x} \end{pmatrix} \quad (21b)$$

where $\bar{x} = (\bar{x}_1, \bar{x}_2) \in \mathbb{R}^2$ is the state of the average hybrid system, and $\bar{\gamma}$ is the same gravitational constant of (11). We can argue, at least visually based on Figure 3, that there exist $T, \epsilon > 0$ such that the solutions of the oscillating hybrid dynamics (10) and the nonoscillating hybrid dynamics (21) in general do not satisfy bounds of the form (6) or (7), no matter how large we select ω in (10b). For example, this behavior is evident in Figure 3(b), where after approximately 15 s, a significant mismatch emerges between the jump times of the red and black trajectories. Specifically, following a jump in the black trajectory, the red trajectory fails to closely track it due to the discontinuity, regardless of how large ω is chosen. This occurs because the red trajectory has not yet undergone its corresponding jump at that moment, preventing the satisfaction of inequalities of the form (6) or (7) at that time. \square

To overcome the issues described in the previous example and to study the closeness of solutions of hybrid systems of the form (17), we consider an equivalent notion that compares the *distance* between solutions defined under the same jump index j for which both solutions are defined [8, Def. 5.23]; that is, we use the distance between *their graphs*. Specifically, given constants $\tau \geq 0$ and $\epsilon > 0$, two hybrid arcs x and \bar{x} are said to be (τ, ϵ) -close if the following conditions are satisfied:

- 1) For each $(t, j) \in \text{dom}(x)$ with $t + j \leq \tau$, there exists $s \in \mathbb{R}_{\geq 0}$ such that $(s, j) \in \text{dom}(\bar{x})$, $|t - s| \leq \epsilon$, and

$$|x(t, j) - \bar{x}(s, j)| < \epsilon. \quad (22)$$

- 2) For each $(t, j) \in \text{dom}(\bar{x})$ with $t + j \leq \tau$, there exists $s \in \mathbb{R}_{\geq 0}$ such that $(s, j) \in \text{dom}(x)$, $|t - s| \leq \epsilon$, and

$$|\bar{x}(t, j) - x(s, j)| < \epsilon. \quad (23)$$

Figure 10 presents an illustration of this property, which essentially asks that the *graphs* of x and \bar{x} are “close” to each other when truncated in (hybrid) time. This property can also be leveraged to study a type of “graphical” convergence between hybrid arcs, akin to the standard uniform convergence properties studied for solutions of ODEs. In particular, given a sequence of hybrid arcs $\{x_i\}_{i=1}^{\infty}$ with $x_i : \text{dom}(x_i) \rightarrow \mathbb{R}^n$ and a hybrid arc x with $x : \text{dom}(x) \rightarrow \mathbb{R}^n$, the property that for each $\tau \geq 0$ and each $\epsilon > 0$, there exists $i^* \in \mathbb{Z}_{\geq 1}$ such that for all $i > i^*$ the hybrid arcs x_i and x are (τ, ϵ) -close, implies the property that the sequence $\{x_i\}_{i=1}^{\infty}$ converges graphically to x ; that is, the sequence of sets $\{\text{gph } x_i\}_{i=1}^{\infty}$ converges (in the sense of set convergence, see [8, Def. 5.1]) to the set $\{\text{gph } x\}$. The converse implication is also true whenever the sequence $\{x_i\}_{i=1}^{\infty}$ is locally eventually bounded¹ [8, Thm. 5.25].

By working with hybrid time domains, we can now see (at least at this point, visually based on Figure 9) that the hybrid arcs describing the two solutions considered in Example 2 are (τ, ϵ) -close when ω is sufficiently large. However, to formalize this property for a broad class of hybrid systems of interest, including those with additive perturbations on the states and dynamics [as in the perturbed gradient flow (4)], we need to introduce an “inflated” version of hybrid systems of the form (17), which will be particularly useful in the context of averaging (see “Averaging Theory: From ODEs to Hybrid Systems”) and singular perturbation theory (see “Singularly Perturbed HDSs”).

Inflated Hybrid Systems

Given a hybrid system \mathcal{H} of the form (17) with state $x \in \mathbb{R}^n$, a continuous function $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^n$, and $\rho > 0$, consider the following “inflated” hybrid system \mathcal{H}_ρ with state $x \in \mathbb{R}^n$ and data $\mathcal{H}_\rho := \{C_\rho, F_\rho, D_\rho, G_\rho\}$ constructed from the data of (17) as follows:

$$C_\rho := \{x \in \mathbb{R}^n : (x + \rho\sigma(x)\mathbb{B}) \cap C \neq \emptyset\} \quad (24a)$$

$$F_\rho(x) := \overline{c_0} F((x + \rho\sigma(x)\mathbb{B}) \cap C) + \rho\sigma(x)\mathbb{B}, \quad \forall x \in C_\rho \quad (24b)$$

$$D_\rho := \{x \in \mathbb{R}^n : (x + \rho\sigma(x)\mathbb{B}) \cap D \neq \emptyset\} \quad (24c)$$

$$G_\rho(x) := \{v \in \mathbb{R}^n : v \in g + \rho\sigma(g)\mathbb{B}, \\ g \in G((x + \rho\sigma(x)\mathbb{B}) \cap D)\}, \quad \forall x \in D_\rho. \quad (24d)$$

¹A sequence of hybrid arcs is said to be locally eventually bounded if for any $m > 0$, there exists $i_0 \in \mathbb{Z}_{\geq 1}$ and a compact set K such that for all $i > i_0$ and all $(t, j) \in \text{dom}(x_i)$ with $t + j \leq m$, we have that $x_i(t, j) \in K$.

Averaging Theory: From ODEs to Hybrid Systems

Dynamical systems with time-varying oscillating vector fields are traditionally studied using averaging theory. Classic references on averaging theory for differential equations include [56], [57], [S6], and [S7]. For HDSs, averaging tools have also been studied in [69], [100], [102], [103], and [S8]. Next, we review some results on averaging theory for hybrid systems, adapted from [69], [102], and [S7].

HYBRID SYSTEMS WITH OSCILLATING FLOW MAPS

Consider an HDS of the form (9), with states $z = (x, p) \in \mathbb{R}^n \times \mathbb{R}^m$ and $\tau \in \mathbb{R}_{\geq 0}$, and dynamics

$$(x, p, \tau) \in C := C_x \times C_p \times \mathbb{R}_{\geq 0}, \quad \begin{cases} \dot{x} = f(x, p, \tau) \\ \dot{p} \in F_p(p) \\ \dot{\tau} = \frac{1}{\varepsilon} \end{cases} \quad (\text{S15a})$$

$$(x, p, \tau) \in D := D_x \times D_p \times \mathbb{R}_{\geq 0}, \quad \begin{cases} x^+ = g(x, p) \\ p^+ \in G_p(p) \\ \tau^+ = \tau \end{cases} \quad (\text{S15b})$$

where $\varepsilon > 0$ is a small parameter, $C_x, D_x \subset \mathbb{R}^n$, $C_p, D_p \subset \mathbb{R}^m$. The state p can model logic states, timers, and oscillators as well as other auxiliary states. We assume that system (S15) satisfies the hybrid basic conditions. In (S15a), the small parameter ε introduces a timescale separation between the dynamics of τ and those of (x, p) . When the vector field f is periodic with respect to τ , the rapid fluctuations in τ result in fast oscillatory behaviors in \dot{x} , which can be averaged out to approximate the dynamics of x using a nonoscillating vector field. The following definition formalizes this notion of average

Definition 4 (The Average Map)

The function f is said to have an *average map* $\bar{f}(\cdot)$ if, for each compact set $K_0 \subset \mathbb{R}^n \times \mathbb{R}^m$, there exists a continuous nonincreasing function γ satisfying $\lim_{\mathcal{T} \rightarrow \infty} \gamma(\mathcal{T}) = 0$, such that for all $(x, p, \tau, T) \in K_0 \times \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$

$$\left| \frac{1}{\mathcal{T}} \int_{\tau}^{\tau+\mathcal{T}} (f(x, p, s) - \bar{f}(x, p)) ds \right| \leq \gamma(\mathcal{T}). \quad (\text{S16})$$

The function $\gamma(\cdot)$ is called the *convergence function*.

When f is periodic, as is usually the case in ES, for each compact set K_0 , there exists $c_0 > 0$ such that we can take $\gamma(\mathcal{T}) = c_0/(1 + \mathcal{T})$, and it suffices to verify the integral (S16) in the interval $[0, \mathcal{T}]$, where \mathcal{T} is the period of the oscillations. In the rest of this section, we make the additional standing assumptions that f is periodic with respect to τ and that \bar{f} is continuous. However, we note that the periodicity of f can be relaxed by working with generalized averages in (S16).

By using the average map \bar{f} , we can study (S15) by inspecting a “simpler” nonoscillating system, with states (\bar{x}, \bar{p}) , and the following *average hybrid dynamics*:

$$(\bar{x}, \bar{p}) \in C := C_x \times C_p, \quad \begin{cases} \dot{\bar{x}} = \bar{f}(\bar{x}, \bar{p}) \\ \dot{\bar{p}} \in F_p(\bar{p}) \end{cases} \quad (\text{S17a})$$

$$(x, p) \in D := D_x \times D_p, \quad \begin{cases} \bar{x}^+ = g(\bar{x}, \bar{p}) \\ \bar{p}^+ \in G_p(\bar{p}) \end{cases}. \quad (\text{S17b})$$

Given that hybrid systems usually lack the uniqueness of solutions property and might also exhibit solutions with jumps, standard results on closeness of solutions, established in the literature of averaging for differential equations, do not trivially

This inflated hybrid system can capture different types of *perturbations* acting on the flow and jump sets of the nominal hybrid dynamics as well as in the nominal flow and jump maps. For example, the inflated flow map F_p and the jump map G_p can capture additive disturbances acting on the states and dynamics of the flow and jump maps of system (17). Similarly, the inflated flow set C_p and jump set

D_p can model measurement noise or other small additive disturbances acting on the states of the nominal system (17). As shown in the following example, the inflated system (24b) can also be used for the study of continuous-time seeking systems, including system (3).

Example 8 (Inflations in Smooth Seeking Systems)

Consider the standard ES smooth dynamics (3), which can be written as (9) using an auxiliary state τ and dynamics

$$\dot{u} = f_\varepsilon(\hat{u}, \tau) := -\frac{2}{\varepsilon_a} J(\hat{u} + \varepsilon_a \sin(\tau)) \sin(\tau), \quad \dot{\tau} = \omega \quad (25)$$

where $\omega = (2\pi/\varepsilon_\omega)$, $\varepsilon = (\varepsilon_a, \varepsilon_\omega) \in \mathbb{R}_{>0}$ are tunable parameters, and where the flow set is given by the condition $(\hat{u}, \tau) \in C := \mathbb{R} \times \mathbb{R}_{\geq 0}$. For small values of ε_a , a Taylor expansion of J around \hat{u} leads to

$$f_\varepsilon(\hat{u}, \tau) = -\left(\frac{2}{\varepsilon_a} J(\hat{u}) \sin(\tau) + 2 \frac{\partial J(\hat{u})}{\partial \hat{u}} \sin(\tau)^2 + O(\varepsilon_a) \right)$$




FIGURE 10 (a) and (b) A graphical illustration of the property of (τ, ε) -closeness in two hybrid arcs x and \bar{x} .

extend to hybrid systems. Instead, we can establish that for every solution (x, ρ, τ) of the original hybrid system (S15), there exists some solution $(\bar{x}, \bar{\rho})$ of the average hybrid dynamics (S17) that is (T, ϵ) -close. In particular, the following theorem follows from the results in [100], [102], and [S7].

Theorem 1 (Closeness of Solutions via Averaging)

Consider the hybrid systems (S15) and (S17). Suppose that system (S17) has no finite escape times. Then, for each compact set of initial conditions $K_0 \subset \mathbb{R}^n$, $T \in \mathbb{R}_{>0}$ and $\epsilon \in \mathbb{R}_{>0}$, there exists $\epsilon^* > 0$ such that for all $\epsilon \in (0, \epsilon^*)$ and for all solutions (x, ρ, τ) to (S15) from K_0 , there exists a solution $(\bar{x}, \bar{\rho})$ to (S17) such that $(\bar{x}, \bar{\rho})$ and (x, ρ) are (T, ϵ) -close.

The result of Theorem 1 is a closeness-of-solutions property for hybrid systems that parallels similar properties used for the study of smooth ES systems. By using this property as well as suitable uniform stability properties in the average hybrid dynamics, the following theorem can be established.

Theorem 2 (SGPAS via Averaging in HDS)

Consider the hybrid systems (S15) and (S17). Suppose that there exists a compact set $\mathcal{A} \subset \mathbb{R}^{n+m}$ that is UGAS for system (S17). Then, system (S15) renders the set $\mathcal{A} \times \mathbb{R}_{\geq 0}$ SGPAS as $\epsilon \rightarrow 0^+$.

The stability result from Theorem 2 can be applied to analyze hybrid set-seeking dynamics of the form (S15). This result remains valid even when the stability properties of the averaged dynamics (S17) are only of semi-global practical nature, with respect to a tunable parameter δ [69, Thm. 7]. This is commonly the case in ES control, where δ can represent the

residual terms that appear in the average dynamics and that are of the same order as the amplitude of the dither signals, see also (4). In the context of hybrid set-seeking dynamics, δ may encompass additional parameters, offering greater flexibility in the design of the algorithms.

AVERAGING USING TIME-INVARIANT MODELS

Theorems 1 and 2 can be extended to other types of models, including hybrid systems of the form

$$(x, \rho, \tau) \in C := C_x \times C_\rho \times \mathbb{R}_{\geq 0}, \begin{cases} \dot{x} \in F_x(x, \rho) \\ \dot{\rho} = f_\rho(\rho, x, \tau) \\ \dot{\tau} = \frac{1}{\epsilon}, \end{cases} \quad (\text{S18a})$$

$$(x, \rho, \tau) \in D := D_x \times D_\rho \times \mathbb{R}_{\geq 0}, \begin{cases} x^+ \in G_x(x, \rho) \\ \rho^+ \in G_\rho(\rho, x, \tau) \\ \tau^+ = \tau \end{cases} \quad (\text{S18b})$$

where all the mappings and sets are selected to satisfy the hybrid basic conditions. Here, the dynamics of x are set-valued, but the average map can still be computed as in (S16) using f_ρ instead of f .

REFERENCES

[S6] J. Sanders and F. Verhulst, *Averaging Methods in Nonlinear Dynamical Systems*. New York, NY, USA: Springer-Verlag, 1985.
 [S7] A. R. Teel and D. Nesić, "Averaging for a class of hybrid systems," *Dyn. Continuous, Discrete Impulsive Syst.*, vol. 17, no. 6, pp. 829–851, 2010.
 [S8] W. Wang, D. Nesić, and A. R. Teel, "Input-to-state stability for a class of hybrid dynamical systems via averaging," *Math. Control, Signals, Syst.*, vol. 23, no. 4, pp. 223–256, 2012, doi: 10.1007/s00498-011-0070-y.

where, for a given compact set $K \subset \mathbb{R}$, the notation $O(\epsilon_a)$ indicates high-order terms $o_h(\tau, \hat{u})$ that satisfy, for some $k > 0$, the inequality $|o_h(\tau, \hat{u})| \leq k\epsilon_a$ for all $\tau \in \mathbb{R}_{\geq 0}$ and all $\hat{u} \in K$. Since $\int_0^{2\pi/\omega} \sin(\tau) d\tau = 0$ and $(\omega/2\pi) \int_0^{2\pi/\omega} \sin(\tau)^2 d\tau = (1/2)$, it follows that the average map f_{ave} of f_ϵ , given by (2), is precisely system (4), which satisfies

$$\begin{aligned} f_{ave}(\bar{u}) &= -\frac{\partial J(\hat{u})}{\partial \hat{u}} + O(\epsilon_a) \in \overline{CO} \left(\frac{\partial J(\bar{u} + k\epsilon_a \mathbb{B})}{\partial \bar{u}} \right) + k\epsilon_a \mathbb{B} \\ &=: F_{k\epsilon_a}(\bar{u}) \end{aligned}$$

for all $\bar{u} \in K$, where $F_{k\epsilon_a}$ is precisely the set-valued map defined in (24b) with $F = -(\partial J / \partial \hat{u})$, $C = \mathbb{R}$, a constant function $\sigma = k$, and $\rho = \epsilon_a$. The previous inclusion implies that, on compact sets K , the set of solutions of system $\dot{\bar{u}} = f_{ave}(\bar{u})$ is contained in the set of solutions of the differential inclusion $\dot{\bar{u}} \in F_{k, \epsilon_a}(\bar{u})$. Hence, by verifying the appropriate properties for all solutions of this differential inclusion (which is just an inflation of the model-based gradient flow), we can

guarantee that these properties will also hold for all solutions of $\dot{\bar{u}} = f_{ave}(\bar{u})$ □

By using the construction of the inflated hybrid system (24), the following key technical property, corresponding to [8, Prop. 6.34], opens the door to a perturbation-based methodology to study hybrid set-seeking systems, thus generalizing the approach used to study the smooth ES algorithm (3). We recall that throughout this article, we assume that the hybrid system \mathcal{H} given by (17) always satisfies the hybrid basic conditions by design.

Lemma 3 (Key Technical Lemma)

Suppose that $x_0 \in \mathbb{R}^n$ is such that each maximal solution to the nominal hybrid system \mathcal{H} , given by (17), from x_0 is complete or bounded. Then, for all $\tau \geq 0$ and $\epsilon > 0$, there exists $\rho > 0$ such that for each solution x_ρ to the perturbed HDS (24), with $|x_\rho(0, 0) - x_0| \leq \rho$, there exists a solution x to \mathcal{H} , with $x(0, 0) = x_0$, such that x and x_ρ are (τ, ϵ) -close. □

Lemma 3 also applies to differential equations simply by taking $D = \emptyset$, $C = \mathbb{R}^n$, and F being a continuous function. For instance, for the smooth ES system considered in Example 8, we can use Lemma 3 to directly conclude that, on compact sets, for each solution $\tilde{u}_{\varepsilon_n}$ of the differential inclusion $\dot{\tilde{u}}_{\varepsilon_n} \in F_{k,\varepsilon_n}(\tilde{u}_{\varepsilon_n})$, there exists a solution \tilde{u} of the model-based gradient flow $\dot{\tilde{u}} = -(\partial J(\tilde{u})/\partial \tilde{u})$ such that $\tilde{u}_{\varepsilon_n}$ and \tilde{u} are (τ, ε) -close. Since, as shown before, on compact sets we have $f_{\text{ave}} \in F_{k,\varepsilon_n}$, it follows that the (τ, ε) -closeness property also holds for the solutions of the average dynamics of (25) and the solutions of the target gradient flow. This analytical methodology extends beyond gradient flows and smooth dynamics to general hybrid set-seeking systems, enabling the establishment of closeness-of-solutions properties via perturbation theory for hybrid systems. In this context, the use of hybrid time domains and graphical convergence provides a foundational mathematical framework for hybrid set-seeking control.

Stability Notions for Hybrid Set-Seeking Systems

Since hybrid systems usually incorporate logic modes, timers, clocks, and other auxiliary states that do not settle into a point but rather evolve on bounded sets, we study the uniform stability properties of hybrid set-seeking systems with respect to general compact sets.

Definition 2 (Uniform Global Asymptotic Stability)

A compact set $\mathcal{A} \subset \mathbb{R}^n$ is said to be *uniformly globally asymptotically stable* (UGAS) for system (17) [or (19)] if there exists a class- \mathcal{KL} function β such that any maximal solution to \mathcal{H} satisfies the bound

$$|x(t, j)|_{\mathcal{A}} \leq \beta(|x(0, 0)|_{\mathcal{A}}, t + j), \quad \forall (t, j) \in \text{dom}(x). \quad (26)$$

If there exist positive constants $c_1, c_2 > 0$ such that $\beta(r, s) = c_1 r e^{-c_2 s}$, then \mathcal{A} is said to be *uniformly globally exponentially stable* (UGES). \square

When $\mathcal{A} = \{0\}$, Definition 2 recovers the standard asymptotic stability notions studied for differential equations [57, Ch. 4] and smooth ES systems [44], [64]. The use of \mathcal{KL} bounds in (26) provides uniformity in terms of rates of convergence and overshoots. Note that Definition 2 requires that *all* solutions to the HDS (17) satisfy the bound (26) at all times in the domain of the solution. However, it does not require that all solutions be complete. For hybrid systems, the completeness of solutions is usually established independently.

Example 9 (Dynamics with Vanishing Coefficients)

Under suitable assumptions, gradient flows and other optimization dynamics typically satisfy the UGAS property in Definition 2, relative to the set of minimizers of a cost function J . However, this assumption does not hold for all gradient systems, even if J is strongly convex. For instance, systems with vanishing coefficients of the form

$$\dot{x} = -\alpha(\tau)\nabla J(x), \quad \text{or} \quad \ddot{x} + \alpha(\tau)\dot{x} + \nabla J(x) = 0 \quad (27)$$

for which $\lim_{\tau \rightarrow \infty} \alpha(\tau) = 0$ typically exhibit nonuniform convergence properties with respect to $\tau(0)$, thus precluding bounds of the form (26) [92]. However, under the appropriate hybrid regularizations of τ , for example, time-triggered [69], [93] or event-triggered restarting [94], the resulting hybrid dynamics can satisfy (26) for some $\beta \in \mathcal{KL}$, see also “Set-Seeking Dynamics with Momentum and Resets.” \square

Since ES systems rely on perturbation theory, the following stability property is commonly used in the analysis of seeking algorithms.

Definition 3 (Semiglobal Practical Stability)

Consider an HDS \mathcal{H}_ε , where $\varepsilon > 0$ is a tunable parameter. System \mathcal{H}_ε is said to render a compact set \mathcal{A} *semiglobally practically asymptotically stable* (SGPAS) if there exists a class- \mathcal{KL} function β such that for each $\Delta > \nu > 0$, there exists $\varepsilon^* > 0$ such that for all $\varepsilon \in (0, \varepsilon^*)$ and all initial conditions $|x(0, 0)| \leq \Delta$, any maximal solution to \mathcal{H}_ε satisfies the bound

$$|x(t, j)|_{\mathcal{A}} \leq \beta(|x(0, 0)|_{\mathcal{A}}, t + j) + \nu, \quad \forall (t, j) \in \text{dom}(x). \quad (28)$$

If there exist positive constants $c_1, c_2 > 0$ such that $\beta(r, s) = c_1 r e^{-c_2 s}$, then \mathcal{A} is said to be *semiglobally practically exponentially stable* (SGPES). \square

The SGPAS property can be extended to hybrid systems that depend on multiple parameters $\varepsilon = (\varepsilon_1, \varepsilon_2, \dots, \varepsilon_m)$. In this case, we use the notation SGPAS as $(\varepsilon_1, \varepsilon_2, \dots, \varepsilon_m) \rightarrow 0^+$ to denote that (28) holds with the parameters ε_i sufficiently small and being selected sequentially, starting from ε_1 and ending with ε_m . While the sequential tuning of certain parameters in ES can be relaxed in some cases [66], to simplify our presentation, we do not study such relaxations in this article. Instead, we observe that for most seeking dynamics based on averaging—though not all (see [95], [96], and [97])—whether modeled as differential equations or hybrid systems, the property of SGPAS is typically the strongest guarantee one can expect. For smooth seeking systems, such as in (3), the small parameters in ε are related to the dither amplitude ε_n , the inverse of the dither frequency ω and, when additional dynamics such as filters are used, the corresponding inverse of the filter gain. For hybrid seeking systems, we will allow ε to also include parameters that are intrinsic to the hybrid dynamics, for example, η_d in the “source-seeking and surveillance” application discussed in Example 3.

PART 2: HYBRID SET-SEEKING SYSTEMS FOR STATIC MAPS

Equipped with the modeling and analysis tools introduced in Part 1, we now consider a broad class of hybrid set-seeking systems for solving model-free decision-making

problems. The primary objective in ES is to regulate the plant input u using only measurements of the output y to drive u toward the solution of an *application-dependent* optimization or decision-making problem of the form

$$\text{optimize } J(u), \text{ subject to } u \in \hat{U} \quad (29)$$

where the set $\hat{U} \subset \mathbb{R}^n$ captures the constraints on the input, and $J(\cdot)$ corresponds to the cost function, which is assumed to be sufficiently smooth. We assume that the set of solutions to the problem (29) is not empty, and we denote it by $O \subset \mathbb{R}^n$, which is also assumed to be compact. The fact that O might not be a singleton further motivates the terminology “set-seeking” as opposed to the more traditional names “extremum-seeking” or “equilibrium seeking.”

Remark 1

The decision-making problem (29) can also be formulated to study multiagent networked systems with individual outputs y_i and cost functions J_i , for all $i \in \{1, \dots, N\}$, where N is the number of subsystems. This case is relevant for distributed optimization problems as well as for Nash equilibrium (NE)-seeking problems in game theoretic scenarios; see, for instance, [51], [104], [105], [106], and [107] for different examples of seeking dynamics studied in these settings.

For the remainder of this section, we assume that the plant to be optimized is characterized by a static map J , as illustrated in the scheme of Figure 11. In this scheme, we use suggestive notation in which the data of the hybrid

controller $\mathcal{H} = \{C, F, D, G\}$ are allowed to depend on y , the output signal of the plant. This notation indicates that the particular structure of the *data* of \mathcal{H} is designed specifically for each application of interest. However, the general structure of the hybrid set-seeking controllers remains largely consistent across applications.

To investigate hybrid set-seeking systems with dynamic inclusions in the loop, we begin by extending the smooth seeking dynamics (3) to the setting of hybrid systems. Specifically, similar to Example 3, and to avoid dealing with time-varying differential inclusions whose averaged maps may be difficult to compute explicitly, we consider a class of time-invariant hybrid seeking systems consisting of a *dynamic derivative estimator* and a *dither generator*, with states $(\xi, \mu) \in \mathbb{R}^n \times \mathbb{R}^n$, and a *hybrid decision-making algorithm*, with internal state $x_{u,z} := (\hat{u}, \hat{z}) \in \mathbb{R}^n \times \mathbb{R}^r$. Figure 12 shows a block diagram representation of the set-seeking controllers. Our objective is to design the system’s core components so that, using only output measurements $J(u)$, the plant input $u \in \mathbb{R}^n$ is driven to a small neighborhood of the solution set of (29). The design must also ensure that the resulting closed-loop system satisfies the hybrid basic conditions.

Autonomous Derivative Estimators

The dynamic derivative estimators make use of n linear oscillators of the form (14), with overall state $\mu := (\mu_1, \dots, \mu_n) \in \mathbb{S}^n \subset \mathbb{R}^{2n}$, satisfying $\mu_i = (\mu_{i,1}, \mu_{i,2})^\top \in \mathcal{S}^1$, for all $i \in \{1, \dots, n\}$, and dynamics

$$\dot{\mu} = \Phi(\omega)\mu, \quad \mu \in \mathbb{S}^n \quad (30)$$

where the ω -parameterized block matrix $\Phi_\omega: \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$ is given by

$$\Phi_\omega := \begin{bmatrix} \omega_1 \mathbf{R}_o & 0 & \dots & 0 \\ 0 & \omega_2 \mathbf{R}_o & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \omega_n \mathbf{R}_o \end{bmatrix} \quad (31)$$

with $\omega = (\omega_1, \dots, \omega_n)$. The entries of ω are selected as $\omega_i = (\kappa_i / \epsilon_\omega)$, where $\epsilon_\omega > 0$ is a tunable constant, and the κ_i ’s are rational numbers that satisfy $\kappa_i \neq \kappa_j$ for $i \neq j$ and $(i, j) \in \{1, \dots, n\}$.

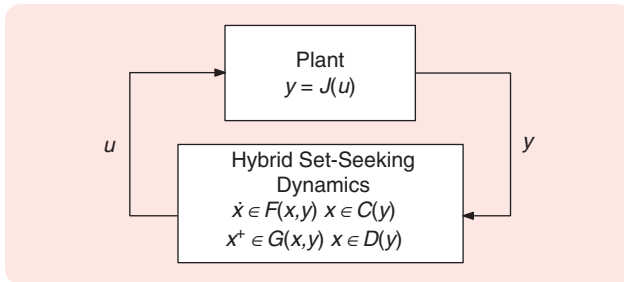


FIGURE 11 A high-level conceptual feedback scheme of hybrid set-seeking systems for static maps.

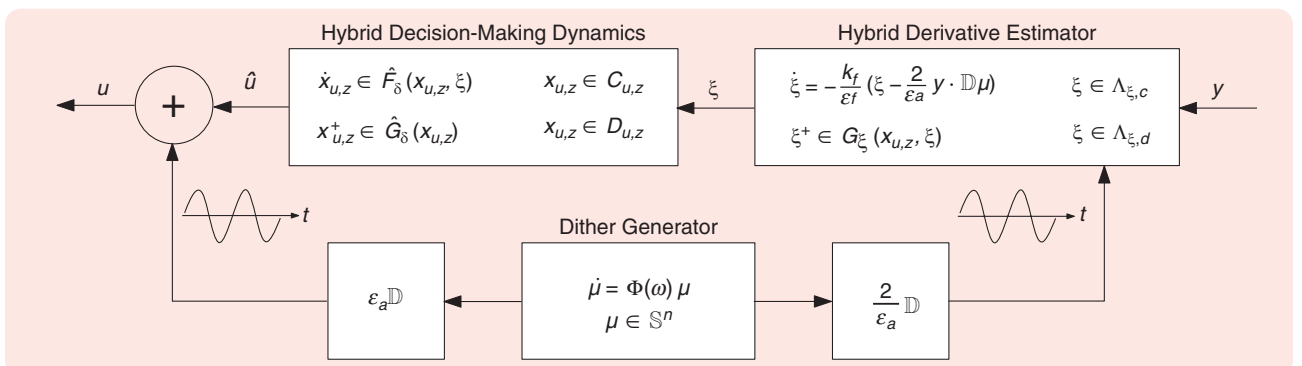


FIGURE 12 The main components of hybrid set-seeking systems with periodic probing.

The role of system (30) is to generate a vector of n periodic probing signals with different frequency for the purpose of real-time *exploration*. This is achieved by extracting the odd entries of the vector μ , given by $\mu_{i,1}(t) = \cos(\omega_i t)\mu_{2i-1}(0) + \sin(\omega_i t)\mu_{2i}(0)$, for each $i \in \{1, \dots, n\}$ via multiplication with the matrix $\mathbb{D} := [e_1, \mathbf{0}, e_2, \mathbf{0}, \dots, e_i, \dots, \mathbf{0}, e_n, \mathbf{0}]$, where $\mathbf{0} \in \mathbb{R}^n$ corresponds to a column vector of zeros, and $e_i \in \mathbb{R}^n$ corresponds to the unitary vector with the i th entry equal to one. Since, by the definition of its flow set in (30), the oscillator is restricted to evolve in \mathbb{S}^n , we have $\mu_{2i-1}(0)^2 + \mu_{2i}(0)^2 = 1$. Moreover, the set \mathbb{S}^n is forward invariant, and since the system (30) has no solutions defined outside \mathbb{S}^n , the set is also trivially UGAS for the dynamics (30). We note that although other types of dither signals and oscillators can be considered—including hybrid ones [108]—we focus here, for simplicity, on standard sinusoidal probing signals generated by linear oscillators.

The performance of the derivative estimators in ES systems can usually be enhanced by incorporating filters into the dynamics. For hybrid set seeking, we can start by considering a simple low-pass filter with state $\xi \in \mathbb{R}^n$ and continuous-time dynamics

$$\dot{\xi} = -\frac{k_f}{\varepsilon_f} \left(\xi - \frac{2}{\varepsilon_a} y \cdot \mathbb{D} \mu \right) \quad (32)$$

where $k_f, \varepsilon_f > 0$ are tunable parameters and which is allowed to evolve in a compact set $\Lambda_\xi := \lambda_\xi \mathbb{B} \subset \mathbb{R}^n$, with $\lambda_\xi \in \mathbb{R}_{>0}$ selected sufficiently large to encompass all the complete solutions of interest. As shown later in Theorem 4 and in “Hybrid Set-Seeking Dynamics with Momentum and Resets,” it is also possible to incorporate filters with hybrid dynamics to, for example, improve transient performance. However, we can simply take $\xi^+ = \xi$ when no discrete-time updates are incorporated into the filters. The compactness assumption on Λ_ξ can also be relaxed under additional structure on the dynamics.

Hybrid Decision-Making Dynamics

As mentioned in the “Introduction,” the basic ES algorithm (3) is designed to approximate—on average—the behavior

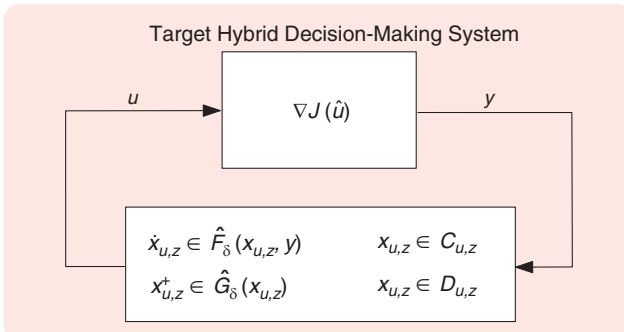


FIGURE 13 A conceptual feedback scheme of the target hybrid decision-making system considered in Assumption 1.

of the model-based “target” gradient flow (5). In the context of hybrid set-seeking systems, we aim to consider more general hybrid target systems for optimization and decision making. Therefore, instead of (3), we consider a class of hybrid seeking dynamics with state $x_{u,z} := (\hat{u}, \hat{z}) \in \mathbb{R}^n \times \mathbb{R}^r$, $r \in \mathbb{Z}_{\geq 0}$, where \hat{u} is the main state of the system and where $z \in \mathbb{R}^r$ is an auxiliary state that can model timers, logic modes, memory states, etc.

The state $x_{u,z}$ evolves according to the following *hybrid decision-making dynamics*:

$$\dot{x}_{u,z} \in \hat{F}_\delta(x_{u,z}, \xi), \quad x_{u,z} \in C_{u,z} \quad (33a)$$

$$x_{u,z}^+ \in \hat{G}_\delta(x_{u,z}), \quad x_{u,z} \in D_{u,z} \quad (33b)$$

where the set-valued mappings $\hat{F}_\delta: \mathbb{R}^{n+r} \times \mathbb{R}^n \rightrightarrows \mathbb{R}^{n+r}$ and $\hat{G}_\delta: \mathbb{R}^{n+r} \rightrightarrows \mathbb{R}^{n+r}$ are allowed to depend on a small tunable parameter $\delta > 0$ that provides flexibility for the design of the controller. It follows that the data of system (33) are given by $\hat{\mathcal{H}}_\delta := \{C_{u,z}, \hat{F}_\delta, D_{u,z}, \hat{G}_\delta\}$; it receives as input the state ξ , generated by the filter (32), and it generates the state \hat{u} , which acts on the plant via the following control law:

$$u = \hat{u} + \varepsilon_a \mathbb{D} \mu. \quad (34)$$

The nominal hybrid decision-making dynamics (33) are designed so that the following assumption holds.

Assumption 1 (Hybrid Decision-Making Dynamics)

The following properties hold:

- a) There exists $\delta^* \in \mathbb{R}_{>0}$ such that for all $\delta \in (0, \delta^*]$
 - 1) The sets $C_{u,z}$ and $D_{u,z}$ are closed and satisfy $P_u(C_{u,z}) \subset \hat{\mathcal{U}}$, $P_u(D_{u,z}) \subset \hat{\mathcal{U}}$, where P_u denotes the projection on the u -component of the set.
 - 2) $\hat{F}_\delta(\cdot, \cdot)$ is OSC, LB, and convex-valued relative to $C_{u,z} \times \mathbb{R}^n$, and $\hat{G}_\delta(\cdot)$ is OSC and LB relative to $D_{u,z}$.
- b) There exists a compact set $\Psi \subset \mathbb{R}^r$ such that the set $\mathcal{A} := \mathcal{O} \times \Psi$ is SGPAS as $\delta \rightarrow 0^+$ for the *target hybrid decision-making system* corresponding to (33) with $\xi := \nabla J(\hat{u})$, where the set \mathcal{O} is the set of solutions to (29).

While condition (a) provides regularity to the data of system $\hat{\mathcal{H}}_\delta$, item (b) indicates that $\hat{\mathcal{H}}_\delta$ should be designed to guarantee that the set of solutions to problem (29) is stabilized under the assumption that the gradient ∇J is known by the hybrid decision-making dynamics. Figure 13 shows a block diagram representation of the target hybrid decision-making system considered in item (b) of Assumption 1. However, we stress that, in practice, the hybrid controller will use only real-time measurements of J , and the exact mathematical forms of J and ∇J are assumed to be unknown. However, qualitative properties of J , for example, quadratic-like structures, convexity, invexity, etc., are typically used to study stability in ES algorithms.

If $\hat{\mathcal{H}}_\delta$ does not depend on any parameter δ , then item (b) in Assumption 1 simply asks that the set \mathcal{A} is UGAS.

If, additionally, system (33) does not depend on any auxiliary state $\hat{z} \in \mathbb{R}^r$, then the set Ψ can be neglected, and we can take $C_{u,z} = C_u \subset \mathbb{R}^n$ and $D_{u,z} = D_u \subset \mathbb{R}^n$. We will use the convention $r = 0$ to specify this case. Finally, note that if, in addition, we have that $D_u = \emptyset$, $C_u = \mathbb{R}^n$, and \hat{F}_δ is a Lipschitz continuous function, then the previous properties reduce to the standard assumptions considered in the modern literature of smooth ES [65], [72], [109]. Local (practical) stability results can also be established when only the local qualitative properties of ∇J or J are known.

Hybrid Set-Seeking: Closed-Loop System

The overall hybrid set-seeking dynamics are obtained by interconnecting the static plant $y = J(u)$, the oscillator (30), the filter (32), and the hybrid decision-making dynamics (33) using the feedback law (34). The resulting hybrid system has an overall state $x := (x_{u,z}, \xi, \mu) \in \mathbb{R}^{r+4n}$ and data given by $\mathcal{H} := \{C, F, D, G\}$, with

$$C := C_{u,z} \times \Lambda_\xi \times \mathbb{S}^n \quad (35a)$$

$$\dot{x} \in F(x) := \begin{pmatrix} \hat{F}_\delta(x_{u,z}, \xi) \\ -\frac{k_f}{\varepsilon_f} \left(\xi - \frac{2}{\varepsilon_a} y \cdot \mathbb{D}\mu \right) \\ \Phi(\omega)\mu \end{pmatrix} \quad (35b)$$

$$D := D_{u,z} \times \Lambda_\xi \times \mathbb{S}^n \quad (35c)$$

$$x^+ \in G(x) := \begin{pmatrix} \hat{G}_\delta(x_{u,z}) \\ \xi \\ \mu \end{pmatrix}. \quad (35d)$$

The stability properties of (35) are captured by the following theorem, which follows by computing the average dynamics of (35) with respect to μ and then applying singular perturbation arguments to study the average hybrid dynamics as a perturbed multitime-scale system with the state ξ evolving on a faster timescale; see Theorem 6 and “Remark 1: Averaging Using Singular Perturbations” in “Singularly Perturbed HDSs.”

Theorem 3

Suppose that Assumption 1 holds. Then, the hybrid set-seeking system (35) with $y = J(u)$ and the feedback control law (34) renders the set $\mathcal{A} \times \Lambda_\xi \times \mathbb{S}^n$ SGPAS as $(\delta, \varepsilon_f, \varepsilon_a, \varepsilon_\omega) \rightarrow 0^+$. \square

The stability result of Theorem 3 states that the trajectories of (35) satisfy bounds of the form (26) with respect to the set $\mathcal{A} \times \Lambda_\xi \times \mathbb{S}^n$. In particular, since Λ_ξ and \mathbb{S}^n are compact, the state $x_{u,z}$ in (35) satisfies

$$|x_{u,z}(t, j)|_{\mathcal{A}} \leq \beta(|x_{u,z}(0, 0)|_{\mathcal{A}}, t + j) + \nu \quad (36)$$

for all $(t, j) \in \text{dom}(x_{u,z})$. Note that Theorem 3 does not make assertions about the completeness of solutions since

these properties need to be established for each particular application under an appropriate design of the data of the hybrid system.

Remark 2

The SGPAS result of Theorem 3 also establishes a tuning order for the parameters $(\delta, \varepsilon_f, \varepsilon_a, \varepsilon_\omega)$. The first parameter to be selected “sufficiently” small is δ , such that item (b) in Assumption 1 holds. Subsequently, the parameters ε_f , ε_a , and ε_ω should be selected to be sufficiently small, in that order. \square

Corollary 1

Suppose that Assumption 1 holds and the sets $C_{u,z}$ and $D_{u,z}$ are bounded. Then, the hybrid dynamics (35) with $y = J(u)$ and control law (34) render the set $\mathcal{A} \times \Lambda_\xi \times \mathbb{S}^n$ Globally Practically Asymptotically Stable (GPAS) as $(\delta, \varepsilon_f, \varepsilon_a, \varepsilon_\omega) \rightarrow 0^+$. \square

Achieving global stability when the flow and jump sets are bounded is particularly relevant for ES problems defined on smooth compact manifolds. In such spaces, standard ES algorithms with target decision-making dynamics characterized by smooth ODEs can achieve only local, or, at best, almost global asymptotic stability. This follows because such spaces are not contractible, a property that is needed for globally asymptotic stability in smooth dynamical systems [110]. However, hybrid set-seeking systems can overcome these limitations and achieve global bounds of the form (36), see [111] and [112].

Remark 3

When $D_{u,z} := \emptyset$, system (35) does not experience jumps, and the set-seeking dynamics reduce to the continuous-time system

$$C := C_{u,z} \times \Lambda_\xi \times \mathbb{S}^n \quad (37a)$$

$$\dot{x} \in F(x) := \begin{pmatrix} \hat{F}_\delta(x_{u,z}, \xi) \\ -\frac{k_f}{\varepsilon_f} \left(\xi - \frac{2}{\varepsilon_a} y \cdot \mathbb{D}\mu \right) \\ \Phi(\omega)\mu \end{pmatrix}. \quad (37b)$$

Set-valued ES systems of the form (37) can emerge in applications where the optimizing vector field is discontinuous. In such cases, \hat{F}_δ will correspond to the Krasovskii regularization of the vector field (see “Krasovskii Solutions of ODEs” in “Continuous-Time Set-Valued Dynamical Systems”). If, on the other hand, \hat{F}_δ is a Lipschitz continuous function, and $C_{u,z} = \mathbb{R}^n$, then (37) recovers the standard framework for the study of gradient-based smooth ES dynamics [45], [66], [65]. \square

One of the advantages of modeling seeking systems with autonomous dither generators of the form (30) is that the resulting closed-loop system (35) satisfies the hybrid basic conditions and has suitable stability properties with respect to a compact set (see Theorem 3). As a consequence, one can directly study the inflated hybrid system (24)

Singularly Perturbed HDSs

Stability properties of ES systems rely on multiple timescale separations induced in the closed-loop system via the appropriate tuning of the control parameters. Singular perturbation theory is a well-established field in the literature of differential equations that enables the study of systems for which certain states evolve in a faster timescale and converge to a quasi-steady-state manifold, where the original system can be “reduced” [57], [58]. Here, we review some extensions of singular perturbations to hybrid systems [102], [S6], [S7] that fit the models considered in this article. Further details can be found in [S7].

HYBRID SYSTEMS WITH “FAST” AND “SLOW” STATES

Consider a hybrid system with states $(x, \theta) \in \mathbb{R}^{n_1} \times \mathbb{R}^{n_2}$ and the following dynamics:

$$(x, \theta) \in C := C_x \times \Theta, \quad \begin{cases} \dot{x} \in F_x(x, \theta) \\ \varepsilon \dot{\theta} \in F_\theta(x, \theta) \end{cases} \quad (\text{S19a})$$

$$(x, \theta) \in D := D_x \times \Theta, \quad (x^+, \theta^+) \in G(x, \theta) \quad (\text{S19b})$$

where the set-valued mappings $F_x: \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \rightrightarrows \mathbb{R}^{n_1}$ and $F_\theta: \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \rightrightarrows \mathbb{R}^{n_2}$ characterize the continuous-time dynamics of x and θ , respectively. Similarly, $G: \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \rightrightarrows \mathbb{R}^{n_1+n_2}$ describes the jump map of the system. The flow set C is composed of the Cartesian product of the sets $C_x \subset \mathbb{R}^{n_1}$ and $C_z \subset \mathbb{R}^{n_2}$, while the jump set D is composed of the sets $D_x \subset \mathbb{R}^{n_1}$ and $D_\theta \subset \mathbb{R}^{n_2}$. We make the assumption that system (S19) satisfies the hybrid basic conditions and that Θ is a compact set.

In system (S19), $\varepsilon \in \mathbb{R}_{>0}$ is a small parameter that induces a timescale separation between the continuous-time dynamics of the state $x \in \mathbb{R}^{n_1}$ and the continuous-time dynamics of

$\theta \in \mathbb{R}^{n_2}$. In particular, when ε is small, $x \in \mathbb{R}^{n_1}$ behaves as a “slow” state in the system, while $\theta \in \mathbb{R}^{n_2}$ behaves as a “fast” state. When the fast dynamics in system (S23) define a suitable “quasi-steady” state, the behavior of the state x can be approximately predicted by a *reduced* system. For example, this is the case when θ converges to a continuously differentiable quasi-steady-state manifold $m: \mathbb{R}^n \rightarrow \mathbb{R}^m$, as is usually the case in the literature of differential equations. However, in the context of hybrid systems, it is common to have fast dynamics for which the quasi-steady state is actually characterized by a set-valued mapping $M: \mathbb{R}^n \rightrightarrows \mathbb{R}^m$. In such cases, we assume that M is OSC and LB and that for each $\rho > 0$, the set

$$M_\rho = \{(x, \theta) : \theta \in M(x), x \in \rho\mathbb{B}\} \quad (\text{S20})$$

is UGAS for the following *boundary layer* dynamics:

$$(x, \theta) \in (C_x \cap \rho\mathbb{B}) \times \Theta, \quad \dot{x} = 0, \quad \dot{\theta} \in F_\theta(x, \theta). \quad (\text{S21})$$

Note that in (S21), the state x is taken as a constant from any initial condition on $C_x \cap \rho\mathbb{B}$.

Using the mapping M , the *reduced hybrid system*, with state $x_r \in \mathbb{R}^n$, can be defined using the following dynamics:

$$x \in C_x, \quad \dot{x} \in F_r(x) := \overline{\text{co}}\{f : f \in F_x(x, \theta), \theta \in M(x)\} \quad (\text{S22a})$$

$$x \in D_x, \quad x^+ \in G_r(x) := \{f_1 : (f_1, f_2) \in G_x(x, \theta), \theta \in \Theta\}. \quad (\text{S22b})$$

By construction, and by the assumptions on the data of system (S19), the reduced hybrid system (S22) satisfies the hybrid basic conditions.

obtained from the data of the hybrid seeking dynamics (35) to establish the following *robustness* result via [8, Thm. 7.21].

Corollary 2

Suppose that Assumption 1 holds. For each compact set of initial conditions $K_0 \subset \mathbb{R}^{n+r}$, let Theorem 3 generate sufficiently small parameters $(\delta, \varepsilon_f, \varepsilon_a, \varepsilon_w)$ such that property (36) holds. Then, the *inflated* hybrid set-seeking system (24) generated from the data of system (35), with control law (34), renders the set $\mathcal{A} \times \Lambda_\xi \times \mathbb{S}^n$ SGPAS as $\rho \rightarrow 0^+$ (relative to K_0). \square

Since, as discussed in Part 1 of this article, the inflated system (24) can capture a variety of additive disturbances and perturbations acting on the states and dynamics of the algorithms, the result of Corollary 2 essentially implies that the hybrid set-seeking algorithms (35) will retain their (semiglobal practical) stability properties under a broad class of small additive disturbances and perturbations acting on the states and dynamics of the algorithms.

Remark 4

To account for situations where the plant output is influenced by external events, such as sporadic feedback measurements or external malicious signals, the filter in (35b) can be modified to adopt the more general structure

$$\dot{\xi} = -\frac{k_f}{\varepsilon_f} \left(\xi - \frac{2}{\varepsilon_a} \gamma(x_{u,z}) y D \mu \right) \quad (38)$$

where $\gamma(\cdot)$ is a continuous function. In this case, item (b) of Assumption 1 should hold with $\xi = \gamma(x_{u,z}) \nabla J(\hat{u})$. \square

Hybrid Set-Seeking Dynamics With Hybrid Filters

By removing any explicit dependence of \hat{F}_δ on the oscillatory signal μ , the nonhybrid low-pass filter considered in (35) facilitates the computation of the average hybrid dynamics for a broad class of set-valued decision-making dynamics (33). However, this comes at the price of having an additional timescale in the system induced by

Theorem 6 (SGPAS in Singularly Perturbed HDSs)

Suppose that there exists a compact set $\mathcal{A} \subset \mathbb{R}^n$ that is UGAS for the reduced HDS (S22). Then, the singularly perturbed HDS (S19) renders the set $\mathcal{A} \times \Theta$ SGPAS as $\varepsilon \rightarrow 0^+$.

When the hybrid system (S19) possesses additional structure, stronger results, such as UGAS or UGES, can be established. For instance, in the context of hybrid ES with hybrid filters operating on a faster timescale, the flow and jump sets in (S19) take the form $C = C_x \times C_\theta$ and $D = D_x \times D_\theta$, where C_θ and D_θ are not necessarily compact. By imposing the appropriate Lyapunov-based conditions on the reduced dynamics, boundary layer dynamics, and certain interconnection properties, UGAS or UGES can usually be established for these dynamics, provided that ε is sufficiently small [102, Sec. 5]. This approach mirrors the classical composite Lyapunov method based on quadratic-like Lyapunov functions used to analyze singularly perturbed differential equations [S11].

Remark 1: Averaging Using Singular Perturbations

There are close connections between averaging and certain models of singularly perturbed systems. For example, when the fast state in (S23) does not settle to a point but instead oscillates, it is still possible to define a well-posed reduced hybrid system. For example, when the continuous-time dynamics in (S23) have the form

$$(x, \theta) \in C := C_x \times \Theta, \quad \begin{cases} \dot{x} = f_x(x, \theta) \\ \varepsilon \dot{\theta} = f_\theta(x, \theta) \end{cases} \quad (\text{S23})$$

and the jump map in (S19b) is of the form $G(x, \theta) = G_x(x) \times \{\theta\}$, we can define a hybrid reduced system using the data

$$\bar{x} \in C_x, \quad \dot{\bar{x}} = \bar{f}(\bar{x}) \quad (\text{S24a})$$

$$\bar{x} \in D_x, \quad \bar{x}^+ \in G(\bar{x}) \quad (\text{S24b})$$

where \bar{f} is a continuous function that satisfies

$$\left| \frac{1}{T} \int_0^T (f_x(x, \theta(s)) - \bar{f}(x, p)) ds \right| \leq \gamma_K(T) \quad (\text{S25})$$

for all $T \in \mathbb{R}_{>0}$ all $x \in C_x \cap K$, all compact sets $K \subset \mathbb{R}^n$, and all solutions $\theta: [0, T] \rightarrow \mathbb{R}^m$ of the dynamics $\dot{\theta}(t) = f_\theta(x, \theta(t))$, where $\gamma_K(\cdot)$ is a continuous nonincreasing function that satisfies $\lim_{T \rightarrow \infty} \gamma_K(T) = 0$ and that is allowed to depend on K . Note that the definition of \bar{f} in (S25) is similar to the average map defined in (S16). Indeed, one can study oscillatory hybrid systems with fast states generated by time-invariant systems via averaging using the model (S24). This observation is important to exploit robustness properties in well-posed hybrid systems with compact attractors.

REFERENCES

[S9] W. Wang, A. Teel, and D. NeŠić, "Analysis for a class of singularly perturbed hybrid systems via averaging," *Automatica*, vol. 48, no. 6, pp. 1057–1068, 2012, doi: [10.1016/j.automatica.2012.03.013](https://doi.org/10.1016/j.automatica.2012.03.013).
 [S10] R. G. Sanfelice and A. R. Teel, "On singular perturbations due to fast actuators in hybrid control systems," *Automatica*, vol. 47, no. 4, pp. 692–701, 2011, doi: [10.1016/j.automatica.2011.01.055](https://doi.org/10.1016/j.automatica.2011.01.055).
 [S11] A. Saberi and H. Khalil, "Quadratic-type Lyapunov functions for singularly perturbed systems," *IEEE Trans. Autom. Control*, vol. 29, no. 6, pp. 542–550, Jun. 1984, doi: [10.1109/TAC.1984.1103586](https://doi.org/10.1109/TAC.1984.1103586).

selecting $\varepsilon_f > 0$ sufficiently small and by restricting ξ to evolve in a compact set for the purpose of stability analysis. Nevertheless, it is possible to relax these assumptions by replacing the flow and jump sets of (35) with the more general sets

$$C := C_{u,z} \times \Lambda_{\xi,c} \times \mathbb{S}^n, \quad D := D_{u,z} \times \Lambda_{\xi,d} \times \mathbb{S}^n \quad (\text{39a})$$

where $\Lambda_{\xi,c}, \Lambda_{\xi,d} \subset \mathbb{R}^n$ are closed sets, and by considering the following flow and jump maps:

$$\dot{x} \in F(x) := \begin{pmatrix} \hat{F}_\delta(x_{u,z}, \xi) \\ -k_f \left(\xi - \frac{2}{\varepsilon_a} y \cdot \mathbb{D}\mu \right) \\ \Phi(\omega)\mu \end{pmatrix} \quad (\text{39b})$$

$$x^+ \in G(x) := \begin{pmatrix} \hat{G}_\delta(x_{u,z}, \xi) \\ G_\xi(x_{u,z}, \xi) \\ \mu \end{pmatrix} \quad (\text{39c})$$

where the set-valued maps \hat{G}_δ and G_ξ are both assumed to be OSC and LB relative to $D_{u,z} \times \Lambda_{\xi,d}$. In this case, we consider the following assumption.

Assumption 2

Condition (a) of [Assumption 1](#) holds, and for the target hybrid decision-making dynamics

$$\begin{pmatrix} \dot{x}_{u,z} \\ \dot{\xi} \end{pmatrix} \in \begin{pmatrix} \hat{F}_\delta(x_{u,z}, \xi) \\ -k_f(\xi - \kappa \nabla J(\hat{u})) \end{pmatrix}, \quad (x_{u,z}, \xi) \in C_{u,z} \times \Lambda_{\xi,c} \quad (\text{40a})$$

$$\begin{pmatrix} x_{u,z}^+ \\ \xi^+ \end{pmatrix} \in \begin{pmatrix} \hat{G}_\delta(x_{u,z}, \xi) \\ G_\xi(x_{u,z}, \xi) \end{pmatrix}, \quad (x_{u,z}, \xi) \in D_{u,z} \times \Lambda_{\xi,d} \quad (\text{40b})$$

where $k_f > 0$, there exist $\kappa \in \mathbb{R}$ and compact sets $\Psi_z \subset \mathbb{R}^r$ and $\Psi_f \subset \mathbb{R}^n$ such that the set $\mathcal{A} := \mathcal{O} \times \Psi_z \times \Psi_f$ is SGPAS as $\delta \rightarrow 0^+$.

With [Assumption 2](#) at hand, we can establish the following theorem.

Theorem 4

Suppose that [Assumption 2](#) holds. Then, the hybrid set-seeking system (39) with $y = J(u)$ and control law (34) renders the set $\mathcal{A} \times \mathbb{S}^n$ SGPAS as $(\delta, \varepsilon_a, \varepsilon_\omega) \rightarrow 0^+$. \square

Seeking systems with hybrid filters is useful for incorporating resets or restarting mechanisms to reduce overshoot, improve transient performance [69], [94], [113], or regularize accelerated gradient flows with dynamic damping, which can speed up convergence when the cost function exhibits low curvature [69], [114], [115]. In such systems, the continuous-time dynamics of ξ can be appropriately modified without significantly altering the structure of (40).

Hybrid Set-Seeking Dynamics With Growing Timers

We now consider the simpler hybrid set-seeking systems that dispense with the filters and for which the time variation in the dither signal is generated by an exogenous timer that grows unbounded, as in systems (3) and (19); see also “Averaging Theory: From ODEs to Hybrid Systems.” In this case, we can consider hybrid set-seeking systems with control law (34) and dynamics

$$(x_{u,z}, \tau) \in C_{u,z} \times \mathbb{R}_{\geq 0}, \quad \begin{cases} \dot{x}_{u,z} = \hat{F}_\delta(x_{u,z}, y \cdot \mu_{\varepsilon_a}(\tau)) \\ \dot{\tau} = \frac{1}{\varepsilon_\omega} \end{cases} \quad (41a)$$

$$(x_{u,z}, \tau) \in D_{u,z} \times \mathbb{R}_{\geq 0}, \quad \begin{cases} x_{u,z}^+ = \hat{G}_\delta(x_{u,z}) \\ \tau^+ = \tau \end{cases} \quad (41b)$$

where $\mu_{\varepsilon_a}(\tau) := (2/\varepsilon_a)(\sin(2\pi\kappa_1\tau), \sin(2\pi\kappa_2\tau), \dots, \sin(2\pi\kappa_n\tau))$, and the parameters $\kappa_i > 0$ are rational and satisfy $\kappa_i \neq \kappa_j$ for all $i \neq j \in \{1, 2, \dots, n\}$. Note that the “bouncing” seeking system studied in Example 2 has the form of (41). When $D_{u,z} = \emptyset$, $C_{u,z} = \mathbb{R}^n$, and $\hat{F}_\delta = -y\mu_{\varepsilon_a}(\tau)$ we recover the most basic continuous-time ES scheme studied in the literature [45], [66].

The stability properties of the hybrid set-system (41) are studied under the following assumption, which leverages the notion of the *average map* introduced in “Averaging Theory: From ODEs to Hybrid Systems.”

Assumption 3

There exists $\delta^* \in \mathbb{R}_{>0}$ such that for all $\delta \in (0, \delta^*]$, the following properties hold:

- Condition a)-1 of [Assumption 1](#) is satisfied, and the functions \hat{F}_δ and \hat{G}_δ are continuous.
- There exists a compact set $\Psi \subset \mathbb{R}^r$ such that the set $\mathcal{A} := \mathcal{O} \times \Psi$ is SGPAS as $\delta \rightarrow 0^+$ for the target hybrid decision-making dynamics

$$\dot{x}_{u,z} = \hat{F}_\delta(x_{u,z}, \nabla J(\hat{u})), \quad x_{u,z} \in C_{u,z} \quad (42a)$$

$$x_{u,z}^+ = \hat{G}_\delta(x_{u,z}), \quad x_{u,z} \in D_{u,z}. \quad (42b)$$

- For each compact set $K \subset \mathbb{R}^n$, there exists $\varepsilon_a^* > 0$ such that for all $\varepsilon_a \in (0, \varepsilon_a^*)$, the average map of \hat{F}_δ in (41),

denoted \bar{F}_δ , satisfies $\bar{F}_\delta(\bar{x}_{u,z}) \in F_{\varepsilon_a}(\bar{x}_{u,z})$ for all $\bar{x}_{u,z} \in K$, where F_{ε_a} is the inflated set-valued map constructed as in (24b) from the data $F(x_{u,z}) := \hat{F}_\delta(x_{u,z}, \nabla J(\hat{u}))$.

The conditions of [Assumption 3](#) essentially ask that the flow map in (41) is designed such that, as $\varepsilon_\omega \rightarrow 0^+$, and on average, the hybrid dynamics behave as an ε_a -perturbed hybrid system for which the unperturbed nominal dynamics (42) are able to solve the decision-making problem (29).

A direct application of averaging tools—see Theorems 1 and 2 in “Averaging Theory: From ODEs to Hybrid Systems”—along with robustness results for hybrid systems satisfying the hybrid basic conditions (see [8, Thm. 7.21]), leads to the following theorem concerning the closed set $\mathcal{A} \times \mathbb{R}_{\geq 0}$.

Theorem 5

Suppose that [Assumption 3](#) holds. Then, the hybrid set-seeking system (41) with $y = J(u)$ and control law

$$u = \hat{u} + \frac{\varepsilon_a^2}{2} \mu_{\varepsilon_a}(\tau)$$

renders the set $\mathcal{A} \times \mathbb{R}_{\geq 0}$ SGPAS as $(\delta, \varepsilon_a, \varepsilon_\omega) \rightarrow 0^+$. \square

Hybrid Set-Seeking Systems With Estimates of Higher Derivatives

A variety of decision-making dynamics benefit from incorporating higher-order derivatives of the cost function to improve transient performance. For example, Newton methods utilize the second derivative of J to eliminate convergence dependence on the cost function’s curvature. In smooth ES systems, it has been shown [64], [65], [72] that higher-order derivatives can also be embedded into the controller through suitable design and manipulation of the dither signals μ and the output function y . Naturally, similar extensions are possible in the context of hybrid set-seeking systems by appropriately modifying the flow maps in (35b) or (41a). While we do not explore these manipulations in detail, we refer the reader to [64], [65], [72], [116], and [117] as well as to an example in the next section, where an ES algorithm switches between Newton- and gradient-based dynamics to enhance performance.

EXAMPLES AND APPLICATIONS OF HYBRID SET-SEEKING SYSTEMS

The study of adaptive seeking systems with logic modes, event-based rules, and set-valued dynamics opens new avenues for research at the intersection of real-time model-free control and computer science, with applications in CPSs. In particular, equipped with the previous theoretical tools, in this section, we now study different types of hybrid set-seeking algorithms, including systems that can be modeled as logic-based switching ES, state-based switched ES, and reset-based ES. By working with well-posed hybrid systems, we show how different

“if-then” rules can be systematically incorporated into set-seeking algorithms while providing stability and robustness guarantees that parallel those existing for continuous-time systems.

Constrained Seeking via Set-Valued Dynamics

By considering (35) with $x_{uz} = \hat{u}$, $r = 0$, $\xi = \nabla J$ and sets $C_{u,z} = \hat{U}$, $D_{u,z} = \emptyset$, $C_z = D_z = \Psi = \hat{G}_\delta = \emptyset$, we can apply Theorem 3 to establish stability properties for a broad class of novel continuous-time set-valued set-seeking systems, where the target decision-making dynamics evolve according to

$$\dot{\hat{u}} \in f(\hat{u}, \nabla J(\hat{u})), \quad \hat{u} \in \hat{U}. \quad (43)$$

In particular, set-valued algorithms of the form (43) are common in decision-making problems with constraints due to the use of inner parametric optimization loops [118], [119], [120], projections [121], or state-dependent switching rules [122]. The following example considers a set-valued seeking algorithm inspired by population dynamics [118] and also illustrates the discussion on Remark 1.

Example 10 Nash Equilibrium-Seeking in Population Games)

Consider a multiagent system with three agents having payoffs of the form

$$J_1 = u_1(-u_3 + u_2), J_2 = u_2(-u_3 + u_1), J_3 = u_3(-u_2 + u_1) \quad (44)$$

and actions $u = (u_1, u_2, u_3) \in \mathbb{R}^3$ restricted to evolve in the simplex set $\hat{U} := \{u \in \mathbb{R}_{\geq 0}^3 : u_1 + u_2 + u_3 = 1\}$, which is compact and convex. The agents seek to optimize their own payoff J_i by controlling their own action u_i , subject to the coupled constraint imposed by the set \hat{U} , resulting in a generalized NE-seeking problem. Since the vector of partial gradients $\nabla J = (\nabla_1 J_1, \nabla_2 J_2, \nabla_3 J_3)$ has entries given by $\nabla_1 J_1 = -u_3 + u_2$, $\nabla_2 J_2 = -u_3 + u_1$, and $\nabla_3 J_3 = -u_2 + u_1$, the NE-seeking problem can also be studied as a rock-paper-scissors population game with three strategies [118]. The unique NE of this game is given by $u^* = ((1/3), (1/3), (1/3))$. To learn u^* using only measurements of J_i , players can implement (35) with the dynamics

$$\dot{\hat{u}} \in -k\hat{u} + k\{w \in \hat{U} : w^\top \xi = \max_{\hat{u} \in \hat{U}} \hat{u}^\top \xi\}, \quad \hat{u} \in \hat{U} \quad (45)$$

where $k > 0$. By construction, the set-valued dynamics (45) satisfy condition (a) of Assumption 1. Moreover, item (b) of Assumption 1 can be verified by studying the target system $\dot{\hat{u}} \in -\hat{u} + \{w \in \hat{U} : w^\top \nabla J = \max_{\hat{u} \in \hat{U}} \hat{u}^\top \nabla J\}$, $\hat{u} \in \hat{U}$, which renders UGAS the NE u^* . Indeed, this stability property holds for a broader class of population games that satisfy $(u_a - u_b)^\top (\nabla J(u_a) - \nabla J(u_b)) \leq 0$ for all $u_a, u_b \in \hat{U}$ [118, Thm. 7.2.7], which allows us to verify Assumption 1 without the explicit knowledge of the costs J_i . Figure 14 shows three different trajectories generated by the set-seeking dynamics (35) using (45), evolving on the set \hat{U} , and converging to

a neighborhood of the NE u^* . Note that the slow-frequency oscillations shown in the inset in Figure 14(a) are not induced by the proving signals μ , but instead by the discontinuous nature of the dynamics (45). On the other hand, the oscillatory behavior on the state ξ , induced by the probing signals, is clearly shown in Figure 14(b). For this simulation, the parameters were selected as $\varepsilon_a = 0.01$, $k = 0.005$, $k_f/\varepsilon_f = 10$, and $\omega = (18.5, 20, 25.3) \times 10^4$.

Set-valued dynamical systems are also common in the analysis of projected gradient-based algorithms. The following example illustrates this application in a seeking algorithm that satisfies Assumption 2, where the filter and the decision-making dynamics are designed to operate on the same timescale.

Example 11 (Set Seeking With Projected Dynamics)

Consider a model-free optimization problem of the form (29), where J is strictly convex, and \hat{U} is a closed and convex set with nonempty interior, describing safety constraints that need to be satisfied for all time. To tackle this problem, we can consider the set-seeking system (39) with $D = \emptyset$ and

$$\dot{\hat{u}} = k\mathcal{P}_{\hat{U}}(\hat{u})(-\xi), \quad \hat{u} \in \hat{U} \quad (46)$$

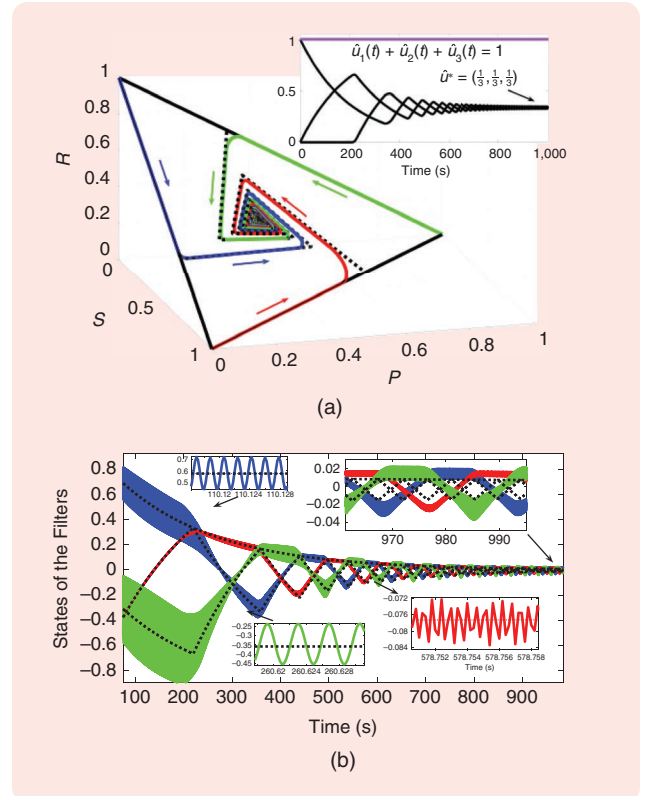


FIGURE 14 The trajectories of the set-seeking system studied in Example 10. (a) Three simulations with different initial conditions $\hat{u}(0)$, evolving on the simplex \hat{U} and converging to a neighborhood of u^* . (b) The time evolution of the filter states for one particular initialization of the algorithm.

where $k > 0$ and $\mathcal{P}_{T_{\hat{U}}(\cdot)}$ projects the vector $-\xi$ onto the tangent cone of \hat{U} at \hat{u} . In general, this projection leads to discontinuous vector fields, necessitating tools from nonsmooth analysis (see “Continuous-Time Set-Valued Dynamical Systems”): namely, the use of Krasovskii regularizations (S5). Indeed, under mild regularity assumptions on J and \hat{U} (see [70]), the Krasovskii solutions to (35) and (46) coincide with their standard (Carathéodory) solutions. Therefore, the resulting target average system (40), given by

$$\dot{\hat{u}} = \mathbb{P}_{T_{\hat{U}}(\hat{u})}(-\xi), \quad \dot{\xi} = -k_f(\xi - \nabla J(\hat{u})), \quad \hat{u} \in \hat{U} \quad (47)$$

renders the set $\mathcal{A} = \{u^* \times \{\nabla J(u^*)\}\} \in \hat{U} \times \mathbb{R}^n$ UGAS for k_f sufficiently small and $\kappa = 1$; see [70, Lemma 10]. It follows that Assumption 2 is satisfied with $\Lambda_{\xi,c} = \Lambda_{\xi,d} = \mathbb{R}^n$, and $k_f > 0$ sufficiently small [70, Lemma 10]. Therefore, by Theorem 4, the model-free set-seeking dynamics (39) assure convergence to a

neighborhood of the optimal solution u^* in \hat{U} . Figure 15 shows a trajectory generated by (39) using the dynamics (46), approximately tracking the time-varying minimizer of a slowly varying cost function of the form $J(u) = (u_1 - q_1)^2 + (u_2 - q_2)^2$, where $t \mapsto q(t) := (q_1(t), q_2(t))$ is a slowly varying parameter and where the feasible set \hat{U} corresponds to a 45°-rotated box centered at the origin. As observed, the trajectories \hat{u} generated by the set-seeking dynamics remain in \hat{U} for all time. The simulation used the parameters $\varepsilon_a = 0.01$, $\omega = (100, 150)$, and $k_f = 1$, $k = 0.2$. The state q was generated by a linear oscillator of the form (14) with frequency equal to 0.001. We note that ES systems with projected dynamics can also be studied using Lipschitz vector fields [70, Sec. IV.A] as well as smooth Riemannian projections for decision-making problems defined on manifolds [106], [112].

Example 12 (Seeking with Unknown Constraints)

In Examples 10 and 11, the constraints that describe the decision-making problem (29) were known a priori by the practitioner. However, in certain applications, the precise mathematical form of the function that describes the constraints is also unknown. In particular, consider the constrained ES problem (29), where J is strictly convex and

$$\hat{U} = \{u \in \mathbb{R}^n : c(u) \leq 0\} \quad (48)$$

where $c : \mathbb{R}^n \rightarrow \mathbb{R}$ is an unknown smooth convex function that is accessible only via evaluations and such that \hat{U} has nonempty interior. To tackle the seeking problem with unknown constraints, we can consider the dynamics (35) with $D = \emptyset$ and the set-valued rule

$$\dot{\hat{u}} \in k \cdot \begin{cases} \{-\xi_J\} & \text{if } c(\hat{u}) < 0 \\ K(\xi_J, \xi_c) & \text{if } c(\hat{u}) = 0 \\ \{-\xi_c\} & \text{if } c(\hat{u}) > 0 \end{cases} \quad (49)$$

where $k > 0$, $K(\xi_J, \xi_c) := \{z \in \mathbb{R}^n : z = \lambda \xi_J + (1 - \lambda) \xi_c, \lambda \in [0, 1]\}$, and where the states ξ_J and ξ_c are generated by low-pass filters of the form (32), with inputs $y = J(u)$ and $y = c(u)$, respectively. It follows that condition (a) of Assumption 1 holds [122], and the target decision-making system (33) associated with (49) is given by

$$\dot{\hat{u}} \in k \cdot \begin{cases} \{-\nabla J(\hat{u})\} & \text{if } c(\hat{u}) < 0 \\ K(\hat{u}) & \text{if } c(\hat{u}) = 0 \\ \{-\nabla c(\hat{u})\} & \text{if } c(\hat{u}) > 0 \end{cases} \quad (50)$$

where $K(\hat{u}) := \{z \in \mathbb{R}^n : z = \lambda \nabla J(\hat{u}) + (1 - \lambda) \nabla c(\hat{u}), \lambda \in [0, 1]\}$, which also satisfies item (b) of Assumption 1 [122]. Therefore, by Theorem 3, the trajectories of (35) with dynamics (49) converge to a neighborhood of u^* , the optimizer of J in \hat{U} . Figure 16 depicts a trajectory of the seeking dynamics generated by (49). The trajectory is initialized outside the feasible set and tracks the minimizer of a slowly varying cost function $J(u)$. In this case, $c(u) = a^\top u + b$, where $a = [-1 - 2]$,

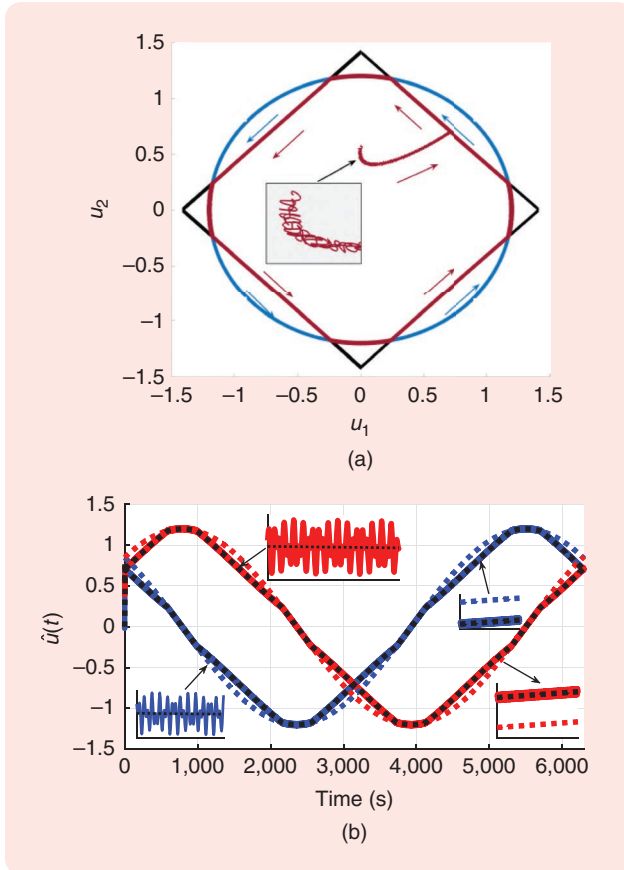


FIGURE 15 The trajectories of the set-seeking system studied in Example 11. (a) The evolution of the trajectories \hat{u} (red) within the feasible set \hat{U} , corresponding to the rotated square. The blue trajectory represents the slowly varying minimizer of the cost function J in \mathbb{R}^2 . As shown, the red trajectories closely follow the blue trajectory inside the feasible set. (b) The trajectories of \hat{u}_1 (blue) and \hat{u}_2 (red), the optimal solution u^* (black dotted), and the time-varying maximizers in \mathbb{R}^2 (dotted red and blue). As shown in the inset, the trajectory \hat{u} closely tracks the constrained optimal solution u^* .

$b = 2$, and J is the same as in Example 11. The parameters used in the simulation were $\varepsilon_a = 0.01$, $\omega = (100, 250)$, $k_f = 1$, and $k = 0.1$. As shown, the trajectory exhibits two phases. Initially, it converges to the feasible set while disregarding feedback from the cost function. Once inside the feasible set, the seeking dynamics drive it to a neighborhood of the optimal solution u^* . As is typical, the discontinuous nature of (49) induces chattering along the boundary of the feasible set \hat{U} . In the next sections, we show how to avoid this chattering by introducing time regularization via resetting timers or spatial regularization via switching conditions that induce hysteresis.

Seeking With Arbitrarily Fast Switching Dynamics

Let $\mathcal{Q} := \{1, \dots, q_{\max}\}$, where $q_{\max} \in \mathbb{Z}_{\geq 1}$, and consider a collection of set-valued mappings $f_q: \mathbb{R}^n \times \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, with main state $\hat{u} \in \mathbb{R}^n$, auxiliary state $q \in \mathcal{Q}$, and target decision-making dynamics (33) with $\xi = \nabla J$, given by

$$\dot{\hat{u}} \in f_q(\hat{u}, \nabla J(\hat{u})), \quad \dot{q} = 0, \quad (\hat{u}, q) \in \hat{U} \times \mathcal{Q} \quad (51a)$$

$$\hat{u}^+ = \hat{u}, \quad q^+ \in \mathcal{Q}, \quad (\hat{u}, q) \in \hat{U} \times \mathcal{Q} \setminus \{q\}. \quad (51b)$$

The update rule (51a) describes a gradient-based differential inclusion parameterized by a constant q . On the other hand, the dynamics (51b) model the switches of q to a different mode in the set \mathcal{Q} while keeping the main state \hat{u} constant. Figure 17 provides a typical graphical representation of logic-based systems that switch between different modes. Such types of representation are connected to the well-known hybrid automata, which can also be written as an HDS (17); see [8, Ex. 1.10]. In this sense, the tools presented in this article can also be used for the synthesis and analysis of seeking systems based on hybrid automata. Nonsmooth systems of the form (51a) can also be considered for the study of ES algorithms with finite-time or fixed-time (practical) convergence properties [123], [124].

Let the data of (51) satisfy the conditions in Assumption 1)-(a) so that it can be used as a suitable target hybrid




FIGURE 16 The trajectories \hat{u} of the seeking dynamics in Example 12. After converging to the feasible set, the trajectories approximately track the slowly varying optimizer u^* .

decision-making system. The following proposition, adapted from [91] and [125], establishes that item (b) in Assumption 1 also holds under arbitrarily fast switching of q , provided that there exists a common Lyapunov function that certifies the stability of (51a) with respect to the set \mathcal{O} across all modes $q \in \mathcal{Q}$.

Proposition 1 (Seeking Under Fast Switching)

Suppose that for each $q \in \mathcal{Q}$, the mapping f_q in (51a) is OSC, LB, and convex-valued. Suppose also that there exists $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$, a continuous positive definite function $W: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, and a continuously differentiable function $V: \text{dom}(V) \rightarrow \mathbb{R}_{\geq 0}$, with $\hat{U} \subset \text{dom}(V)$, such that for all $\hat{u} \in \hat{U}$, all $\tilde{f} \in f_q$, and all $q \in \mathcal{Q}$:

$$\alpha_1(|\hat{u}|_0) \leq V(\hat{u}) \leq \alpha_2(|\hat{u}|_0) \quad (52a)$$

$$\langle \nabla V(\hat{u}), \tilde{f} \rangle \leq -W(|\hat{u}|_0). \quad (52b)$$

Then, the differential inclusion

$$\dot{\hat{u}} \in \hat{F}(\hat{u}, \nabla J(\hat{u})) := \overline{\text{co}} \bigcup_{q \in \mathcal{Q}} f_q(\hat{u}, \nabla J(\hat{u})), \quad \hat{u} \in \hat{U} \quad (53)$$

has the structure of (33) with the state $x_{uz} = \hat{u}$, $r = 0$, $\xi = \nabla J$, and sets $C_{u,z} = \hat{U}$, $D_{u,z} = \emptyset$, $C_z = D_z = \Psi = \hat{G}_\delta = \emptyset$. Moreover, system (53) satisfies Assumption 1, and for each non-purely discrete solution (\hat{u}, q) of the switching system (51), \hat{u} is also a solution to (53). \square

The main takeaway from Proposition 1 is that seeking algorithms can be designed and analyzed to switch arbitrarily fast between a finite number of modes, provided that their continuous-time target decision-making dynamics share a common Lyapunov function. In this case, all (non-discrete) solutions of the switching decision-making dynamics (51) can be generated by the differential inclusion (53). Note that common Lyapunov functions satisfying the conditions of Proposition 1 emerge frequently in optimization algorithms since the cost function J usually serves as a common Lyapunov function for different optimization dynamics. For example, this is the case in the standard gradient ascent and Newton methods as well as related dynamics, such as the Newton variable structure algorithm or the




FIGURE 17 Logic-based ES allows one to switch between different algorithms or “operating modes.”

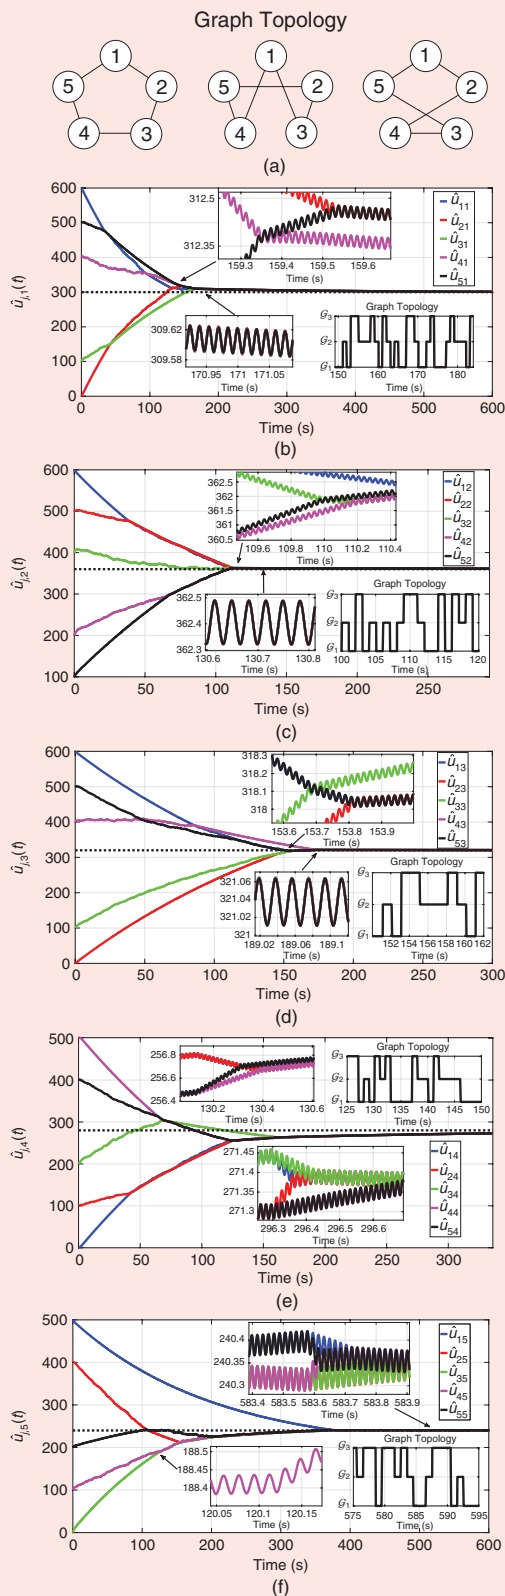


FIGURE 18 (a)–(f) The trajectories of the hybrid set-seeking algorithm (55) with communication graphs allowed to switch every 0.1 s.

continuous Jacobian matrix transpose algorithm, for quadratic cost functions (see [126, Ch. 2]). A common Lyapunov function also emerges frequently in algorithms for distributed optimization and learning in multiagent systems with undirected switching communication topologies [127], [128].

Example 13 (Switching Multiagent Seeking Systems)

Consider a multiagent system where N agents seek to cooperatively solve the following consensus-optimization problem:

$$\min J(u) = \sum_{i=1}^N J_i(u_i), \quad u_i \in \mathbb{R}^n, \quad s.t. \quad u_1 = u_2 = \dots = u_N \quad (54)$$

Agents can communicate only to neighbors, who are characterized by an undirected graph that is allowed to change topology arbitrarily fast but which is connected at all times. Since (54) is a well-studied problem in the literature whenever the gradients ∇J_i are known to the agents, there is a plethora of potential target nominal dynamics (33) that could be considered for the design of an ES system [129]. To illustrate the use of Krasovskii solutions, we consider the dynamics

$$\dot{\hat{u}}_i \in \hat{F}_{\delta,i}(\hat{u}, \xi_i) = \frac{k}{\delta} \sum_{j \in N_i(t)} \overline{\text{sign}}(\hat{u}_j - \hat{u}_i) - k\gamma \xi_i \quad (55)$$

where $\overline{\text{sign}}$ is the set-valued map defined in (S5); that is

$$\overline{\text{sign}}(z) = \begin{cases} \{1\} & \text{if } z > 0 \\ [-1, 1] & \text{if } z = 0 \\ \{-1\} & \text{if } z < 0 \end{cases} \quad (56)$$

To generate ξ_i , each agent implements its own autonomous derivative estimator (30) using its own cost function J_i as input. In this way, agents share information with neighbors only via the dynamics (55). By indexing all possible graph topologies using a set of logic modes $Q = \{1, 2, \dots, q_{\max}\}$, the resulting system has the form (51), where each fixed logic mode q represents a particular graph topology. Figure 18 presents a numerical example with five agents, showing the evolution in time of the states \hat{u}_i , which converge to the optimal solution. The cost functions used in the simulations are given by $J_i = (1/2)u_i^\top u_i - 100b_i^\top u_i$, where $b_1 = (1, 2, 3, 4, 5)^\top$, $b_2 = (5, 4, 3, 2, 1)^\top$, $b_3 = (2, 3, 4, 5, 1)^\top$, $b_4 = (3, 4, 5, 1, 2)^\top$, $b_5 = (4, 5, 1, 2, 3)^\top$, and the ES parameters are $\varepsilon_a = 0.8$, $k = 0.05$, $k_f = 0.1$, $\varepsilon_l = 0.1$, $\omega = 500(5/6, 3/8, 2/3, 4/5, 7/10)$, $\alpha = 15$, and $\gamma = 0.1$. It can be verified that, in this example, the target nominal system (55) with $\xi_i = \nabla J_i$ and sufficiently large δ admits a common Lyapunov function across all connected undirected graphs [129]. It follows that Proposition 1 holds, and therefore, by Theorem 3, the set-seeking dynamics achieve practical convergence to the set of solutions of (54).

Seeking With Average Dwell-Time Switching

When a common Lyapunov function does not exist for the target switching system (51), it is still possible to stabilize \mathcal{O} in a semiglobal practical sense by restricting how fast the

state q switches between modes. For example, this can be achieved by imposing a dwell-time or an average dwell-time constraint on the switching signal. As discussed in Example 3, these types of constraints can be enforced by using a hybrid automaton with state $\tau_1 \in \mathbb{R}_{\geq 0}$ and dynamics

$$\dot{\tau}_1 \in [0, \eta_1], \quad \tau_1 \in [0, N_0] \quad (57a)$$

$$\tau_1^+ = \tau_1 - 1, \quad \tau_1 \in [1, N_0] \quad (57b)$$

where $\eta_1 \in \mathbb{R}_{> 0}$ and $N_0 \in \mathbb{R}_{\geq 1}$. In particular, every solution to (57) has a hybrid time domain E with the property that for each of its elements $(s, i), (t, j) \in E$ with $s + i \leq t + j$, the following inequality holds:

$$N(t, s) := j - i \leq \eta_1(t - s) + N_0 \quad (58)$$

where $N(t, s)$ denotes the number of jumps in the interval $[s, t]$; see [80]. Therefore, when considering the joint hybrid dynamics (51) and (57), the hybrid time domains of the resulting system also satisfy inequality (58), thus effectively imposing a bound on how frequently q can switch in (51). When $N_0 = 1$, inequality (58) reduces to the well-known *dwell-time* condition; see [85]. When $N_0 \in \mathbb{Z}_{> 1}$, inequality (58) describes an *average dwell-time* condition. It turns out that any hybrid time domain satisfying (58) can be generated by the hybrid automaton (57); see [130, Proposition 1.1]. Therefore, these dynamics can be used to study switching set-seeking systems under dwell-time and average dwell-time constraints on the switching signal.

Proposition 2 (Seeking With Slow Switching)

Suppose that for each $q \in \mathcal{Q}$, the mapping f_q in (51a) is OSC, LB, and convex-valued. Moreover, suppose that for each $q \in \mathcal{Q}$, the \hat{u} -dynamics in system (51a) render UGAS the set \mathcal{O} . Then, the HDS generated by (51) and (57) has the structure of (33) with $r = 2$, $\delta = \eta_1$, $\hat{z} := (q, \tau)^\top$, $C_{u,z} := \hat{\mathbb{U}} \times \mathcal{Q} \times [0, N_0]$, $D_{u,z} := \hat{\mathbb{U}} \times \mathcal{Q} \times [1, N_0]$, $\Psi := \mathcal{Q} \times [0, N_0]$, $\hat{F}_\delta := f_q \times \{0\} \times [0, \eta_1]$, and $\hat{G}_\delta := \{\hat{u}\} \times \mathcal{Q} \setminus \{q\} \times [\tau_1 - 1]$. Moreover, this HDS satisfies all the items in Assumption 1. \square

The result of Proposition 2, which follows from [8, Corollary 7.28], enables the study of switching seeking dynamics where each seeking mode is able to individually stabilize the set \mathcal{O} and where the switching between modes is “sufficiently slow” on average. Such types of problems might emerge in multiagent systems with switching directed communication graphs for which a common Lyapunov function does not exist as well as in seeking problems with slow variations in the landscape of the cost function or problems that require a persistent switching of the adaptation gains to achieve a desired performance.

The result of Proposition 2 can also be extended to ES systems with switching cost functions, even when the cost functions do not share the same minimizer $u_q^* \in \mathbb{R}^n$, as in Example 3. In this case, when $f_q(\cdot)$ is single-valued and

continuous, for a sufficiently large compact set K of initial conditions, the (practical) stability properties of the resulting switching dynamics (51) and (57) can be studied with respect to the Omega-limit set from K , denoted $\Omega(K)$, which is given by $\Omega(K) = \bigcup_{q \in \mathcal{Q}} \Omega_q$, where

$$\Omega_q = \bigcap_{j \in \mathbb{Z}_{\geq 0}} \overline{R_0 \left(\left(\{u_q^*\} + \frac{1}{j+1} \mathbb{B} \right) \times \{q\} \times [0, N_0] \right)} \quad (59)$$

and where $R_0(K)$ denotes the reachable set from K . An illustration of this set is presented in Figure 6 for the source-seeking problem with persistent surveillance. For further details on switching systems with distinct equilibrium points, we refer the reader to [86].

Seeking With Switching Unstable Modes

In many applications, such as source-seeking missions, seeking algorithms are often implemented in highly adversarial and dynamic environments. In these settings, intermittent communication, sensing, or actuation failures are unavoidable due to external factors. Such intermittent failures can lead to unstable behaviors, which, however, can be corrected if the system is able to return to its nominal operating condition “sufficiently” often. This scenario can be analyzed using the framework of hybrid set-seeking systems by considering both stable and unstable modes in (51). Specifically, the compact set $\mathcal{Q} := \{1, \dots, q_{\max}\}$ can be partitioned as $\mathcal{Q} = \mathcal{Q}_u \cup \mathcal{Q}_s$, where the modes $q \in \mathcal{Q}_s$ represent stable dynamics and the modes $q \in \mathcal{Q}_u$ represent unstable dynamics. For this type of system, stability can be guaranteed as long as the total activation time $T(s, t)$ of unstable modes during any time interval (s, t) satisfies the bound

$$T(s, t) := \int_s^t \mathbb{I}_{\mathcal{Q}_u}(q(r, j(r))) dr \leq T_0 + \eta_2(t - s) \quad (60)$$

where $\eta_2 \in [0, 1)$, $T_0 \in \mathbb{R}_{\geq 0}$, and $j(r)$ is the minimum j such that (r, j) belong to the domain of the solution. Similarly to (57) and (58), the average activation-time constraint (60) can also be imposed on the switching signal q by using the following hybrid automaton with state $\tau_2 \in \mathbb{R}_{\geq 0}$ and dynamics:

$$\dot{\tau}_2 \in [0, \eta_2] - \mathbb{I}_{\mathcal{Q}_u}(q), \quad \tau_2 \in [0, T_0] \quad (61a)$$

$$\tau_2^+ = \tau_2, \quad \tau_2 \in [0, T_0] \quad (61b)$$

where $\mathbb{I}_{\mathcal{Q}_u}(\cdot)$ corresponds to the standard indicator function. In particular, every solution to the HDS given by (51) and (61) has a hybrid time domain E , such that for each of its elements $(s, i), (t, j) \in E$ with $s + i \leq t + j$ and signal $q: E \rightarrow \mathcal{Q}$, we have that the bound (60) holds. In fact, each hybrid time domain E with elements $(s, i), (t, j) \in E$ with $s + i \leq t + j$, satisfying (60), can be generated by the hybrid dynamics (51) and (61) [91, Lemma 7]. Therefore, by combining dynamics (51), (57), and (61), we can design different types

of hybrid seeking algorithms (33) that satisfy **Assumption 1** even when the decision-making algorithms operate under a persistent loss of communication, sensing, or actuation.

Proposition 3 (Seeking with Unstable Dynamics)

Suppose that for each $q \in \mathcal{Q}$, the mapping f_q in (51) is OSC, LB, and convex-valued and that there exist $\chi \in \mathbb{R}_{\geq 1}$, $\alpha_{1,q}, \alpha_{2,q} \in \mathcal{K}_{\infty}$, $\lambda_s \in \mathbb{R}_{>0}$, $\lambda_u \in \mathbb{R}_{>0}$ and smooth functions $V_q: \hat{\mathcal{U}} \rightarrow \mathbb{R}_{\geq 0}$, such that the following inequalities hold:

$$\alpha_{1,q}(|\hat{u}|_0) \leq V_q(\hat{u}) \leq \alpha_{2,q}(|\hat{u}|_0), \quad \forall (q, \hat{u}) \in \mathcal{Q} \times \hat{\mathcal{U}}. \quad (62a)$$

$$\langle \nabla V_{q_s}, \tilde{f} \rangle \leq -\lambda_s V_{q_s}(\hat{u}), \quad \forall (q_s, \hat{u}, \tilde{f}) \in \mathcal{Q}_s \times \hat{\mathcal{U}} \times f_{q_s}. \quad (62b)$$

$$\langle \nabla V_{q_u}, \tilde{f} \rangle \leq \lambda_u V_{q_u}(\hat{u}), \quad \forall (q_u, \hat{u}, \tilde{f}) \in \mathcal{Q}_u \times \hat{\mathcal{U}} \times f_{q_u}. \quad (62c)$$

$$V_p(\hat{u}) \leq \chi V_q(\hat{u}), \quad \forall (q, p), \hat{u} \in \mathcal{Q} \times \hat{\mathcal{U}}. \quad (62d)$$

If the previous parameters satisfy the following inequality:

$$\lambda_s > \eta_1 \log(\mu) + \eta_2(\lambda_s + \lambda_u) \quad (63)$$

then, the HDS composed by (51), (57), and (61) has the structure of (33) with $r=3$ $\hat{z} := (q, \tau_1, \tau_2)$, $C_u = D_u = \hat{\mathcal{U}}$, $C_z = \mathcal{Q} \times [0, N_0] \times [0, T_0]$, $D_z = \mathcal{Q} \times [1, N_0] \times [0, T_0]$, and $\Psi := \mathcal{Q} \times [0, N_0] \times [0, T_0]$. Moreover, this HDS satisfies the items in **Assumption 1**, and the solutions' hybrid time domains satisfy the bounds (58) and (60). \square

The Lyapunov-based conditions (62) and the inequality (63) are common in the literature of switching systems with unstable modes. In this sense, Proposition 3 simply recasts such conditions using the hybrid automaton (57) and the hybrid monitor (61) such that, when interconnected with the target nominal dynamics (51), the resulting system has the form (33) and satisfies all the conditions needed for the analysis of switching set-seeking systems with unstable modes. The verification of the assumptions in Proposition 3 can be straightforward for some problems and algorithms, particularly when the cost J is quadratic, which leads to linear gradients ∇J . In fact, conditions (62a) and (62b) simply ask for UGAS of \mathcal{O} for each mode f_{q_s} , where $q_s \in \mathcal{Q}_s$. Similarly, condition (62c) simply asks that solutions of the differential inclusions associated with the unstable modes $q \in \mathcal{Q}_u$ have no finite escape times. Additionally, condition (62d)

is easily satisfied if the Lyapunov functions are quadratic. In this case, conservative estimates of the parameters λ_s , λ_u , and μ can be used to design the parameters η_1 and η_2 such that (63) holds.

To illustrate Proposition 3 as well as the discussion of Remark 4, we consider an ES algorithm with standard gradient dynamics $\dot{u} = -\xi$, operating in an environment where the sign of the measurements of y is corrupted in real time. Specifically, we consider a quadratic cost function $J(u) = 0.01u^T Q u + b^T u$, where $Q = [1, 1/2; 1/2, 3/2]$, $b = (0.1, 0.1)$, and we define $q \in \mathcal{Q} = \{-1, 1\}$, where $q = -1$ indicates that the output has been corrupted. See Figure 19 for a block diagram representation. Figure 20 shows the time evolution of the resulting hybrid set-seeking dynamics with two different average activation times of the mode $q = -1$, using $\omega = (810, 420)$, $\varepsilon_a = 0.1$, and $\varepsilon_f = 1$. As observed, when the activation time is sufficiently short (that is, η_2 is small), the seeking system successfully converges to the optimal point. In contrast, when η_2 is too large, the system becomes unstable. For related applications of ES in the context of deceptive attacks on gradient estimates, we refer the reader to [131]. Note that by redefining the set \mathcal{Q} as $\mathcal{Q} = \{0, 1\}$, the previous model can also capture ES systems operating under sporadic feedback. In this case, source seeking can still




FIGURE 19 A scheme of an ES controller under persistent multiplicative attacks ($q \in \{-1, 1\}$) on the output measurements.





FIGURE 20 ES dynamics in \mathbb{R}^2 , under persistent attacks on the gradient estimation mechanism. The intensity of the attacks is modeled using (61). As shown in (a), when η_2 is sufficiently small, the ES dynamics preserve stability. (b) shows the unstable behavior that emerges when η_2 is moderate.

be achieved, provided that the condition (60) is satisfied. Multiagent seeking problems involving persistent adversarial agents can also be analyzed using similar techniques; see [91, Sec. 6.3].

Seeking With Slowly Varying or Jumping Parameters

In many practical seeking problems, the parameters associated with the plant or the algorithm vary slowly over time. For instance, this occurs when the response map J is time varying or exhibits mild jumps on some of its parameters. Although slowly varying seeking problems have been studied under the assumption that the target decision-making algorithm is input-to-state stable (ISS) with respect to the time derivatives of the varying signals in the cost [123], [132], [133], the ISS property may not necessarily hold for a broad class of decision-making algorithms (33). Nevertheless, such problems can still be addressed using the hybrid set-seeking systems framework studied in this article. In particular, by letting q represent the time-varying parameter (or its derivative, depending on the application), redefining \mathcal{Q} in (51) as a compact but not necessarily finite set, and by allowing the flow dynamics in the hybrid system (51) to depend explicitly on q , that is, $\dot{\hat{u}} \in f(\hat{u}, \nabla J(\hat{u}), q)$, the slowly varying and weakly jumping behavior of q within \mathcal{Q} can be effectively modeled via the inclusions

$$\dot{q} \in \eta_3 \mathbb{B} \quad \text{and} \quad q^+ \in q + \eta_3 \mathbb{B}, \quad \text{where} \quad \eta_3 \in \mathbb{R}_{>0} \quad (64)$$

which replace the q -dynamics in (51). In this scenario, we assume that for each $q \in \mathcal{Q}$, there exists an optimal compact set $\mathcal{O}_q \subset \hat{\mathcal{U}}$ that minimizes J , and we employ a simple average dwell-time automaton of the form (57) to rule out purely discrete-time solutions in the hybrid system. The following proposition is a direct consequence of [8, Corollary 7.27].

Proposition 4 (Seeking Under Slowly Varying Parameters)

Let $\mathcal{O}_q \subset \hat{\mathcal{U}}$, and suppose that for each $q \in \mathcal{Q}$, the mapping f_q in (51a) satisfies condition (a) in Assumption 1 with the flow map given by $f(\hat{u}, \nabla J(\hat{u}), q)$ and the q -dynamics given by (51). Suppose also that for each $q \in \mathcal{Q}$, the dynamics $\dot{\hat{u}} \in f(\hat{u}, \nabla J(\hat{u}), q)$ render the set \mathcal{O}_q UGAS. Then, this HDS combined with the dynamics (57) and (64) has the structure of (33) with $r = 2$, $\delta = \eta_3$, $\hat{z} := (q, \tau_1)$, $C_u = D_u = \hat{\mathcal{U}}$, $C_z = \mathcal{Q} \times [0, N_0]$, $D_z := \mathcal{Q} \times [1, N_0]$, $\mathcal{O} := \{(\hat{u}, q) : \hat{u} \in \mathcal{O}_q, q \in \mathcal{Q}\}$, and $\Psi := \mathcal{Q} \times [0, N_0]$. Moreover, this HDS satisfies all the items in Assumption 1. \square

An application of Proposition 4 was already illustrated in Example 11 and in Figure 15, where a slowly varying parameter q defines the cost function J . In particular, when $\dot{q} = 0$, the state q remains constant, and the trajectories \hat{u} of the system converge to a neighborhood of the optimal point $u^*(q)$. Conversely, when $\dot{q} \in \eta_3 \mathbb{B}$, with \mathcal{Q} compact and $\eta_3 > 0$ sufficiently small, the tracking error $\tilde{u} = \hat{u} - u^*(q)$ converges to a neighborhood of zero.

Seeking in Noncontractible Spaces: Source Seeking With Obstacles

Smooth controllers, and dynamical systems in general, face fundamental limitations in stabilization problems when operating in complex spaces, particularly when obstacles are present. This is a consequence of the fact that obstacles usually preclude the contractibility of the operational space, which is a necessary condition for global stabilization [110] under smooth feedback laws. For example, it has been shown that in obstacle avoidance problems, different types of controllers, including those based on navigation functions or control barrier functions, can generate “deadlocks” [134] and undesirable spurious equilibria in the closed-loop system [135]. While it is well known that such undesirable points—often forming a set of measure zero—can be “pushed out” to the boundary of the feasible set [136], introducing even arbitrarily small perturbations from measurement noise or external disturbances in the closed-loop system can cause these problematic points in the state space to no longer be of measure zero, thereby precluding (almost) global convergence guarantees. However, these obstructions to robust global stabilization under obstacles, which also apply to optimization dynamics, can be overcome by using hybrid set-seeking systems [83], [111], [112], [137]. For example, consider a vehicle operating on the plane according to the simple point-mass dynamics

$$\dot{p} = \alpha, \quad p \in \mathbb{R}^2 \quad (65)$$

where $p = (p_1, p_2)$ is the position, and $\alpha \in \mathbb{R}^2$ is the control input. The main goal of the vehicle is to “discover” in real time the location $\mathcal{O} = \{u\} \in \mathbb{R}^2$ where a potential field J attains its maximum value, using only intensity measurements, while simultaneously avoiding an obstacle $\mathcal{N} \subset \mathbb{R}^2$ in the plane. As discussed in [110] and [137], this problem cannot be robustly solved using target decision-making algorithms (33) characterized by smooth feedback dynamics due to the topological obstructions introduced by the obstacle. Therefore, to achieve source seeking and obstacle avoidance, we consider a hybrid source-seeking controller with the feedback law given by

$$\alpha = \varepsilon_a \omega \mathbf{R}_0 \mu + \xi. \quad (66)$$

Assuming that we have access to measurements of J at the points p [to be used by the filter (32)] and $p - \varepsilon_a \mu$ (to be used in the flow and jump sets), which is possible via the collocation of sensors, we consider the change of variable $\hat{u} = p - \varepsilon_a \mu$ to obtain the dynamics $\dot{\hat{u}} = \xi$, which puts the system into the form (35b) with $\hat{F} = \xi$ and $y = J(\hat{u} + \varepsilon_a \mu)$. To design the target hybrid decision-making dynamics (33), and inspired by [137], we let $q \in \mathcal{Q} = \{1, 2\}$ be a logic state that remains constant during flows and that is used to construct the following mode-dependent function:

$$\hat{J}_q(\hat{u}) = -J(\hat{u}) + B_q(\hat{u}) \quad (67)$$

where B_q is a mode-dependent barrier function to be designed. We restrict our attention to bounded obstacles $\mathcal{N} \subset \mathbb{R}^n$ for which there exists $p_0 = (p_{0,1}, p_{0,2}) \in \mathbb{R}^2$, $\rho \in \mathbb{R}_{>0}$ and $\delta \in \mathbb{R}_{>0}$ such that $\mathcal{N} \subset p_0 + \rho\mathbb{B}$ and $(p_0 + 2\rho\sqrt{2}\mathbb{B}) \cap (u^* + \delta\mathbb{B}) = \emptyset$; that is, the obstacle \mathcal{N} is contained in a ball of radius ρ , centered at the point p_0 , and located sufficiently far away from the target p^* . To achieve obstacle avoidance, we divide the space into two different regions, each region indexed by a mode $q \in \{1, 2\}$ and having a corresponding mode-dependent signal $\hat{J}_q(\cdot)$, defined in (67). The main goal is to design the regions so that the vehicle can synergistically use the two potential fields $\hat{J}_q(\cdot)$ to safely navigate the space. To formalize this idea, consider the set

$$\mathcal{B}_{p_0, \rho} := \{p : \|p - p_0\|_1 \leq 2\rho\sqrt{2}\} \quad (68)$$

which satisfies $\{p_0\} + \rho\mathbb{B} \subset \mathcal{B}_{p_0, \rho} \subset \{p_0\} + 2\rho\sqrt{2}\mathbb{B}$. Additionally, consider the following subsets of \mathbb{R}^2 :

$$\begin{aligned} L_{1a} &:= \{p : p_2 < -p_1 + p_{0,2} + p_{0,1} - 2\rho\sqrt{2}\} \\ L_{1b} &:= \{p : p_2 < p_1 + p_{0,2} + p_{0,1} - 2\rho\sqrt{2}\} \\ L_{2a} &:= \{p : p_2 > p_1 + p_{0,2} + p_{0,1} + 2\rho\sqrt{2}\} \\ L_{2b} &:= \{p : p_2 > -p_1 + p_{0,2} + p_{0,1} + 2\rho\sqrt{2}\} \end{aligned}$$

and let $\mathcal{L}_1 := L_{1a} \cup L_{1b}$, $\mathcal{L}_2 := L_{2a} \cup L_{2b}$, and $\mathcal{U} := \mathcal{L}_1 \cup \mathcal{L}_2$. Figure 21 illustrates the construction of the sets \mathcal{L}_1 and \mathcal{L}_2 , which satisfy $u^* \in \mathcal{L}_1 \cap \mathcal{L}_2$ and also $\mathcal{L}_1 \cap \mathcal{N} = \emptyset$, $\mathcal{L}_2 \cap \mathcal{N} = \emptyset$. In fact, $\mathcal{U} = \mathbb{R}^2 \setminus \mathcal{B}_{p_0, \rho}$. The mode-dependent barrier function B_q is designed such that the function $W(\hat{u}, q) := \hat{J}_q(\hat{u}) + J(\hat{u}^*)$ satisfies the following assumption.

Assumption 4 (Mode-Dependent Barrier Functions)

- a) For each $q \in \mathcal{Q}$, the map $\hat{J}_q : \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is continuously differentiable in \mathcal{L}_q , and as $\hat{u} \rightarrow \infty$ or $\hat{u} \rightarrow \text{bd}(\mathcal{L}_q)$, we have that $\hat{J}_q(\hat{u}) \rightarrow \infty$. Moreover, for every $\hat{u} \in \mathbb{R}^2 \setminus \mathcal{L}_q$, we define $J_q(\hat{u}) := \infty$.

- b) There exist functions $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$ and a proper indicator² \tilde{w} of u^* on \mathcal{U} , such that

$$\alpha_1(\tilde{w}(\hat{u})) \leq \min_{s \in \mathcal{Q}} W(\hat{u}, s) \leq \alpha_2(\tilde{w}(\hat{u})), \quad \forall \hat{u} \in \mathcal{L}.$$

- c) There exists a positive definite function ρ , such that for each $q \in \mathcal{Q}$:

$$\|\nabla \hat{J}_q(\hat{u})\|^2 \geq \rho(W(\hat{u}, q))$$

for all $\hat{u} \in \mathcal{L}_q$. \square

The previous properties essentially guarantee that in each region \mathcal{L}_q , a gradient-based feedback law of the form $\dot{\hat{u}} = -\nabla \hat{J}_q(\hat{u})$ can steer the vehicle toward points where the virtual potential \hat{J}_q has smaller values, without leaving the region \mathcal{L}_q . Since it is not possible to satisfy property (c) in the entire space using only one barrier function, we ask that the condition holds only in each of the regions \mathcal{L}_q via a different potential \hat{J}_q .

By using the measurements of \hat{J}_q at the points \hat{u} , we can now introduce the flow and jump sets for the hybrid set-seeking system (Figure 21), with state (\hat{u}, q)

$$C_{u,q} := \{(\hat{u}, q) \in \overline{\mathcal{U}} \times \mathcal{Q} : \hat{J}_q(\hat{u}) \leq \chi \hat{J}_{3-q}(\hat{u})\} \quad (69a)$$

$$D_{u,q} := \{(\hat{u}, q) \in \overline{\mathcal{U}} \times \mathcal{Q} : \hat{J}_q(\hat{u}) \geq (\chi - \lambda) \hat{J}_{3-q}(\hat{u})\} \quad (69b)$$

where $\chi \in (1, \infty)$ and $\lambda \in (0, \chi - 1)$ are tunable parameters. The set $C_{u,q}$ describes the points in the space where the controller uses the potential field \hat{J}_q , while the set $D_{u,q}$ indicates the “switching zones” for the controller, where the controller toggles the logic state using $q^+ = 3 - q$. By construction, this switching behavior will take place whenever the value of the current function \hat{J}_q exceeds a threshold compared to the

²A function $\tilde{w} : \mathcal{U} \rightarrow \mathbb{R}_{>0}$ is a proper indicator on the open set \mathcal{U} if it is continuous and $\tilde{w}(x_i) \rightarrow \infty$ when $i \rightarrow \infty$ if either $|x_i| \rightarrow \infty$ or the sequence $\{x_i\}_{i=1}^\infty$ approaches the boundary of \mathcal{U} .




FIGURE 21 (a) and (b) The trajectories of a vehicle in \mathbb{R}^2 , controlled via a hybrid source-seeking algorithm for obstacle avoidance. The trajectories are shown over the virtual partitions \mathcal{L}_1 and \mathcal{L}_2 . The system is initialized in \mathcal{L}_1 with $q = 1$ but eventually switches to mode $q = 2$ to evolve in \mathcal{L}_2 until it converges to a neighborhood of the source.

other function \hat{J}_{3-q} . In particular, the construction of the flow and jump sets imposes a *hysteresis* property on the switching controller, which precludes Zeno behavior. The red and blue lines in Figure 21 illustrate the boundaries of $D_{u,q}$ and $C_{u,q}$, respectively. The following proposition shows that the resulting nominal target hybrid controller has the form (33) and satisfies all the properties of Assumption 1 relative to $\bar{\mathcal{U}} \times \mathcal{Q}$.

Proposition 5 (Source Seeking and Obstacle Avoidance)

Consider the hybrid set-seeking system (35), with sets $C_{u,q}$ and $D_{u,q}$ given by (69), function $\hat{F}_\delta = \xi$, and input $y = \hat{J}_q$. Suppose that Assumption 4 holds. Then, system (35) satisfies condition (a) in Assumption 1 with $r = 1$, $\delta = 0$, $\hat{z} = q$, $\Psi = \mathcal{Q}$, and $\mathcal{O} = \{u^*\}$. Moreover, item (b) holds with the basin of attraction $\mathcal{U} \times \mathcal{Q}$. \square

To further illustrate the stability properties of the resulting hybrid source-seeking system for obstacle avoidance, we show in Figure 21 the constructions of the partitions \mathcal{L}_1 and \mathcal{L}_2 as well as the trajectory of the vehicle (65) under the control law (66), evolving over the two virtual partitions. As observed, the hybrid controller can (practically) stabilize the source of the signal J by switching the potential field after approximately 30 s of flow while avoiding the obstacle. We note that this simple example can be extended to other vehicle dynamics, including nonholonomic vehicles [83], and to the multiobstacle case by considering other types of target hybrid systems [83], [138], [139]. In such cases, the target hybrid decision-making dynamics (33) will be modified, but the structure and stability properties of the hybrid set-seeking dynamics (35) can be studied using similar tools.

Seeking With Momentum and Resets

Reset control [140], [141], [142] is a well-known technique used to improve transient performance in dynamical systems that incorporate integrators in the feedback loop. The main idea behind the incorporation of resets is to dissipate energy at the appropriate times to decrease overshoots and improve transient performance [142]. Resets can also be used to model restarting methods [113], [114], commonly employed in optimization routines. By leveraging tools from hybrid systems, set-seeking algorithms with momentum and restarting can be designed to reduce overshoots and improve transient performance [69], [90], [116]. To illustrate this idea, we can consider the following target hybrid decision-making dynamics with state $(\hat{u}, q, \tau) \in \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}_{\geq 0}$ and data

$$C_{u,z} = \mathbb{R}^n \times \mathbb{R}^n \times [T_0, T] \quad (70a)$$

$$\hat{F}(\hat{u}, q, \tau) = \begin{pmatrix} \frac{2}{\tau}(q - \hat{u}) - k_1 \nabla J(\hat{u}) \\ -2k_2 \tau \nabla J(\hat{u}) \\ \rho \end{pmatrix} \quad (70b)$$

$$D_{u,z} = \mathbb{R}^n \times \mathbb{R}^n \times \{T\} \quad (70c)$$

$$\hat{G}(\hat{u}, q, \tau) = \begin{pmatrix} \hat{u} \\ \alpha \hat{u} + (1 - \alpha)q \\ T_0 \end{pmatrix} \quad (70d)$$

where $k_1 \geq 0$ and $T_0, T, k_2, \rho > 0$ are tunable parameters that satisfy $T > T_0$. The parameter $\alpha \in \{0, 1\}$ is selected a priori to indicate whether the state q is reset during jumps. When the timer τ satisfies $\tau = T$, the discrete-time dynamics (70d) reset the timer back to T_0 . If, additionally, the parameter α satisfies $\alpha = 1$, then the state q is also reset to the current value of \hat{u} . Figure 22 is a block diagram representation of the hybrid set-seeking dynamics (35) associated with the target system (70) for the case $k_1 = 0$. The low-pass filter is omitted for simplicity.

Using the change of coordinates $s = \hat{u}$ and $q = s + (1/2)\tau(\dot{s} + k_1 \nabla J)$, the flows induced by (70b) are related to the second-order ODE

$$\ddot{s} + \frac{(2 + \dot{\tau})}{\tau} \dot{s} + 4k_2 \nabla J(s) + k_1 \left(\nabla^2 J(s)^\top \dot{s} + \frac{\dot{\tau}}{\tau} \nabla J(s) \right) = 0$$

which has been studied in the context of accelerated optimization algorithms under different choices of $\dot{\tau}$ and parameters k_1, k_2 [93], [114], [143], [144], [145]. For example, when $k_1 = 0$, the previous dynamics can also be written in state-space representation using $s_1 = s$ and $s_2 = \dot{s}$, leading to

$$\dot{s}_1 = s_2, \quad \dot{s}_2 = -\frac{c_1}{\tau} s_2 - c_2 \nabla J(s_1), \quad \dot{\tau} = \rho \quad (71)$$

for suitable constants $c_1, c_2 > 0$. Note that when $\nabla J = \xi$, the system (71) has a structure similar to (40a) when $\hat{F}_\delta = \xi$. To avoid the issues discussed in Example 9, we can consider the flow and jump sets described in (70c) and (70a), respectively, and the jump map

$$s_1^+ = s_1, \quad s_2^+ = (1 - \alpha)s_2, \quad \tau^+ = T_0 \quad (72)$$

which has a similar structure as (40b), with $s_1 = \hat{u}$, $\hat{z} = \tau$, $s_2 = \xi$, $\hat{G}_\delta = (s_1, T_0)$, and $G_\xi = (1 - \alpha)s_2$. Therefore, this system can be viewed as a target hybrid decision-making system equipped with a hybrid filter that incorporates jumps to periodically reset the momentum to zero. Under a suitable choice of parameters, this hybrid set-seeking




FIGURE 22 A block diagram representation of the hybrid set-seeking system with momentum, with (70) as the target hybrid decision-making dynamics with $k_1 = 0$. Here, the low-pass filter is omitted.

as in [72] so that it has entries satisfying $N_{11} = (16/\varepsilon_a^2)(\mu_1^2 - (1/2))$, $N_{22} = (16/\varepsilon_a^2)(\mu_3^2 - (1/2))$, $N_{12} = (4/\varepsilon_a^2)\mu_1\mu_3$, and $N_{12} = N_{21}$. To define the sets C and D , we follow a uniting control approach [9], where the logic state \hat{z} indicates the current algorithm used by the controller. In particular, when the value of the cost J is sufficiently close to the optimal value J^* , we use $\hat{z} = 0$. On the other hand, when the value of J is greater than J^* , we use $\hat{z} = 1$. While this approach might require knowledge of J^* , for many ES applications where the cost to be minimized corresponds to some error function, it is known that $J^* = 0$. On the other hand, if J^* is unknown, a rough estimate can be used for the initial construction of the controller.

Building on these ideas, the switching zones for \hat{z} are defined using tunable thresholds specified by constants c_0 and c_{10} , with $c_0 > c_{10} > 0$. To characterize these thresholds, we define the sets

$$C_0 := \{\hat{u} \in \mathbb{R}^2 : J(\hat{u}) - J^* \leq c_0\} \quad (75a)$$

$$C_1 := \overline{\mathbb{R}^2 \setminus \{\hat{u} \in \mathbb{R}^2 : J(\hat{u}) - J^* < c_{10}\}} \quad (75b)$$

as well as the sets

$$D_0 := \overline{\mathbb{R}^2 \setminus \{\hat{u} \in \mathbb{R}^2 : J(\hat{u}) - J^* < c_0\}} \quad (76a)$$

$$D_1 := \{\hat{u} \in \mathbb{R}^2 : J(\hat{u}) - J^* \leq c_{10}\}. \quad (76b)$$

Using the constructions (75) and (76), the flow and jump sets for the state $x_{u,z}$ are given by

$$C_{u,z} := (C_0 \times \{0\}) \cup (C_1 \times \{1\}) \quad (77a)$$

$$D_{u,z} := (D_0 \times \{0\}) \cup (D_1 \times \{1\}) \quad (77b)$$

and the discrete-time dynamics of $x_{u,z}$ are given by

$$\hat{G}(x_{u,z}) = \begin{pmatrix} \hat{u} \\ 1 - \hat{z} \end{pmatrix}. \quad (78)$$

This scheme can be seen as an ES with a “hybrid supervisor” that decides when to switch, thereby precluding Zeno behavior owing to the hysteresis mechanism induced by the design of the flow and jump sets. A diagram of the switching zones in \mathbb{R}^2 is shown in Figure 25.

Proposition 7 (Switched Newton-Gradient Dynamics)

Consider the hybrid set-seeking dynamics (35), with $y = J(u)$, flow map (73), data (74)–(78), and feedback law (34). Then, condition (a) of Assumption 1 holds with $r = 1$ and $\delta = 0$. Moreover, if J is smooth, and strongly convex, then item (b) also holds for the respective target hybrid decision-making dynamics with $\xi_1 = \nabla J$, $\xi_2 = \nabla^2 J^{-1} \nabla J$, $\mathcal{O} = \{u^*\}$ and $\Psi = \{0\}$. \square

We illustrate the Newton-gradient switched seeking dynamics using a simple 2D numerical example for a quadratic cost function J , with the results shown in Figure 26. The figure depicts in green a diverging trajectory that occurs when Newton-like ES dynamics are applied far from the optimal point using a moderate frequency ω . As shown by the blue trajectory, incorporating the hybrid uniting mechanism introduced in this section not only eliminates the instability but also significantly improves the transient performance of the controller—outperforming the more robust yet slower gradient-based ES dynamics, represented by the red trajectory. Figure 27 highlights this improvement in transient behavior, reducing the convergence time by more than half.

PART 3: HYBRID SET-SEEKING SYSTEMS FOR DYNAMIC PLANTS

In the previous section, we presented results and illustrative examples of hybrid set-seeking systems applied to plants modeled as static maps with output $y = J(u)$. In Part 3, we extend these results to settings in which the plant exhibits dynamics. By leveraging singular perturbation results for hybrid inclusions (see “Singularly Perturbed HDSs”), the standard methodology used to incorporate dynamic plants into smooth ES controllers can be naturally extended to the hybrid systems framework. This extension enables real-time set seeking in dynamical systems.

Model of the Plant

Consider a plant with state $\theta \in \mathbb{R}^m$, input $u \in \mathbb{R}^n$, and output $y \in \mathbb{R}$, modeled by the constrained differential inclusion

$$\dot{\theta} \in P(\theta, u), \quad (\theta, u) \in \Lambda_\theta \times \mathbb{U}, \quad y = \varphi(\theta, u) \quad (79)$$




FIGURE 25 (a)–(c) The representation of sets defined in (75) and (76) as well as the region of the space where hysteresis is induced in the switched controller.

where $P: \mathbb{R}^m \times \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ is a set-valued mapping, $\varphi: \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}$ is an output function, θ evolves in the compact set $\Lambda_\theta := \lambda_\theta \mathbb{B} \subset \mathbb{R}^m$, $\lambda_\theta \in \mathbb{R}_{>0}$ can be taken arbitrarily large to encompass any solution of interest, and u evolves in the closed set $\mathbb{U} := \hat{\mathbb{U}} + \mathbb{B}$, where $\hat{\mathbb{U}} \subset \mathbb{R}^n$. The unitary inflation on $\hat{\mathbb{U}}$ is motivated by the fact that the control signal u will satisfy $u(t, j) \in \hat{\mathbb{U}} + \varepsilon_a \mathbb{B} \subset \mathbb{U}$ for all (t, j) in the domain of the solutions, where $\varepsilon_a \in (0, 1)$ is a tunable parameter.

To have enough regularity in the closed-loop system, we assume that $\hat{\mathbb{U}}$ is a closed set, that $P(\cdot, \cdot)$ is OSC, LB, and convex-valued relative to $\mathbb{R}^m \times \mathbb{U}$, and that $\varphi(\cdot, \cdot)$ is continuously differentiable. In this sense, the model (79) is quite general as it captures differential equations with a continuous right-hand side as well as discontinuous plants modeled by their corresponding Krasovskii regularizations (see “Krasovskii Solutions of ODEs” in “Continuous-Time Set-Valued Dynamical Systems”). Common examples of plants with a discontinuous right-hand side include mechanical systems with Coulomb friction [154, Ch. 12], systems arbitrarily switching between a finite number of continuous vector fields [8, Example 2.14], and plants with uncertain models and internal discontinuous feedback controllers, such as those based on sliding mode control [57].

Model-Free Feedback Optimization

For dynamic plants, the goal in ES is to regulate the input u toward the set of points that optimizes the steady-state input-to-output map of the plant. To ensure a well-defined optimization problem, we assume that the plant (79) exhibits suitable stability properties for each fixed input u . This can be achieved by first designing an inner controller that stabilizes (79) with respect to an external reference. Specifically, we assume that there exists a set-valued mapping $H: \mathbb{R}^n \rightrightarrows \mathbb{R}^m$ that is OSC and LB relative to \mathbb{U} , such that $H(\mathbb{U}) \subset \Lambda_\theta$, and for each $\rho > 0$, the compact set

$$\mathbb{M}_\rho := \{(\theta, u) : \theta \in H(u), u \in \mathbb{U} \cap \rho \mathbb{B}\} \quad (80)$$




FIGURE 26 The trajectory \hat{u} of the hybrid set-seeking system (shown in blue) switching between Newton-based and gradient-based ES compared to the trajectories obtained via standard gradient-based ES (shown in red) and standard Newton-based ES (shown in green).

is UGAS for the dynamics

$$(\theta, u) \in \Lambda_\theta \times (\mathbb{U} \cap \rho \mathbb{B}), \quad \dot{\theta} \in P(\theta, u), \quad \dot{u} = 0. \quad (81)$$

In other words, the previous property establishes that for each fixed input u , the plant dynamics (79) have a well-defined (potentially set-valued) quasi-steady state, characterized by the set-valued mapping H . This property generalizes the classic assumptions made for ES in differential equations, for example, [66, Assumptions 1 and 2], for the case that the plant under control is given by a constrained set-valued dynamical system.

Example 14 (Oscillator with Coulomb Friction)

Consider a simple harmonic oscillator with Coulomb friction and a controllable velocity offset, given in open loop by the discontinuous dynamics

$$\dot{\theta}_1 = \theta_2 - u, \quad \dot{\theta}_2 = -\frac{B}{M} \operatorname{sgn}(\theta_2 - u) - \frac{K}{M} \theta_1, \quad \dot{u} = 0 \quad (82)$$

with $\Lambda_\theta = \lambda_\theta \mathbb{B} \subset \mathbb{R}^2$, $\mathbb{U} = \mathbb{R}$, $(B, K, M) \in \mathbb{R}_{>0}^3$, and $\lambda_\theta > (B/K) > 0$ is selected sufficiently large to characterize all the complete solutions of interest. The Krasovskii regularization of this system affects only the dynamics of θ_2 and is given by

$$\dot{\theta}_2 \in P_2(\theta, u) = \begin{cases} -\frac{B}{M} - \frac{K}{M} \theta_1 & \text{if } \theta_2 > u \\ \left[-\frac{B}{M}, \frac{B}{M}\right] - \frac{K}{M} \theta_1 & \text{if } \theta_2 = u \\ \frac{B}{M} - \frac{K}{M} \theta_1 & \text{if } \theta_2 < u. \end{cases} \quad (83)$$

In this case, for each $\rho > 0$, system (81) renders the set $\mathbb{M}_\rho := \{(\theta, u) : \theta \in [-(B/K), (B/K)] \times \{u\}, u \in \mathbb{R} \cap \rho \mathbb{B}\}$ UGAS [155, Section 3], thus satisfying the assumptions made on (79). \square

Example 15 (Fast Switching Plant)

Consider an open-loop switched linear system given by $\dot{\theta} = p_q(\theta, u) := A_q \theta + B_q u$, $\dot{u} = 0$, $\mathbb{U} := \mathbb{R}$, $(\epsilon, \lambda_\theta) \in \mathbb{R}_{>0}^2$, with matrices




FIGURE 27 The evolution in time of trajectories \hat{u} for the hybrid set-seeking system switching between Newton-based and gradient-based ES (shown in blue). The red trajectory corresponds to the standard gradient-based ES.

$$A_q = \begin{bmatrix} -1 & \frac{3}{2} - \frac{5}{4}q \\ -\frac{9}{4} + \frac{5}{4}q & -1 \end{bmatrix}, B_q = \begin{bmatrix} 1 \\ \frac{9}{4} - \frac{5}{4}q \end{bmatrix} \quad (84)$$

where $q \in \{1, 2\}$. Under the arbitrarily fast switching of q , this system is conveniently modeled by the differential inclusion

$$\dot{\theta} \in \{\alpha p_1(\theta, u) + (1 - \alpha)p_2(\theta, u), \alpha \in [0, 1]\}, \dot{u} = 0, \quad (85)$$

which is OSC, LB, and convex-valued. For each $\rho > 0$, system (85) with the flow set $\Lambda_\theta \times (\mathbb{U} \cap \rho\mathbb{B})$ renders UGAS the set $\mathbb{M}_\rho := \{(\theta, u) : \theta \in \{u\} \times \{0\}, u \in \mathbb{R} \cap \rho\mathbb{B}\}$, thus satisfying the required assumption. \square

To ensure a well-defined single-valued steady-state model-free optimization problem, we also assume that for each $u \in \mathbb{U}$ and any pair $\theta, \theta' \in H(u)$, the condition $\varphi(\theta, u) = \varphi(\theta', u)$ holds. In Example 14, any continuous output function $\varphi(\cdot, \cdot)$ that depends only on θ_2 in its first argument will satisfy this assumption. On the other hand, in Example 15, any continuous output function satisfies this assumption. If H is single-valued, as in the traditional literature of ES, then this assumption is also satisfied.

We can now define the *response map* of the plant (79)

$$J(u) := \{\varphi(\theta, u) : \theta \in H(u)\} \quad (86)$$

which is assumed to be continuously differentiable. As in the static case, the set-seeking problem is given by

$$\text{optimize } J(u), \text{ subject to } u \in \hat{\mathbb{U}}. \quad (87)$$

For consistency, we denote again the set of solutions to (87) as $\mathcal{O} \subset \mathbb{R}^n$, which is assumed to be nonempty and compact.

Closed-Loop System and Stability Properties

To interconnect the plant (79) with the hybrid set-seeking dynamics (35) via the feedback law (34), a timescale separation usually needs to be induced between the flows of both systems. To achieve this separation, the flow map (35b) is multiplied by a small gain $k > 0$, leading to the following closed-loop hybrid system:

$$C := C_{u,z} \times \Lambda_\xi \times \mathbb{S}^n \times \Lambda_\theta \quad (88a)$$

$$\dot{x} \in F(x) := \begin{pmatrix} k\hat{F}_\delta(x_{u,z}, \xi) \\ -k\frac{k_f}{\varepsilon_f} \left(\xi - \frac{2}{\varepsilon_a} \varphi(\theta, \hat{u} + \varepsilon_a \mathbb{D}\mu) \cdot \mathbb{D}\mu \right) \\ k\Phi(\omega)\mu \\ P(\theta, \hat{u} + \varepsilon_a \mathbb{D}\mu) \end{pmatrix} \quad (88b)$$

$$D := D_{u,z} \times \Lambda_\xi \times \mathbb{S}^n \times \Lambda_\theta \quad (88c)$$

$$x^+ \in G(x) := \begin{pmatrix} \hat{G}_\delta(x_{u,z}) \\ \xi \\ \mu \\ \theta \end{pmatrix} \quad (88d)$$

where $x := (x_{u,z}, \xi, \mu, \theta) \in \mathbb{R}^{r+4n+m}$, with tunable parameters $(k, \varepsilon_a, \varepsilon_f, \varepsilon_\omega)$. Figure 28 shows a high-level block diagram of the interconnection between the plant and the controller. The dependence of the hybrid dynamics on the output y highlights that the specific structure of the components is application dependent.

To study the stability properties of the closed-loop system (88) via singular perturbation theory, we introduce a new continuous-time variable given by $s = kt$. Since $(d/dt)x = k(d/ds)x$, the continuous-time dynamics (88b) in the new s timescale become

$$\frac{d}{ds}x \in F(x) := \begin{pmatrix} \hat{F}_\delta(x_{u,z}, \xi) \\ -\frac{k_f}{\varepsilon_f} \left(\xi - \frac{2}{\varepsilon_a} \varphi(\theta, \hat{u} + \varepsilon_a \mathbb{D}\mu) \cdot \mathbb{D}\mu \right) \\ \Phi(\omega)\mu \\ \frac{1}{k}P(\theta, \hat{u} + \varepsilon_a \mathbb{D}\mu) \end{pmatrix} \quad (89)$$

which, combined with (88a), (88c), and (88d), generates a singularly perturbed hybrid system of the form (S19) with $\varepsilon = k$ (see “Singularly Perturbed HDSs”). We denote this system as \mathcal{H}_s to emphasize that the continuous-time dynamics evolve in the s timescale. For this system, it follows that the reduced hybrid dynamics (S22) are precisely given by the hybrid set-seeking systems (35) with $y = J(u)$, whose stability properties were studied in Theorems 3 and 4. Since the boundary layer dynamics of \mathcal{H}_s are precisely given by the plant dynamics (81), we can use Theorem 6 (see “Singularly Perturbed HDSs”) to obtain the following stability result for hybrid set-seeking systems with plants in the loop [91]:

Theorem 7 (Stability Properties in Dynamic Maps)

Consider the HDS \mathcal{H}_s , and suppose that Assumption 1 holds for the hybrid reduced set-seeking dynamics (35). Then, the set $\mathcal{A} \times \Lambda_\xi \times \mathbb{S}^n \times \Lambda_\theta$ is SGPAS as $(\delta, \varepsilon_f, \varepsilon_a, \varepsilon_\omega, k) \rightarrow 0^+$.

Based on Theorem 7 and the material presented in Parts 2 and 3 of this article, several comments are in order.

- » The semiglobal practical stability properties of Theorem 7 are established with respect to compact sets, where the main state $x_{u,z}$ is steered toward a neighborhood of $\mathcal{A} = \mathcal{O} \times \Psi$; that is, \hat{u} converges to a neighborhood of the optimal solution to the problem (87). Since the actual input u is defined via the feedback law (34), it will also converge to a neighborhood of order ε_a of the set \mathcal{O} , which is consistent with the classic results in ES.
- » It is possible to refine the convergence statement for the state θ , such that $\theta(t, j) \rightarrow H(\mathcal{O})$ as $(t, j) \rightarrow \infty$. Without additional structure or regularity conditions, this property holds only if jumps do not dominate the hybrid system’s dynamics—that is, the system must not exhibit Zeno behavior or purely discrete solutions. Such pathological behaviors can be ruled out by applying any of the temporal regularization strategies

The analytical development of hybrid set-seeking systems builds on the traditional tools used for ES control: averaging and singular perturbation theory.

discussed in Part 2 using dynamic timers, which impose dwell-time or average dwell-time constraints on the jump times.

- » Similarly to the static map case, the result of Theorem 7 asserts suitable \mathcal{KL} bounds of the form (28) for all solutions of the system. In this sense, we do not insist on the uniqueness of solutions for any of the proposed controllers. Instead, we ask that every solution satisfies the desired property. When multiple solutions exist from a given initial condition, the practitioner can select the most meaningful solution for the application of interest based on application-dependent considerations.
- » As discussed in Part 2, the result of Theorem 7 can be extended to cover set-seeking systems that use hybrid filters and hybrid dither generators as part of the feedback loops. Similarly, more complex plants can be incorporated by using averaging tools for hybrid systems that exhibit hybrid boundary layer dynamics. Due to space constraints, we do not elaborate on these extensions here, and instead, we refer readers interested in these subjects to references [43], [98], [100], [102], and [108].

We finish this section with a numerical example in the context of Nash equilibrium-seeking in noncooperative games, and by highlighting the fact that the stability results presented in Theorems 3, 4, and 7 are amenable to standard discretization techniques, such as (consistent) Euler methods or Runge-Kutta methods (c.f., Corollary 2). This followed by the fact that the proposed hybrid controllers are robust to small perturbations, including those that emerge when the flow map is discretized. For further details on this subject, we refer the reader to [69, Sec. 4] and [88].

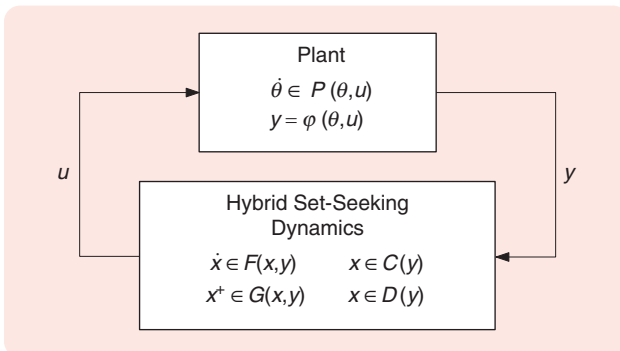


FIGURE 28 A dynamic plant interconnected with hybrid set-seeking dynamics for model-free feedback optimization.

Example 16 (Nash Seeking With Intermittent Updates)

Consider a two-player noncooperative game, where each player $i \in \{1, 2\}$ is characterized by a switching plant with state $\theta_i = (\theta_{i,1}, \theta_{i,2})$, and dynamics of the form (84), that is

$$\dot{\theta}_{i,1} = -\theta_{i,1} + \left(\frac{3}{2} - \frac{5}{4}p_i\right)\theta_{i,2} + u_i \quad (90)$$

$$\dot{\theta}_{i,2} = -\theta_{i,2} + \left(-\frac{9}{4} + \frac{5}{4}p_i\right)\theta_{i,1} + \left(\frac{9}{4} - \frac{5}{4}p_i\right)u_i \quad (91)$$

where $p_i: \text{dom}(q) \rightarrow \{1, 2\}$ is the switching signal, which, for simplicity, is assumed to be the same for both players. Each player has a quadratic output function φ_i that depends on the first state of each player, that is, $y_i = \varphi_i(\theta_1) = \bar{\theta}_1^T Q_i \bar{\theta}_1 + \bar{\theta}_1^T b_i + c_i$, where $\bar{\theta}_1 = (\theta_{1,1}, \theta_{2,1})$, $Q_i \in \mathbb{R}^{2 \times 2}$, $b_i \in \mathbb{R}^2$ and $c_i \in \mathbb{R}$ are selected as in [91, Sec.6.3]. The goal of the players is to individually maximize their own steady-state input-to-output cost function (86) using only individual output measurements y_i . However, unlike traditional Nash-seeking problems where players continuously update their actions [51], [107], we study the scenario where the inputs to the player dynamics are updated only sporadically, which is common under computational or communication constraints. To model this behavior, we assign to each player a logic state $q_i \in \{0, 1\}$ that multiplies the dynamics of the player's states \hat{u}_i , $i \in \{1, 2\}$. In particular, we consider the simple Nash-seeking rule $\dot{\hat{u}}_i = q_i \xi_i$, where ξ_i is generated by an individual low-pass filter implemented by each player.

Figure 29 shows a block diagram representation of the closed-loop system of each player $i \in \{1, 2\}$ under the seeking dynamics with switching logic states q_i and switching plant with logic states p_i . To study the stability properties of the overall switching system, we let $\ell \in \{1, 2, 3, 4\}$ be a logic state that indexes the four possible combinations of the pair (q_1, q_2) , where $\ell = 1$ corresponds to $q_1 = q_2 = 1$. Using $\ell \in Q = \{1, 2, 3, 4\}$ as a logic state and $Q = Q_s \cup Q_u$, where $Q_s = \{1\}$ and $Q_u = \{2, 3, 4\}$, we impose an average dwell-time bound of the form (58) on the jump times of ℓ^+ using the hybrid automaton (57) and an average activation bound (60) using the hybrid monitor (61). Additionally, since the switching plant admits a common Lyapunov function, we leverage Proposition 1 and consider the differential inclusion (53) obtained from the two modes of the plant. The resulting system has the form (88), with main states $x_{u,z} = (\hat{u}, \ell, \tau_1, \tau_2)$. Figure 30 shows the trajectories of the state variables θ_1 for the players under the Nash-seeking dynamics with intermittent updates. The red-shaded regions indicate time intervals during which player 2 does not update its action, while the

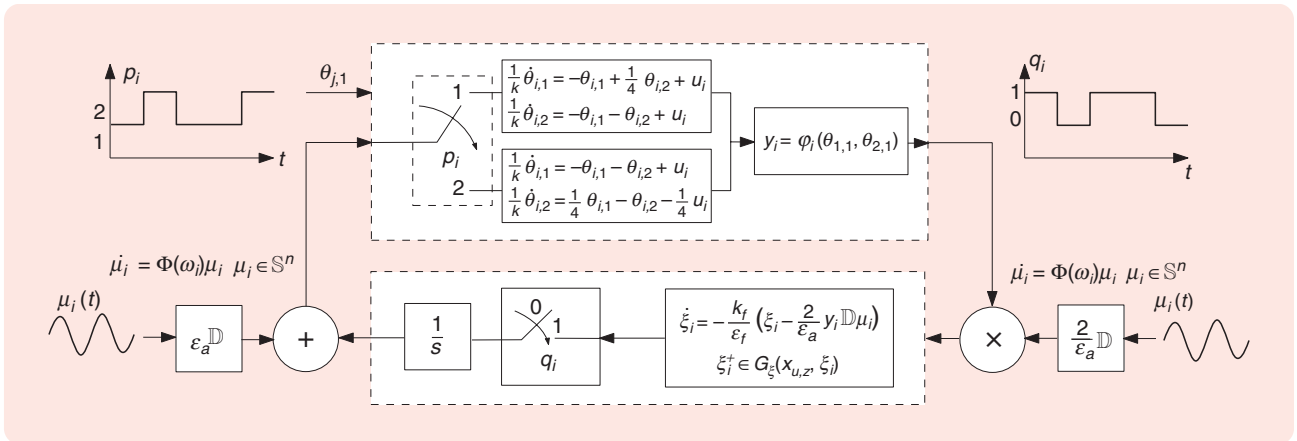


FIGURE 29 A block diagram representation of Nash equilibrium-seeking dynamics under intermittent updates for players characterized by switching dynamical systems. The logic mode q_i indicates if the dynamics \dot{u}_i are inactive ($\dot{u}_i = 0$) or active ($\dot{u}_i = \xi$).

blue-shaded region corresponds to periods when player 1 does not update \hat{u}_1 . As observed, for a moderate value of η_2 , the trajectories still converge to a neighborhood of the NE, indicated by the dotted lines. However, the convergence time increases as $\eta_2 \rightarrow 1$. For related results in scenarios where q_i is allowed to take negative values—causing the nominal Nash-seeking dynamics to become unstable for bounded periods of time (that is, in the presence of adversarial players)—we refer the reader to [91, Sec.6.3].

CONCLUSION

Hybrid set-seeking systems extend traditional continuous-time ES algorithms by enabling feedback schemes, algorithms, and plants that incorporate both continuous-time and discrete-time dynamics and whose stability properties are studied with respect to general sets as opposed to equilibrium points. These systems naturally arise in applications that require switching between multiple controllers to meet stringent transient or steady-state specifications that cannot be achieved with smooth feedback alone. They are also relevant for “intelligent” controllers that use “if-then” rules. Hybrid set-seeking systems may also emerge when the seeking dynamics are smooth but the plant under control exhibits hybrid behavior, such as jumps or impacts. Other potential applications include developing general real-time model-free optimization algorithms for CPSs, where digital components are integrated into physical plants, such as in sampled-data systems and event-triggered control. The analytical development of hybrid set-seeking systems builds on the traditional tools used for ES control: averaging and singular perturbation theory. By extending these tools to hybrid systems, new families of model-free hybrid set-seeking systems can be constructed to emulate the model-based hybrid decision-making algorithm. In this sense, working with hybrid time domains and dynamics that satisfy the appropriate regularity conditions allows for the property of “closeness of solutions,” which is crucial in the study of smooth ES systems, to be established for hybrid set-seeking systems as




FIGURE 30 The evolution in time of the states $\theta_{i,1}$, for each player $i \in \{1, 2\}$ under intermittent updates of the Nash equilibrium-seeking dynamics. The colors (light blue and light red) of the shaded regions indicate times when $q_i = 0$ for each player i .

well. By leveraging both this property and the uniform stability of the target dynamics, semiglobal stability can be demonstrated for general hybrid set-seeking dynamics.

The approaches studied in this article rely on conventional averaging techniques, often referred to as “first-order” averaging. However, over the past decade, a rich body of work in ES has emerged based on Lie-bracket and second-order averaging techniques; see [48], [95], [101], [133], [144], [151], [152], and [153]. These methods are often better suited for applications involving geometric constraints, such as mechanical systems and other applications defined on manifolds. For a monograph on Lie-bracket-based ES with applications to model-free stabilization, we refer the reader to [146]. For hybrid tools similar to those discussed in this article, which are applicable to Lie-bracket averaging-based systems, we refer readers to [103] and [156].

The examples studied in this article only scratch the surface of the broader potential of hybrid ES and hybrid set-seeking for model-free control and optimization. These approaches open the door to novel controllers and algorithms capable of incorporating “if-then” logic through hybrid systems tools. Furthermore, such algorithms are expected to benefit from

techniques developed in related fields, including computer science, signal processing, and communications.

Finally, in addition to the development of these novel algorithms, several open problems remain in the context of hybrid seeking systems. These include the development of a comprehensive stability framework for hybrid seeking systems that incorporate stochastic behaviors, as well as developing tools for interconnected networked seeking systems, that is, “open ES systems”—an underexplored area with numerous potential applications in complex engineering, biological, and sociotechnical systems where multiple individual agents implement different adaptation and optimization algorithms for real-time decision making.

ACKNOWLEDGMENT

Many of the developments reported in this article are a result of research supported by the Air Force Office of Scientific Research under Grant FA9550-22-1-0211 and the National Science Foundation under Grants CNS 1947613, ECCS CAREER 2305756, and CMMI 2228791. The opening image was generated using OpenAI-ChatGPT, GPT-5.3, and edited by the authors.

AUTHOR INFORMATION

Jorge I. Poveda (poveda@ucsd.edu) received the B.S. degrees in electronics engineering and mechanical engineering, both from the University of Los Andes, Bogotá, Colombia, in 2012. He received the M.Sc. and Ph.D. degrees in electrical and computer engineering from the University of California, Santa Barbara, in 2016 and 2018, respectively. Subsequently, he was a postdoctoral fellow at Harvard University in 2018 and an assistant professor at the University of Colorado (Boulder). Since 2022, he has been with the Department of Electrical and Computer Engineering at the University of California, San Diego, San Diego, CA 92093 USA, where he is currently an associate professor. He is the recipient of the CRII and CAREER Awards from NSF, the Young Investigator Awards from AFOSR and SHPE, the Donald P. Eckman Award, and the CCDC Outstanding Scholar Fellowship and Best Ph.D. Dissertation awards from UCSB. He has also received several Best Paper and Best Paper Finalist awards with his students and colleagues, including recognitions from IEEE Transactions on Control of Network Systems (2023), the American Control Conference (2024, 2026), the IEEE Conference on Decision and Control (2017, 2021), and the IFAC Conference on Analysis and Design of Hybrid Systems (2024). His research interests are in hybrid control, adaptive systems, and multiagent decision-making algorithms.

Andrew R. Teel received the A.B. degree in engineering sciences from Dartmouth College in Hanover, NH, USA, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley, in 1989

and 1992, respectively. After receiving his Ph.D. degree, he was a postdoctoral fellow at the Ecole des Mines de Paris in Fontainebleau, France. In 1992, he joined the faculty of the Electrical Engineering Department at the University of Minnesota, where he was an assistant professor until 1997. Subsequently, he joined the faculty of the Electrical and Computer Engineering Department at the University of California, Santa Barbara, Santa Barbara, CA 93106-9560 USA, where he is currently a Distinguished Professor and the director of the Center for Control, Dynamical Systems, and Computation. His research interests are in nonlinear and hybrid dynamical systems, with a focus on stability analysis and control design. He has received NSF Research Initiation and CAREER Awards, the 1998 IEEE Leon K. Kirchmayer Prize Paper Award, the 1998 George S. Axelby Outstanding Paper Award, and was the recipient of the first SIAM Control and Systems Theory Prize in 1998. He was the recipient of the 1999 Donald P. Eckman Award and the 2001 O. Hugo Schuck Best Paper Award, both given by the American Automatic Control Council, and also received the 2010 IEEE Control Systems Magazine Outstanding Paper Award. In 2016, he received the Certificate of Excellent Achievements from the IFAC Technical Committee on Nonlinear Control Systems. In 2020, he and his coauthors received the HSCC Test-of-Time Award. He served as the editor-in-chief of *Automatica* from 2017 to 2025. He is a Fellow of the IEEE and of IFAC.

REFERENCES

- [1] A. M. Farid, “A hybrid dynamic system model for multimodal transportation electrification,” *IEEE Trans. Control Syst. Technol.*, vol. 25, no. 3, pp. 940–951, May 2017, doi: [10.1109/TCST.2016.2579602](https://doi.org/10.1109/TCST.2016.2579602).
- [2] D. E. Ochoa, F. Galarza-Jimenez, F. Wilches-Bernal, D. A. Schoenwald, and J. I. Poveda, “Control systems for low-inertia power grids: A survey on virtual power plants,” *IEEE Access*, vol. 11, pp. 20,560–20,581, 2023, doi: [10.1109/ACCESS.2023.3249151](https://doi.org/10.1109/ACCESS.2023.3249151).
- [3] A. D. Ames and I. Poulakakis, *Hybrid Zero Dynamics Control of Legged Robots*. Oxford, U.K.: Butterworth-Heinemann, 2017.
- [4] M. Senesky, G. Eirea, and T. J. Koo, “Hybrid modelling and control of power electronics,” in *Proc. Int. Workshop Hybrid Syst.: Computation Control*, Berlin, Germany: Springer-Verlag, 2003, pp. 450–465.
- [5] D. Stewart, “Rigid-body dynamics with friction and impact,” *SIAM Rev.*, vol. 42, no. 1, pp. 3–39, 2000, doi: [10.1137/S0036144599360110](https://doi.org/10.1137/S0036144599360110).
- [6] B. Lennartson, M. Tittus, B. Egardt, and S. Pettersson, “Hybrid systems in process control,” *IEEE Control Syst. Mag.*, vol. 16, no. 5, pp. 45–56, Oct. 1996, doi: [10.1109/37.537208](https://doi.org/10.1109/37.537208).
- [7] D. L. Pepyne and C. G. Cassandras, “Optimal control of hybrid systems in manufacturing,” *Proc. IEEE*, vol. 88, no. 7, pp. 1108–1123, Jul. 2000, doi: [10.1109/5.871312](https://doi.org/10.1109/5.871312).
- [8] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton, NJ, USA: Princeton Univ. Press, 2012.
- [9] R. Sanfelice, *Hybrid Feedback Control*. Princeton, NJ, USA: Princeton Univ. Press, 2021.
- [10] J. Lygeros, C. Tomlin, S. Sastry, “Hybrid systems: Modeling, analysis and control,” *Electron. Res. Lab., Univ. of California, Berkeley, CA, USA, Tech. Rep. UCB/ERL M*, 2008.
- [11] T. A. Henzinger, “The theory of hybrid automata,” in *Proc. 11th Annu. IEEE Symp. Logic Comput. Sci.*, Piscataway, NJ, USA: IEEE Press, pp. 278–292, 1996, doi: [10.1109/LICS.1996.561342](https://doi.org/10.1109/LICS.1996.561342).
- [12] J. Lygeros, K. H. Johansson, S. N. Simic, J. Zhang, and S. S. Sastry, “Dynamical properties of hybrid automata,” *IEEE Trans. Autom. Control*, vol. 48, no. 1, pp. 2–17, Jan. 2003, doi: [10.1109/TAC.2002.806650](https://doi.org/10.1109/TAC.2002.806650).

- [13] L. Tavernini, "Differential automata and their discrete simulators," *Nonlinear Anal. Theory Methods Appl.*, vol. 11, no. 6, pp. 665–683, 1987, doi: [10.1016/0362-546X\(87\)90034-4](https://doi.org/10.1016/0362-546X(87)90034-4).
- [14] A. R. Teel and D. Popović, "Solving smooth and nonsmooth multivariable extremum seeking problems by the methods of nonlinear programming," in *Proc. Amer. Control Conf.*, 2001, pp. 2394–2399, doi: [10.1109/ACC.2001.946111](https://doi.org/10.1109/ACC.2001.946111).
- [15] T. Pan, Z. Ji, and Z. Jiang, "Maximum power point tracking of wind energy conversion systems based on sliding mode extremum seeking control," in *Proc. IEEE Energy Conf.*, Atlanta, GA, USA, pp. 1–5, 2008, doi: [10.1109/ENERGY.2008.4781032](https://doi.org/10.1109/ENERGY.2008.4781032).
- [16] N. J. Killingsworth and M. Krstic, "PID tuning using extremum seeking: Online, model-free performance optimization," *IEEE Control Syst. Mag.*, vol. 26, no. 1, pp. 70–79, Feb. 2006, doi: [10.1109/MCS.2006.1580155](https://doi.org/10.1109/MCS.2006.1580155).
- [17] J. Cochran, E. Kanso, S. D. Kelly, H. Xiong, and M. Krstić, "Source seeking for two nonholonomic models of fish locomotion," *IEEE Trans. Robot.*, vol. 25, no. 5, pp. 1166–1176, Oct. 2009, doi: [10.1109/TRO.2009.2026500](https://doi.org/10.1109/TRO.2009.2026500).
- [18] H. Yu and U. Ozguner, "Extremum-seeking control strategy for ABS system with time delay," in *Proc. Am. Control Conf.*, 2002, pp. 3753–3758, doi: [10.1109/ACC.2002.1024511](https://doi.org/10.1109/ACC.2002.1024511).
- [19] M. Benosman, "Multi-parametric extremum seeking-based auto-tuning for robust input-output linearization control," in *Proc. 53rd IEEE Conf. Decis. Control*, 2014, pp. 2685–2690, doi: [10.1109/CDC.2014.7039800](https://doi.org/10.1109/CDC.2014.7039800).
- [20] B. Ren, P. Frihauf, R. J. Rafac, and M. Krstic, "Laser pulse shaping via extremum seeking," *Control Eng. Pract.*, vol. 20, no. 7, pp. 674–683, 2012, doi: [10.1016/j.conengprac.2012.03.006](https://doi.org/10.1016/j.conengprac.2012.03.006).
- [21] Z. Zhi-Dan, H. Hai-Bo, Z. Xin-Jian, C. Guang-Yi, and R. Yuan, "Adaptive maximum power point tracking control of fuel cell power plants," *J. Power Sources*, vol. 176, no. 1, pp. 259–269, 2008.
- [22] M. Benosman and J. I. Poveda, "Robust source seeking and formation learning-based controller," U.S. Patent 10,915,108 B2, 2019.
- [23] Y. Li and E. Seem, "Extremum seeking control with actuator saturation control," Patent: U.S. 2010/0106331 A1, 2010.
- [24] A. Subbaraman and M. Benosman, "Extremum seeking-based iterative learning model predictive control (ESILC-MPC)," in *Proc. 12th IFAC Workshop Adaptation Learn. Control Signal*, 2016, pp. 193–198, doi: [10.1016/j.ifacol.2016.07.950](https://doi.org/10.1016/j.ifacol.2016.07.950).
- [25] M. Abdelgalil, Y. Aboelkassem, and H. Taha, "Sea urchin sperm exploit extremum seeking control to find the egg," *Phys. Rev. E*, vol. 106, no. 6, 2022, Art. no. L062401, doi: [10.1103/PhysRevE.106.L062401](https://doi.org/10.1103/PhysRevE.106.L062401).
- [26] D. E. Ochoa and J. I. Poveda, "High-performance optimal incentive-seeking in transactive control for traffic congestion," *Eur. J. Control*, vol. 68, Nov. 2022, Art. no. 100696, doi: [10.1016/j.ejcon.2022.100696](https://doi.org/10.1016/j.ejcon.2022.100696).
- [27] H. Yu, S. Koga, T. R. Oliveira, and M. Krstic, "Extremum seeking for traffic congestion control with a downstream bottleneck," *J. Dyn. Syst. Meas. Contr.*, vol. 143, no. 3, 2021, Art. no. 031007, doi: [10.1115/1.4404878](https://doi.org/10.1115/1.4404878).
- [28] A. A. Elgohary and S. A. Eisa, "Extremum seeking for controlled vibrational stabilization of mechanical systems: A variation-of-constant averaging approach inspired by flapping insects mechanics," *IEEE Contr. Syst. Lett.*, vol. 9, pp. 264–269, 2025, doi: [10.1109/LCSYS.2025.3570767](https://doi.org/10.1109/LCSYS.2025.3570767).
- [29] R. Suttner, "Extremum-seeking control for a class of mechanical systems," *IEEE Trans. Autom. Control*, vol. 68, no. 2, pp. 1200–1207, Feb. 2023, doi: [10.1109/TAC.2022.3148377](https://doi.org/10.1109/TAC.2022.3148377).
- [30] R. Suttner, "Extremum seeking control for a class of nonholonomic systems," *SIAM J. Control Optim.*, vol. 58, no. 4, pp. 2588–2615, 2020.
- [31] S. Z. Khong, D. Nešić, Y. Tan, and C. Manzie, "Unified framework for sampled-data extremum seeking control: Global optimisation and multi-unit systems," *Automatica*, vol. 49, no. 9, pp. 2720–2733, 2013, doi: [10.1016/j.automatica.2013.06.020](https://doi.org/10.1016/j.automatica.2013.06.020).
- [32] Y. Zhu, E. Fridman, and T. R. Oliveira, "Sampled-data extremum seeking with constant delay: A time-delay approach," *IEEE Trans. Autom. Control*, vol. 68, no. 1, pp. 432–439, 2023, doi: [10.1109/TAC.2022.3140259](https://doi.org/10.1109/TAC.2022.3140259).
- [33] A. Scheinker, "100 years of extremum seeking: A survey," *Automatica*, vol. 161, Mar. 2024, Art. no. 111481, doi: [10.1016/j.automatica.2023.111481](https://doi.org/10.1016/j.automatica.2023.111481).
- [34] Y. Tan, W. H. C. Manzie, D. Nešić, and I. Mareels, "Extremum seeking from 1922 to 2010," in *Proc. 29th Chin. Control Conf.*, 2010, pp. 14–26.
- [35] J. Sternby, "A review of extremum control," Dept. of Autom. Control, Lund Inst. of Technol. (LTH), Lund, Sweden, Tech. Rep. TFRT-7161, 1979.
- [36] K. J. Astrom and B. Wittenmark, *Adaptive Control* (Dover Books on Electrical Engineering), 2nd ed. Mineola, NY, USA: Dover, 2008.
- [37] M. Leblanc, "Sur l'électrification des chemins de fer au moyen de courants alternatifs de fréquence élevée," *Revue générale de l'électricité*, vol. 12, no. 8, pp. 275–277, 1922.
- [38] I. S. Morosonov, "Method of extremum control," *Autom. Remote Control*, vol. 18, pp. 1077–1092, 1957.
- [39] S. M. Meerkov, "Asymptotic methods for investigating quasistationary states in continuous systems of automatic optimization," *Avotomatika i Tekhnika*, no. 11, pp. 119–139, 1967.
- [40] S. M. Meerkov, "The averaging of the trajectories of slow dynamical systems," *Differentsial'nye Uravneniya*, vol. 9, no. 9, pp. 1609–1617, 1973.
- [41] A. Krasovskii, *Dynamics of Continuous Self-Tuning Systems*. Moscow, Russia: Fizmatgiz, 1963.
- [42] P. Chinaev, "Self-tuning systems handbook," *Automn. Remote Control*, vol. 22, pp. 1–12, 1969.
- [43] R. J. Kutadinata, W. Moase, and C. Manzie, "Extremum-seeking in singularly perturbed hybrid systems," *IEEE Trans. Autom. Control*, vol. 62, no. 6, pp. 3014–3020, Jun. 2017, doi: [10.1109/TAC.2016.2607282](https://doi.org/10.1109/TAC.2016.2607282).
- [44] M. Krstić and H.-H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, no. 4, pp. 595–601, Apr. 2000, doi: [10.1016/S0005-1098\(99\)00183-1](https://doi.org/10.1016/S0005-1098(99)00183-1).
- [45] K. B. Ariyur and M. Krstić, *Real-Time Optimization by Extremum-Seeking Control*. New York, NY, USA: Wiley, 2003.
- [46] C. Zhang and R. Ordóñez, *Extremum-Seeking Control and Applications: A Numerical Optimization Based Approach*. London, U.K.: Springer-Verlag, 2012.
- [47] A. Scheinker and M. Krstic, *Model-Free Stabilization by Extremum Seeking*. Cham, Switzerland: Springer-Verlag, 2017.
- [48] H. Dürr, M. Stanković, C. Ebenbauer, and K. H. Johansson, "Lie bracket approximation of extremum seeking systems," *Automatica*, vol. 49, no. 6, pp. 1538–1552, 2013, doi: [10.1016/j.automatica.2013.02.016](https://doi.org/10.1016/j.automatica.2013.02.016).
- [49] W. H. Moase and C. Manzie, "Fast extremum-seeking for Wiener-Hammerstein plants," *Automatica*, vol. 48, no. 10, pp. 2433–2443, 2012, doi: [10.1016/j.automatica.2012.06.071](https://doi.org/10.1016/j.automatica.2012.06.071).
- [50] M. Benosman, *Learning-Based Adaptive Control: An Extremum Seeking Approach - Theory and Applications*. Cambridge, MA, USA: Butterworth-Heinemann, 2016.
- [51] P. Frihauf, M. Krstic, and T. Basar, "Nash equilibrium seeking in non-cooperative games," *IEEE Trans. Autom. Control*, vol. 57, no. 5, pp. 1192–1207, May 2012, doi: [10.1109/TAC.2011.2173412](https://doi.org/10.1109/TAC.2011.2173412).
- [52] S. Dougherty and M. Guay, "An extremum-seeking controller for distributed optimization over sensor networks," *IEEE Trans. Autom. Control*, vol. 62, no. 2, pp. 928–933, Feb. 2017, doi: [10.1109/TAC.2016.2566806](https://doi.org/10.1109/TAC.2016.2566806).
- [53] J. Poveda and N. Quijano, "Distributed extremum seeking for real-time resource allocation," in *Proc. Am. Control Conf.*, 2013, pp. 2772–2777.
- [54] J. Feiling, S. Koga, M. Krstić, and T. R. Oliveira, "Gradient extremum seeking for static maps with actuation dynamics governed by diffusion PDEs," *Automatica*, vol. 95, pp. 197–206, Sep. 2018, doi: [10.1016/j.automatica.2018.05.023](https://doi.org/10.1016/j.automatica.2018.05.023).
- [55] S. Liu and M. Krstic, "Stochastic Nash equilibrium seeking for games with general nonlinear payoffs," *SIAM J. Control Optim.*, vol. 49, no. 4, pp. 1659–1679, 2011, doi: [10.1137/100811738](https://doi.org/10.1137/100811738).
- [56] F. Verhulst, *Nonlinear Differential Equations and Dynamical Systems*. Berlin, Germany: Springer-Verlag, 1996.
- [57] H. K. Khalil, *Nonlinear Systems*. Hoboken, NJ, USA: Prentice Hall, 2002.
- [58] P. Kokotović, H. K. Khalil, and J. O'Reilly, *Singular Perturbation Methods in Control: Analysis and Design*. New York, NY, USA: Academic Press, 1986.
- [59] A. R. Teel, L. Moreau, and D. Nešić, "A unified framework for input-to-state stability in systems with two time scales," *IEEE Trans. Autom. Control*, vol. 48, no. 9, pp. 1526–1544, Sep. 2003, doi: [10.1109/TAC.2003.816966](https://doi.org/10.1109/TAC.2003.816966).
- [60] F. Lamnabhi-Lagarrigue et al., "Systems & control for the future of humanity, research agenda: Current and future roles, impact and grand challenges," *Annu. Rev. Control*, vol. 43, pp. 1–64, 2017, doi: [10.1016/j.arcon.2017.04.001](https://doi.org/10.1016/j.arcon.2017.04.001).
- [61] R. G. Sanfelice, *Analysis and Design of Cyber-Physical Systems: A Hybrid Control Systems Approach*. Boca Raton, FL, USA: CRC Press, 2016.
- [62] J. C. Luxat and L. H. Lees, "Stability of peak-holding control systems," *IEEE Trans. Ind. Electron. Control Instrum.*, vol. IECI-18, no. 1, pp. 11–15, Feb. 1971, doi: [10.1109/TIECI.1971.230455](https://doi.org/10.1109/TIECI.1971.230455).
- [63] S. J. Moura and Y. A. Chang, "Lyapunov-based switched extremum seeking for photovoltaic power maximization," *Control Eng. Pract.*, vol. 21, no. 7, pp. 971–980, 2013, doi: [10.1016/j.conengprac.2013.02.009](https://doi.org/10.1016/j.conengprac.2013.02.009).
- [64] D. Nešić, Y. Tan, C. Manzie, A. Mohammadi, and W. Moase, "A unifying framework for analysis and design of extremum seeking controllers," in *Proc. IEEE Chin. Control Decis. Conf.*, 2012, pp. 4274–4285, doi: [10.1109/CCDC.2012.6244001](https://doi.org/10.1109/CCDC.2012.6244001).
- [65] D. Nešić, Y. Tan, W. H. Moase, and C. Manzie, "A unifying approach to extremum seeking: Adaptive schemes based on estimation of derivatives," in *Proc. 49th IEEE Conf. Decis. Control*, 2010, pp. 4625–4630, doi: [10.1109/CDC.2010.5717929](https://doi.org/10.1109/CDC.2010.5717929).

- [66] Y. Tan, D. Nešić, and I. Mareels, "On non-local stability properties of extremum seeking controllers," *Automatica*, vol. 42, no. 6, pp. 889–903, 2006, doi: [10.1016/j.automatica.2006.01.014](https://doi.org/10.1016/j.automatica.2006.01.014).
- [67] N. Mimmo, L. Marconi, and G. Notarstefano, "Uniform nonconvex optimization via extremum seeking," *IEEE Trans. Autom. Control*, vol. 69, no. 12, pp. 8263–8276, Dec. 2024, doi: [10.1109/TAC.2024.3406229](https://doi.org/10.1109/TAC.2024.3406229).
- [68] A. R. Teel and L. Praly, "A smooth Lyapunov function from a Class-KL estimate involving two positive semidefinite functions," *ESAIM: Control, Optim. Calculus Variations*, vol. 5, pp. 313–367, Jun. 2000, doi: [10.1051/cocv:2000113](https://doi.org/10.1051/cocv:2000113).
- [69] J. I. Poveda and N. Li, "Robust hybrid zero-order optimization algorithms with acceleration via averaging in continuous time," *Automatica*, vol. 123, Jan. 2021, Art. no. 109361, doi: [10.1016/j.automatica.2020.109361](https://doi.org/10.1016/j.automatica.2020.109361).
- [70] X. Chen, J. I. Poveda, and N. Li, "Continuous-time zeroth-order dynamics with projection maps: Model-free feedback optimization with safety guarantees," *IEEE Trans. Autom. Control*, vol. 70, no. 8, pp. 5005–5020, Aug. 2025, doi: [10.1109/TAC.2025.3537956](https://doi.org/10.1109/TAC.2025.3537956).
- [71] J. Poveda and N. Quijano, "A Shahshahani gradient based extremum seeking scheme," in *Proc. Conf. Decis. Control (CDC)*, 2012, pp. 5104–5109, doi: [10.1109/CDC.2012.6426134](https://doi.org/10.1109/CDC.2012.6426134).
- [72] A. Ghaffari, M. Krstić, and D. Nešić, "Multivariable newton-based extremum seeking," *Automatica*, vol. 48, no. 8, pp. 1759–1767, 2012, doi: [10.1016/j.automatica.2012.05.059](https://doi.org/10.1016/j.automatica.2012.05.059).
- [73] M. Ye, Q.-L. Han, L. Ding, and S. Xu, "Distributed Nash equilibrium seeking in games with partial decision information: A survey," *Proc. IEEE*, vol. 111, no. 2, pp. 140–157, Feb. 2023, doi: [10.1109/JPROC.2023.3234687](https://doi.org/10.1109/JPROC.2023.3234687).
- [74] R. Goebel, R. G. Sanfelice, and A. R. Teel, "Hybrid dynamical systems," *IEEE Control Syst. Mag.*, vol. 29, no. 2, pp. 28–93, Apr. 2009, doi: [10.1109/MCS.2008.931718](https://doi.org/10.1109/MCS.2008.931718).
- [75] R. K. Goebel, *Set-Valued, Convex, and Nonsmooth Analysis in Dynamics and Control: An Introduction*. Philadelphia, PA, USA: SIAM, 2024.
- [76] D. Liberzon, *Switching in Systems and Control*. Boston, MA, USA: Birkhauser, 2003.
- [77] G. Bianchin, J. I. Poveda, and E. Dall'Anese, "Online optimization of switched LTI systems using continuous-time and hybrid accelerated gradient flows," *Automatica*, vol. 146, Dec. 2022, Art. no. 110579.
- [78] W. M. Haddad, V. Chellaboina, and S. G. Nersesov, *Impulsive and Hybrid Dynamical Systems: Stability, Dissipativity, and Control*. NJ, Princeton, NJ, USA: Princeton Univ. Press, 2006.
- [79] K. Astrom and B. Bernhardsson, "Comparison of periodic and event based sampling for first order stochastic systems," in *Proc. IFAC World Conf.*, 1999, pp. 301–306.
- [80] J. Hespanha and A. Morse, "Stability of switched systems with average dwell-time," in *Proc. 38th IEEE Conf. Decis. Control*, 1999, vol. 3, pp. 2655–2660, doi: [10.1109/CDC.1999.831330](https://doi.org/10.1109/CDC.1999.831330).
- [81] H. Hajieghrary, D. Mox, and M. A. Hsieh, "Information theoretic source seeking strategies for multiagent plume tracking in turbulent fields," *J. Mar. Sci. Eng.*, vol. 5, no. 1, 2017, Art. no. 3, doi: [10.3390/jmse5010003](https://doi.org/10.3390/jmse5010003).
- [82] N. Ghods and M. Krstić, "Multi-agent deployment around a source in one dimension by extremum seeking," in *Proc. Amer. Control Conf.*, 2010, pp. 4794–4799, doi: [10.1109/ACC.2010.5531113](https://doi.org/10.1109/ACC.2010.5531113).
- [83] J. I. Poveda, M. Benosman, A. R. Teel, and R. G. Sanfelice, "Robust coordinated hybrid source seeking with obstacle avoidance in multivehicle autonomous systems," *IEEE Trans. Autom. Control*, vol. 67, no. 2, pp. 706–721, Feb. 2022, doi: [10.1109/TAC.2021.3056365](https://doi.org/10.1109/TAC.2021.3056365).
- [84] C. Zhang, A. Siranosian, and M. Krstić, "Extremum seeking for moderately unstable systems and for autonomous vehicle target tracking without position measurements," *Automatica*, vol. 43, no. 10, pp. 1832–1839, 2007, doi: [10.1016/j.automatica.2007.03.009](https://doi.org/10.1016/j.automatica.2007.03.009).
- [85] A. Morse, "Supervisory control of families of linear set-point controllers-part I: Exact matching," *IEEE Trans. Autom. Control*, vol. 41, no. 10, pp. 1413–1431, Oct. 1996, doi: [10.1109/9.539424](https://doi.org/10.1109/9.539424).
- [86] M. Baradaran and A. R. Teel, "Omega-limit sets and robust stability for switched systems with distinct equilibria," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 2039–2044, 2020, doi: [10.1016/j.ifacol.2020.12.2515](https://doi.org/10.1016/j.ifacol.2020.12.2515).
- [87] R. T. Rockafellar and R. J. Wets, *Variational Analysis*. Berlin, Germany: Springer-Verlag, 1998.
- [88] D. Popović, "Topics in extremum seeking," Ph.D. thesis, Univ. of California, Santa Barbara, CA, USA, 2004.
- [89] S. Z. Khong, Y. Tan, C. Manzie, and D. Nesic, "Multi-agent source seeking via discrete-time extremum seeking control," *Automatica*, vol. 50, no. 9, pp. 2312–2320, 2014, doi: [10.1016/j.automatica.2014.06.009](https://doi.org/10.1016/j.automatica.2014.06.009).
- [90] D. Ochoa and J. I. Poveda, "Momentum-based Nash set-seeking over networks via multitime scale hybrid dynamic inclusions," *IEEE Trans. Autom. Control*, vol. 69, no. 7, pp. 4245–4260, Jul. 2024, doi: [10.1109/TAC.2023.3321901](https://doi.org/10.1109/TAC.2023.3321901).
- [91] J. I. Poveda and A. R. Teel, "A framework for a class of hybrid extremum seeking controllers with dynamic inclusions," *Automatica*, vol. 76, pp. 113–126, Feb. 2017, doi: [10.1016/j.automatica.2016.10.029](https://doi.org/10.1016/j.automatica.2016.10.029).
- [92] J. I. Poveda and A. R. Teel, "The heavy-ball ODE with time-varying damping: Persistence of excitation and uniform asymptotic stability," in *Proc. Amer. Control Conf.*, 2020, pp. 773–778, doi: [10.23919/ACC45564.2020.9147733](https://doi.org/10.23919/ACC45564.2020.9147733).
- [93] J. I. Poveda and N. Li, "Inducing uniform asymptotic stability in non-autonomous accelerated optimization dynamics via hybrid regularization," in *Proc. 58th IEEE Conf. Decis. Control*, 2019, pp. 3000–3005, doi: [10.1109/CDC40024.2019.9030127](https://doi.org/10.1109/CDC40024.2019.9030127).
- [94] A. R. Teel, J. I. Poveda, and J. Le, "First-order optimization algorithms with resets and hamiltonian flows," in *Proc. IEEE 58th Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 5838–5843, doi: [10.1109/CDC40024.2019.9029333](https://doi.org/10.1109/CDC40024.2019.9029333).
- [95] M. Abdelgalil and J. I. Poveda, "Initialization-free Lie-bracket extremum seeking," *Syst. Control Lett.*, vol. 191, Sep. 2024, Art. no. 105881, doi: [10.1016/j.sysconle.2024.105881](https://doi.org/10.1016/j.sysconle.2024.105881).
- [96] R. Suttner, "Extremum seeking control with an adaptive dither signal," *Automatica*, vol. 101, pp. 214–222, Mar. 2019, doi: [10.1016/j.automatica.2018.11.055](https://doi.org/10.1016/j.automatica.2018.11.055).
- [97] M. Weber, B. Gharesifard, and C. Ebenbauer, "Inferring global exponential stability properties using Lie-bracket approximations," in *Proc. IEEE 63rd Conf. Decis. Control*, 2024, pp. 6274–6279, doi: [10.1109/CDC56724.2024.10886836](https://doi.org/10.1109/CDC56724.2024.10886836).
- [98] R. Kutadinata, W. Moase, C. Manzie, L. Zhang, and T. Garoni, "Enhancing the performance of existing urban traffic light control through extremum-seeking," *Transp. Res. Part C: Emerg. Technol.*, vol. 62, pp. 1–20, Jan. 2016, doi: [10.1016/j.trc.2015.10.016](https://doi.org/10.1016/j.trc.2015.10.016).
- [99] R. Sanfelice and A. R. Teel, "Dynamical properties of hybrid systems simulators," *Automatica*, vol. 46, no. 2, pp. 239–248, 2010, doi: [10.1016/j.automatica.2009.09.026](https://doi.org/10.1016/j.automatica.2009.09.026).
- [100] W. Wang, A. R. Teel, and D. Nešić, "Averaging in singularly perturbed hybrid systems with hybrid boundary layer systems," in *Proc. 51st IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2012, pp. 6855–6860, doi: [10.1109/CDC.2012.6425965](https://doi.org/10.1109/CDC.2012.6425965).
- [101] V. Grushkovskaya, H.-B. Durr, C. Ebenbauer, and A. Zuyev, "Extremum seeking for time-varying functions using Lie-bracket approximations," *IFAC-Papers Online*, vol. 50, no. 1, pp. 5522–5528, 2017, doi: [10.1016/j.ifacol.2017.08.1093](https://doi.org/10.1016/j.ifacol.2017.08.1093).
- [102] M. Abdelgalil, D. E. Ochoa, and J. I. Poveda, "Multi-time scale control and optimization via averaging and singular perturbation theory: From ODE's to hybrid dynamical systems," *Annu. Rev. Control*, vol. 56, Jan. 2023, Art. no. 100926, doi: [10.1016/j.arcontrol.2023.100926](https://doi.org/10.1016/j.arcontrol.2023.100926).
- [103] M. Abdelgalil and J. I. Poveda, "On Lie-Bracket averaging for a class of hybrid dynamical systems with applications to model-free control and optimization," *IEEE Trans. Autom. Control*, vol. 70, no. 7, pp. 4655–4670, Jul. 2025, doi: [10.1109/TAC.2025.3529375](https://doi.org/10.1109/TAC.2025.3529375).
- [104] R. Kutadinata, W. H. Moase, and C. Manzie, "Dither re-use in Nash equilibrium seeking," *IEEE Trans. Autom. Control*, vol. 60, no. 5, pp. 1433–1438, May 2015, doi: [10.1109/TAC.2014.2347212](https://doi.org/10.1109/TAC.2014.2347212).
- [105] M. Stanković, K. H. Johansson, and D. M. Stipanović, "Distributed seeking of Nash equilibria with applications to mobile sensor networks," *IEEE Trans. Autom. Control*, vol. 57, no. 4, pp. 904–919, Apr. 2012, doi: [10.1109/TAC.2011.2174678](https://doi.org/10.1109/TAC.2011.2174678).
- [106] J. I. Poveda and N. Quijano, "Shahshahani gradient-like extremum seeking," *Automatica*, vol. 58, pp. 51–59, Aug. 2015, doi: [10.1016/j.automatica.2015.05.002](https://doi.org/10.1016/j.automatica.2015.05.002).
- [107] J. I. Poveda, M. Krstić, and T. Başar, "Fixed-time Nash equilibrium seeking in time-varying networks," *IEEE Trans. Autom. Control*, vol. 68, no. 4, pp. 1954–1969, Apr. 2023, doi: [10.1109/TAC.2022.3168527](https://doi.org/10.1109/TAC.2022.3168527).
- [108] J. I. Poveda, R. Kutadinata, C. Manzie, D. Nešić, A. R. Teel, and C. Liao, "Hybrid extremum seeking for black-box optimization in hybrid plants: An analytical framework," in *Proc. 57th IEEE Conf. Decis. Control*, 2018, pp. 2235–2240, doi: [10.1109/CDC.2018.8618907](https://doi.org/10.1109/CDC.2018.8618907).
- [109] W. Moase, Y. Tan, D. Nešić, and C. Manzie, "Non-local stability of a multi-variable extremum-seeking scheme," in *Proc. IEEE Australian Control Conf.*, 2011, pp. 38–43.
- [110] E. D. Sontag, "Stability and stabilization: Discontinuities and the effect of disturbances," in *Nonlinear Analysis, Differential Equations and*

- Control, F. H. Clarke, R. J. Stern, and G. Sabidussi, Eds., Berlin, Germany: Springer-Verlag, 1999, pp. 551–598.
- [111] T. Strizic, J. I. Poveda, and A. R. Teel, “Hybrid gradient descent for robust global optimization on the circle,” in *Proc. 56th IEEE Conf. Decis. Control*, 2017, pp. 2985–2990, doi: [10.1109/CDC.2017.8264093](https://doi.org/10.1109/CDC.2017.8264093).
- [112] D. Ochoa and J. I. Poveda, “Robust global optimization on smooth compact manifolds via hybrid gradient-free dynamics,” *Automatica*, vol. 171, Jan. 2025, Art. no. 111916, doi: [10.1016/j.automatica.2024.111916](https://doi.org/10.1016/j.automatica.2024.111916).
- [113] B. O’Donoghue and E. J. Candes, “Adaptive restart for accelerated gradient schemes,” *Found. Comput. Math.*, vol. 15, no. 3, pp. 715–732, 2015, doi: [10.1007/s10208-013-9150-3](https://doi.org/10.1007/s10208-013-9150-3).
- [114] W. Su, S. Boyd, and E. Candes, “A differential equation for modeling Nesterov’s accelerated gradient method: Theory and insights,” *J. Mach. Learn. Res.*, vol. 17, no. 153, pp. 1–43, 2016.
- [115] F. Galarza-Jimenez, J. I. Poveda, R. Kutadinata, L. Zhang, and E. Dall’Anese, “Self-optimizing traffic light control using hybrid accelerated extremum seeking,” in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 1941–1946, doi: [10.1109/CDC45484.2021.9683507](https://doi.org/10.1109/CDC45484.2021.9683507).
- [116] F. Galarza-Jimenez, J. I. Poveda, G. Bianchi, and E. Dall’Anese, “Extremum seeking under persistent gradient deception: A switching systems approach,” *IEEE Contr. Syst. Lett.*, vol. 6, pp. 133–138, 2022, doi: [10.1109/LC-SYS.2021.3050451](https://doi.org/10.1109/LC-SYS.2021.3050451).
- [117] J. I. Poveda and M. Krstic, “Fixed-time Newton-like extremum seeking,” *IFAC-PapersOnLine*, vol. 53, no. 2, 2020, pp. 5356–5361, 2020.
- [118] W. Sandholm, *Population Games and Evolutionary Dynamics*. Cambridge, MA, USA: MIT Press, 2010.
- [119] F. H. Clarke, *Optimization and Nonsmooth Analysis*. Philadelphia, PA, USA: SIAM, 1990.
- [120] F. H. Clarke, Y. S. Ledyaev, R. J. Stern, and P. R. Wolenski, *Nonsmooth Analysis and Control Theory (Graduate Texts in Mathematics)*. New York, NY, USA: Springer-Verlag, 1998.
- [121] C. Henry, “An existence theorem for a class of differential equations with multivalued right-hand side,” *J. Math. Anal. Appl.*, vol. 41, no. 1, pp. 179–186, 1973, doi: [10.1016/0022-247X\(73\)90192-3](https://doi.org/10.1016/0022-247X(73)90192-3).
- [122] F. Galarza-Jimenez, J. Poveda, and E. Dall’Anese, “Sliding-seeking control: Model-free optimization with safety constraints,” in *Proc. Learn. Dyn. Control Conf.*, PMLR, 2022, pp. 1100–1111.
- [123] J. I. Poveda and M. Krstić, “Fixed-time seeking and tracking of time-varying extrema,” in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, pp. 108–113, 2021, doi: [10.1109/CDC45484.2021.9683248](https://doi.org/10.1109/CDC45484.2021.9683248).
- [124] J. Poveda and M. Krstic, “Nonsmooth extremum seeking control with user-prescribed fixed-time convergence,” *IEEE Trans. Autom. Control*, vol. 66, no. 12, pp. 6156–6163, Dec. 2021, doi: [10.1109/TAC.2021.3063700](https://doi.org/10.1109/TAC.2021.3063700).
- [125] R. Kamalapurkar, J. A. Rosenfeld, A. Parikh, A. R. Teel, and W. E. Dixon, “Invariance-like results for nonautonomous switched systems,” *IEEE Trans. Autom. Control*, vol. 64, no. 2, pp. 614–627, Feb. 2019, doi: [10.1109/TAC.2018.2838055](https://doi.org/10.1109/TAC.2018.2838055).
- [126] A. Bhaya and E. Kaszkurewicz, *Control Perspectives on Numerical Algorithms and Matrix Problems*. Philadelphia, PA, USA: SIAM, 2006.
- [127] R. Olfati-Saber and P. Murray, “Consensus problems in networks of agents with switching topology and time-delays,” *IEEE Trans. Autom. Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004, doi: [10.1109/TAC.2004.834113](https://doi.org/10.1109/TAC.2004.834113).
- [128] J. I. Poveda, M. Benosman, and A. R. Teel, “Distributed extremum seeking in multi-agent systems with arbitrary switching graphs,” *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 735–740, 2017, doi: [10.1016/j.ifacol.2017.08.240](https://doi.org/10.1016/j.ifacol.2017.08.240).
- [129] P. Lin, W. Ren, and J. A. Farrell, “Distributed continuous-time optimization: Nonuniform gradient gains, finite-time convergence, and convex constraint set,” *IEEE Trans. Autom. Control*, vol. 62, no. 5, pp. 2239–2253, 2017, doi: [10.1109/TAC.2016.2604324](https://doi.org/10.1109/TAC.2016.2604324).
- [130] C. Cai, A. R. Teel, and R. Goebel, “Smooth Lyapunov functions for hybrid systems. Part II: (Pre)asymptotically stable compact sets,” *IEEE Trans. Autom. Control*, vol. 53, no. 3, pp. 734–748, Apr. 2008, doi: [10.1109/TAC.2008.919257](https://doi.org/10.1109/TAC.2008.919257).
- [131] F. Galarza-Jimenez, G. Bianchin, J. I. Poveda, and E. Dall’Anese, “Online optimization of LTI systems under persistent attacks: Stability, tracking, and robustness,” *Nonlinear Anal. Hybrid Syst.*, vol. 44, May 2022, Art. no. 101152.
- [132] A. Scheinker and M. Krstic, “Extremum seeking-based tracking for unknown systems with unknown control directions,” in *Proc. 51st IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2012, pp. 6065–6070, doi: [10.1109/CDC.2012.6426821](https://doi.org/10.1109/CDC.2012.6426821).
- [133] C. Labar, C. Ebenbauer, and L. Marconi, “ISS-like properties in Lie-bracket approximations and application to extremum seeking,” *Automatica*, vol. 136, Feb. 2022, Art. no. 110041, doi: [10.1016/j.automatica.2021.110041](https://doi.org/10.1016/j.automatica.2021.110041).
- [134] M. Jankovic, M. Santillo, and Y. Wang, “Multiagent systems with CBF-based controllers: Collision avoidance and liveness from instability,” *IEEE Trans. Control Syst. Technol.*, vol. 32, no. 2, pp. 705–712, Mar. 2024, doi: [10.1109/TCST.2023.3324531](https://doi.org/10.1109/TCST.2023.3324531).
- [135] M. Marley, R. Skjetne, and A. R. Teel, “Hybrid control barrier functions for continuous-time systems,” *IEEE Trans. Autom. Control*, vol. 69, no. 10, pp. 6605–6619, Oct. 2024, doi: [10.1109/TAC.2024.3374265](https://doi.org/10.1109/TAC.2024.3374265).
- [136] D. E. Koditschek and E. Rimon, “Robot navigation functions on manifolds with boundary,” *Adv. Appl. Math.*, vol. 11, no. 4, pp. 412–442, 1990, doi: [10.1016/0196-8858\(90\)90017-S](https://doi.org/10.1016/0196-8858(90)90017-S).
- [137] R. G. Sanfelice, M. J. Messina, S. E. Tuna, and A. R. Teel, “Robust hybrid controllers for continuous-time systems with applications to obstacle avoidance and regulation to disconnected set of points,” in *Proc. Am. Control Conf.*, 2006, pp. 3352–3357.
- [138] S. Berkane, A. Bisoffi, and D. V. Dimarogonas, “Obstacle avoidance via hybrid feedback,” *IEEE Trans. Autom. Control*, vol. 67, no. 1, pp. 512–519, Jan. 2022, doi: [10.1109/TAC.2021.3086329](https://doi.org/10.1109/TAC.2021.3086329).
- [139] M. Sawant, I. Polushin, and A. Tayebi, “Hybrid feedback control design for nonconvex obstacle avoidance,” *IEEE Trans. Autom. Control*, vol. 69, no. 11, pp. 7508–7523, Nov. 2024, doi: [10.1109/TAC.2024.3388952](https://doi.org/10.1109/TAC.2024.3388952).
- [140] P. Jain and P. Kar, “Non-convex optimization for machine learning,” *Found. Trends Mach. Learn.*, vol. 10, nos. 3–4, pp. 142–363, 2017.
- [141] D. Netic, A. R. Teel, and L. Zaccarian, “Stability and performance of SISO control systems with first-order reset elements,” *IEEE Trans. Autom. Control*, vol. 56, no. 11, pp. 2567–2582, Nov. 2011, doi: [10.1109/TAC.2011.2114436](https://doi.org/10.1109/TAC.2011.2114436).
- [142] G. Zhao, D. Nešić, Y. Tan, and C. Hua, “Overcoming overshoot performance limitations of linear systems with reset control,” *Automatica*, vol. 101, pp. 27–35, Mar. 2019, doi: [10.1016/j.automatica.2018.11.038](https://doi.org/10.1016/j.automatica.2018.11.038).
- [143] J. Zhang, A. Mokhtari, S. Sra, and A. Jadbabaie, “Direct Runge-Kutta discretization achieves acceleration,” in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2018, pp. 3904–3913.
- [144] S. Michalowsky and C. Ebenbauer, “The multidimensional n-th order heavy ball method and its application to extremum seeking,” in *Proc. 53rd IEEE Conf. Decis. Control*, Piscataway, NJ, USA: IEEE Press, 2014, pp. 2660–2666, doi: [10.1109/CDC.2014.7039796](https://doi.org/10.1109/CDC.2014.7039796).
- [145] M. Laborde and A. Oberman, “A Lyapunov analysis for accelerated gradient methods: From deterministic to stochastic case,” in *Proc. 23rd Int. Conf. Artif. Intell. Statist.*, PMLR, 2020, pp. 602–612.
- [146] A. Scheinker and M. Krstić, “Non-C2 Lie-bracket averaging for nonsmooth extremum seekers,” *ASME. J. Dyn. Sys., Meas., Control*, vol. 136, no. 1, 2013, Art. no. 011010.
- [147] S. Krilašević and S. Grammatico, “Learning generalized Nash equilibria in monotone games: A hybrid adaptive extremum seeking control approach,” *Automatica*, vol. 151, May 2023, Art. no. 110931, doi: [10.1016/j.automatica.2023.110931](https://doi.org/10.1016/j.automatica.2023.110931).
- [148] J. I. Poveda and A. R. Teel, “Hybrid mechanisms for robust synchronization and coordination of multi-agent networked sampled-data systems,” *Automatica*, vol. 99, pp. 41–53, Jan. 2019, doi: [10.1016/j.automatica.2018.10.010](https://doi.org/10.1016/j.automatica.2018.10.010).
- [149] J. Martens and I. Sutskever, “Training deep and recurrent networks with hessian-free optimization,” in *Neural Networks: Tricks of the Trade*, 2nd ed. Berlin, Germany: Springer-Verlag, 2012, pp. 479–535.
- [150] H. Attouch and P. Redont, “The second-order in time continuous newton method,” in *Approximation, Optimization and Mathematical Economics*, M. Lassonde, Ed., Heidelberg, Germany: Physica, 2001, pp. 25–36.
- [151] C. Labar, E. Garone, M. Kinnaert, and C. Ebenbauer, “Newton-based extremum seeking: A second-order Lie-bracket approximation approach,” *Automatica*, vol. 105, pp. 356–367, Jul. 2019, doi: [10.1016/j.automatica.2019.04.010](https://doi.org/10.1016/j.automatica.2019.04.010).
- [152] C. Labar, C. Ebenbauer, and L. Marconi, “Extremum seeking with intermittent measurements: A Lie-brackets approach,” *IEEE Trans. Autom. Control*, vol. 67, no. 12, pp. 6968–6974, Dec. 2022, doi: [10.1109/TAC.2022.3171306](https://doi.org/10.1109/TAC.2022.3171306).
- [153] V. Grushkovskaya, A. Zuyev, and C. Ebenbauer, “On a class of generating vector fields for the extremum seeking problem: Lie bracket approximation and stability properties,” *Automatica*, vol. 94, pp. 151–160, Aug. 2018, doi: [10.1016/j.automatica.2018.04.024](https://doi.org/10.1016/j.automatica.2018.04.024).
- [154] Y. V. Orlov, *Discontinuous Systems: Lyapunov Analysis and Robust Synthesis under Uncertainty Conditions*. London, U.K.: Springer-Verlag, 2009.
- [155] D. Shevitz and B. Paden, “Lyapunov stability theory of nonsmooth systems,” *IEEE Trans. Autom. Control*, vol. 39, no. 9, pp. 1910–1914, Sep. 1994, doi: [10.1109/9.317122](https://doi.org/10.1109/9.317122).
- [156] M. Abdelgalil and J. I. Poveda, “Hybrid minimum-seeking in synergistic Lyapunov functions: Robust global stabilization under unknown control directions,” 2024, *arXiv:2408.04882*.

Examples of Extremum Seeking Control in Industry

AUTOTUNING APPLICATIONS FROM ELECTROMAGNETIC BRAKES TO FLUID DYNAMICS

MOUHACINE BENOSMAN 

Extremum seeking control (ESC) is a mathematically rigorous data-driven framework, extensively researched over the last two decades. At its core, ESC is based on designing a closed-loop system aimed at optimizing a cost function, whether static or dynamic, without the need for its explicit expression or gradient, characterizing it as a zero-order optimization method within the broader domain of continuous optimization. This article recalls the foundational principles of ESC, then pivots to ESC's practical applications across various industries (see "Summary"). We highlight ESC's role in dynamically tuning electromagnetic brakes in high-rise elevators, which promotes durability and performance. Furthermore, we detail the application of ESC in real-time enhancement of 4G/5G digital amplifiers—crucial components in the ubiquitous digital environment. Finally, we explore how ESC can be used in fluidic systems, more specifically, heating, ventilation, and air-conditioning (HVAC) systems, demonstrating its capability to dynamically adjust the intricate interplay between temperature and airflow, thereby achieving optimal environmental conditions.

Digital Object Identifier 10.1109/MCS.2025.3646338

Date of current version: 26 March 2026

(GRAPHIC FROM AUTHOR-GENERATED ARTWORK USING OPENAI CHATGPT. FOR IMAGE CREDITS, PLEASE SEE THE 'ACKNOWLEDGMENT' SECTION OF THIS ARTICLE.)

Extremum seeking control is a mathematically rigorous data-driven framework, extensively researched over the last two decades.

ESC was first introduced in France, in 1920, to regulate the electrical voltage of the railway [1]. However, it was not until 2000, eight decades later, that a comprehensive theoretical framework was established to analyze the convergence of ESC for dynamical maps modeled by ordinary differential equations (ODEs) [2]. The groundbreaking work in [2] was subsequently followed by numerous technical papers exploring ESC across various system classes, from stochastic systems [3] to multiagent systems [4] and, more recently, hybrid systems [5]. This extensive research has culminated in a series of outstanding books on the subject [6], [7], [8].

Today, feedback controllers are implemented in a variety of real-time systems. Several types of feedback exist, for example, state or output feedback and linear or nonlinear feedback. However, one common characteristic of all available feedback controllers is the fact that they all rely on some “well chosen” feedback gains. The selection of these feedback gains is often done based on some desired performance. For instance, the gains can be chosen to minimize the overshoot of a linear closed-loop system. Settling time can be another performance target. Minimizing a given finite-time or asymptotic, state, or output tracking error can be of interest in many applications as well.

Over the years, there have been myriad results on feedback gain tuning. Maybe one of the most famous and widely taught techniques is the so-called Ziegler–Nichols rules for proportional-integral-derivative (PID) gain tuning for linear systems [9]. However, such rules apply for the particular class of linear systems under linear PID feedback and are considered heuristic in nature. For more general cases of models and controllers, with a more systematic or autonomous way of tuning feedback gains, the control community started looking at an iterative procedure to autotune feedback gains for closed-loop systems.

Indeed, in the seminal paper [10], the authors introduced the idea that feedback controllers’ parameters could be tuned iteratively to compensate for model uncertainties and that the tuning could be based on measurements, which are directly obtained from the system. This idea of iterative control tuning led to the so-called iterative feedback tuning (IFT) research field, where the goal is to iteratively autotune feedback gains of closed-loop systems, based on the online optimization of a well-defined performance cost function.

There have been a lot of results on IFT in the past 20 years, and it is not the purpose of this article to survey

all existing works in the field. However, these results are mainly dedicated to linear systems controlled with linear feedback, for example, [11], [12], [13], [14], [15], and [16]. Based on these IFT algorithms for linear systems, some extensions to nonlinear systems have been studied. For instance, in [17], the author studied the case of discrete nonlinear systems controlled with linear time-invariant output feedback. The effect of IFT algorithms, developed originally for linear systems, was studied on nonlinear systems by assuming local Taylor approximation of nonlinear dynamics. However, the full analysis of the feedback loop, that is, IFT merged with the linear controller and the nonlinear dynamics, was not reported in this paper. Other feedback gain iterative tuning algorithms were developed for nonlinear systems in [18], [19], and [20]. The algorithms developed in these papers first assume that the closed-loop input and output signals remain bounded during the gain tuning, and then they rely on the numerical estimation of the gradient of a given cost function with respect to the controller gains, which necessitates running the system $n + 2$ times if the dimension of the tuned gain vector is n . This obviously can be a limiting factor if the number of tuned parameters is large.

In this article, we report some recent results that rely on ESC (refer to “ESC: Basic Ideas” for a succinct

Summary

Real-life systems are often faced with the problem of having to adapt to changing operating conditions. For example, the aging of system components or changes in environmental conditions might lead to important drifts of the nominal operating conditions. As a result, it is of paramount importance to consider the robustness and adaptability of control modules meant for real-time deployment. One way of achieving such robustness and adaptability is by fine-tuning the controller in real time when assessing its interaction with the environment. This is mostly known as “iterative feedback tuning” (“IFT”) in the control community. We propose to look at this problem of IFT from the perspective of ESC. Indeed, ESC is a zero-order data-driven optimization method that could be readily leveraged for the IFT of real-time controllers, where the optimization cost function encompasses the feedback tuning performance target. In this article, we present three industrial applications of ESC in the context of IFT.

ESC: Basic Ideas

To provide the reader with an understanding of how ESC methods function, we introduce a few simple ESC algorithms. Let us begin by considering the following general dynamics:

$$\dot{x} = f(x, u) \quad (\text{S1})$$

where $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{R}$ is the scalar control (for simplicity), and $f: \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ is a smooth function. Now, assume that the equation given by (S1) describes the behavior of a real-world system. The objective of the control is to optimize the system's performance, which could range from a simple task, such as regulating a specific output to a constant target value, to a more complex one, such as tracking a desired time-varying trajectory. We model the desired performance with a smooth function $J(x, u): \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$, which can be simplified to $J(u)$ since the state x is influenced by u . To establish convergence results, we must assume the following conditions.

Assumption 1

There exists a smooth function $l: \mathbb{R} \rightarrow \mathbb{R}^n$ such that

$$f(x, u) = 0, \text{ if and only if } x = l(u). \quad (\text{S2})$$

Assumption 2

For each $u \in \mathbb{R}$, the equilibrium $x = l(u)$ of the system (S1) is locally exponentially stable.

Assumption 3

There exists (a maximum) $u^* \in \mathbb{R}$ such that

$$\begin{aligned} (J \circ l)^{(1)}(u^*) &= 0 \\ (J \circ l)^{(2)}(u^*) &< 0. \end{aligned} \quad (\text{S3})$$

Based on these assumptions, we can design some simple extremum seekers with proven convergence bounds. Indeed, one of the simplest ways to maximize J is to use gradient-based ESC, as follows:

$$\dot{u} = k \frac{dJ}{du}, \quad k > 0. \quad (\text{S4})$$

We can analyze the convergence of the ESC algorithm (S4) by using the Lyapunov function

$$V = J(u^*) - J(u) > 0, \text{ for } u \neq u^*. \quad (\text{S5})$$

The derivative of V leads to

$$\dot{V} = \frac{dJ}{du} \dot{u} = -k \left(\frac{dJ}{du} \right)^2 \leq 0. \quad (\text{S6})$$

This proves that algorithm (S4) drives u to the invariant set subject to $dJ/du = 0$, which is (by Assumption 3) equivalent to $u = u^*$. However, as simple as algorithm (S4) might seem, it still requires knowledge of the gradient of J . To overcome this requirement, one can instead use an algorithm motivated by the sliding mode control ideas. For instance, we can define the tracking error

$$e = J(u) - \text{ref}(t) \quad (\text{S7})$$

where "ref" denotes a time function that is monotonically increasing. The idea is that if J tracks ref, then it increases until it reaches an invariant set centered around the equality $dJ/du = 0$. A simple way to achieve this goal is by choosing the following ESC law:

$$\dot{u} = k_1 \text{sgn} \left(\sin \left(\frac{\pi e}{k_2} \right) \right), \quad k_1, k_2 > 0. \quad (\text{S8})$$

This controller is shown to steer u to the set subject to $|dJ/du| < |\dot{\text{ref}}(t)|/k_1$, which can be made arbitrarily small by proper tuning of k_1 .

Another well-known extremum seeking approach is the so-called perturbation-based ESC. It uses a perturbation signal (often sinusoidal) to explore the space of control and steers the control variable toward its local optimum by implicitly following a gradient update. This type of ESC algorithm has been thoroughly analyzed, for example, in [2]. Let us show a simplified version of a sinusoidal disturbance-based ESC algorithm.

$$\begin{aligned} \dot{g} &= a \sin \left(\omega t + \frac{\pi}{2} \right) J(u) \\ u &= g + a \sin \left(\omega t - \frac{\pi}{2} \right), \quad a > 0, \omega > 0. \end{aligned} \quad (\text{S9})$$

It has been shown, using averaging theory and singular perturbation theory, that this simple algorithm, under some simple assumptions (of at least local optimality and the smoothness of J), can (locally) converge to a neighborhood of the optimal control u^* , for example, [2]. Several other ESC algorithms exist; however, it is not the purpose of this article to review all the extremum seeking results. Instead, we refer the interested reader to the recent survey [49].

introduction) to solve the IFT problem for complex industrial systems.

ESC can be classified as model-free adaptive control in the sense that it does not rely on any mathematical model of the system. ESC algorithms are solely based on online measurements collected directly from the system. The term "adaptive" here means that the controller can adapt to and

deal with any uncertainty in the system since the controller does not rely on any specific model.

Recently, ESC has been incorporated into the design of various systems, such as electromagnetic actuators for high-rise elevators and 5G digital power amplifiers (DPAs). Some of the highlights of these results are presented in the following sections.

Electromagnetic Actuators: A Brief Overview

Electromagnetic actuators are devices that convert electrical energy into mechanical motion using the principles of electromagnetism [S1]. These actuators typically consist of a coil of wire that generates a magnetic field when an electric current passes through it and a movable ferromagnetic core, or armature, that responds to this magnetic field. The interaction between the magnetic field and the core produces force or motion, which can be harnessed to perform various mechanical tasks.

HOW ELECTROMAGNETIC ACTUATORS WORK

The basic operation of an electromagnetic actuator involves the generation of a magnetic field by the coil, which in turn attracts or repels the ferromagnetic core. By controlling the electric current through the coil, the magnetic field strength can be varied, allowing precise control over the movement of the core. This movement can be linear, as in solenoids, or rotary, as in electric motors. The force generated by the actuator is proportional to the current, making it a highly controllable and efficient means of producing motion.

APPLICATIONS OF ELECTROMAGNETIC ACTUATORS

Electromagnetic actuators are widely used in various applications across multiple industries. In the automotive industry, they are employed in systems like fuel injectors and brake actuators, where precise control over mechanical movement is essential. In industrial automation, electromagnetic actuators are used in robotic arms and conveyor systems to perform tasks such as lifting, positioning, and sorting.

One notable application is in the field of elevator systems, particularly in high-rise buildings. Here, electromagnetic actuators are integral to the operation of electromagnetic brakes, which are crucial to ensuring the safety and smooth operation of elevators. These brakes rely on electromagnetic forces to engage or release the brake pads, providing reliable control over the elevator's motion.

REFERENCE

[S1] J.-R. Brauer, *Magnetic Actuators and Sensors*, 2nd ed. Hoboken, NJ, USA: Wiley, 2006.

ELECTROMAGNETIC ACTUATORS FOR HIGH-RISE BUILDINGS

Electromagnetic actuators operate on the fundamental principles of electromagnetism, converting electrical energy into mechanical motion (refer to “[Electromagnetic Actuators: A Brief Overview](#)”). Their ability to deliver precise, rapid, and controlled movements makes them ideal for applications where reliability and safety are paramount. In high-rise buildings, where elevator systems must operate continuously and often under significant loads, the advantages of electromagnetic actuators are particularly pronounced.

Electromagnetic actuators provide exceptional precision in controlling the braking mechanism of elevators. In high-rise buildings, where elevators travel at high speeds and over long distances, the ability to precisely control the deceleration and stopping of the elevator car is crucial for safety and passenger comfort. Electromagnetic actuators allow for smooth and controlled braking, reducing the likelihood of abrupt stops that could cause discomfort or injury to passengers.

More specifically, noise is a significant concern in high-rise buildings, particularly in residential or office environments, where quiet operation is desired. Electromagnetic actuators contribute to noise reduction in elevator systems by providing smooth and quiet operation. The absence of fluid movement or mechanical friction, as seen in hydraulic or mechanical systems, means that electromagnetic brakes generate less noise during engagement and disengagement. This is particularly important in

high-rise buildings where multiple elevators operate simultaneously.

Electromagnetic actuators allow for precise control of a moving armature between two target positions. With regard to noise management, the primary objective, known as the “soft landing” of the moving armature, is to ensure minimal contact velocity between the armature and the fixed parts of the actuator (the left and right sides of the electromagnet in [Figure 1](#)). This motion is typically iterative, as the actuator must repeatedly open and close to achieve the desired cyclic motion of a mechanical part attached to it (the “actuated object” in [Figure 1](#)). Such actuators are commonly used in elevator brakes, where the moving mechanical part contacts the elevator sheave's




FIGURE 1 An electromagnetic actuator.

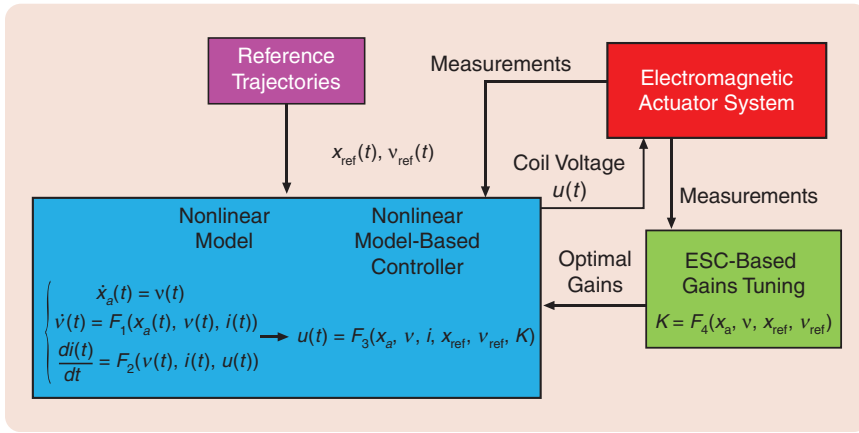


FIGURE 2 An elevator brake: the control loop with ESC gain tuning. Here, F_i represents functional relationships among variables, where detailed mathematical expression is omitted for simplicity.

drum to halt its rotation and consequently stop the vertical motion of the elevator car. Soft landing performance is particularly desirable in this application because it is associated with lower noise levels, which is crucial for deployment in high-rise residential buildings.

To achieve robust soft landing performance that can withstand system aging, we propose the control loop in Figure 2. This ESC-based control architecture can be decomposed into three main components. First, there is a system modeling phase coupled with a reference trajectory generator that produces desired armature trajectories for the system to follow. Second, a model-based controller is designed to track these reference trajectories through a feedback policy that modulates the coil voltage based on real-time measurements, including armature position and velocity. Third, an adaptive ESC-based gain tuning mechanism continuously optimizes the feedback policy's parameters to compensate for system aging. The following sections provide detailed descriptions of each component.

System Modeling

We consider the following nonlinear model for electromagnetic actuators:

$$m \frac{d^2 x_a}{dt^2} = k(x_0 - x_a) - \eta \frac{dx_a}{dt} - \frac{ai^2}{2(b + x_a)^2}$$

$$u = Ri + \frac{a}{b + x_a} \frac{di}{dt} - \frac{ai}{(b + x_a)^2} \frac{dx_a}{dt}, \quad 0 \leq x_a \leq x_f \quad (1)$$

where x_a represents the armature position physically constrained between the initial position of the armature 0 and the maximal position of the armature x_f , dx_a/dt represents the armature velocity, m is the armature mass, k is the spring constant, x_0 is the initial spring length, η is the damping coefficient (assumed to be constant), $ai^2/(2(b + x_a)^2)$ represents the electromagnetic force (EMF) generated by the coil, a and b are two constant parameters

of the coil, R is the resistance of the coil, $L = a/(b + x_a)$ is the coil inductance, and $(ai/(b + x_a)^2)(dx_a/dt)$ represents the back EMF. Finally, i denotes the coil current, and di/dt is its time derivative, and u represents the control voltage applied to the coil. In this model, we do not consider the saturation region of the flux linkage in the magnetic field generated by the coil since we assume that a current and armature motion ranges within the linear region of the flux.

We define x_{ref} as a desired armature position trajectory such that x_{ref} is a smooth (at least C^2) function satisfying the initial/final constraints:

$$x_{ref}(0) = 0, x_{ref}(t_f) = x_f, \quad \dot{x}_{ref}(0) = 0,$$

$\dot{x}_{ref}(t_f) = 0$, where t_f is the desired finite motion time and x_f is a desired final position, which is considered equal to the maximal position of the armature.

We consider the dynamical system (1) with bounded parametric uncertainties in the spring coefficient δk , with $|\delta k| \leq \delta k_{max}$, and the damping coefficient $\delta \eta$, with $|\delta \eta| \leq \delta \eta_{max}$, such that $k = k_{nominal} + \delta k$ and $\eta = \eta_{nominal} + \delta \eta$, where $k_{nominal}$ and $\eta_{nominal}$ are the nominal values of the spring stiffness and the damping coefficient, respectively.

Next, following the input–output linearization and Lyapunov reconstruction methods, for example, [21], we can write the following robust passive controller:

$$u = -\frac{m(b + x_a)}{i} \left(v_s + \frac{k_{nominal}}{m} \dot{x}_a + \frac{\eta_{nominal}}{m} \ddot{x}_a - \frac{Ri^2}{(b + x_a)m} \right)$$

$$+ \frac{m(b + x_a)}{i} \frac{\partial V}{\partial z_3} k \left(\frac{\delta k_{max}}{m} |\dot{x}_a| + \frac{\delta \eta_{max}}{m} |\ddot{x}_a| \right)$$

$$v_s = x_{ref}^{(3)}(t) + K_3 z_3(t) + K_2 z_2(t) + K_1 z_1(t)$$

$$k > 0, K_i < 0, i = 1, 2, 3 \quad (2)$$

where $\mathbf{z} := (z_1, z_2, z_3)^T = (x_a - x_{ref}, \dot{x}_a - \dot{x}_{ref}, \ddot{x}_a - \ddot{x}_{ref})^T$, with $\dot{x}_{ref} = (dx_{ref}(t)/dt)$ and $\ddot{x}_{ref} = (d^2 x_{ref}(t)/dt^2)$, and $V = \mathbf{z}^T P \mathbf{z}$, $P > 0$ is a solution of equation $P\tilde{A} + \tilde{A}^T P = -I$, with

$$\tilde{A} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ K_1 & K_2 & K_3 \end{pmatrix} \quad (3)$$

where K_1, K_2 , and K_3 are chosen such that \tilde{A} is Hurwitz.

IFT of the Controller Gains Using ESC

Considering a cyclic behavior of the actuator such that each iteration happens over a time interval of length t_f , we define the tuning cost function, to be minimized with respect to β , as

$$J(\beta) = C_1 z_1(It_f)^2 + C_2 z_2(It_f)^2 \quad (4)$$

where $I = 1, 2, 3, \dots$ is the number of iterations, $C_1, C_2 > 0$, and the optimization variable is $\beta = [\delta K_1, \delta K_2, \delta K_3, \delta k]^T$, so that the feedback gains are written as

$$\begin{aligned} K_1 &= K_{1\text{nominal}} + \delta K_1 \\ K_2 &= K_{2\text{nominal}} + \delta K_2 \\ K_3 &= K_{3\text{nominal}} + \delta K_3 \\ k &= k_{\text{nominal}} + \delta k \end{aligned} \quad (5)$$

where $K_{1\text{nominal}}, K_{2\text{nominal}}, K_{3\text{nominal}}$ and k_{nominal} are the nominal initial values of the feedback gains in (2).

Following a dither-based ESC approach, such as (S9), the variations of the estimated gains are given by

$$\begin{aligned} \dot{x}_{K_1} &= a_{K_1} \sin\left(\omega_1 t - \frac{\pi}{2}\right) J(z(\beta)) \\ \delta \hat{K}_1(t) &= x_{K_1}(t) + a_{K_1} \sin\left(\omega_1 t + \frac{\pi}{2}\right) \\ \dot{x}_{K_2} &= a_{K_2} \sin\left(\omega_2 t - \frac{\pi}{2}\right) J(z(\beta)) \\ \delta \hat{K}_2(t) &= x_{K_2}(t) + a_{K_2} \sin\left(\omega_2 t + \frac{\pi}{2}\right) \\ \dot{x}_{K_3} &= a_{K_3} \sin\left(\omega_3 t - \frac{\pi}{2}\right) J(z(\beta)) \\ \delta \hat{K}_3(t) &= x_{K_3}(t) + a_{K_3} \sin\left(\omega_3 t + \frac{\pi}{2}\right) \\ \dot{x}_k &= a_k \sin\left(\omega_4 t - \frac{\pi}{2}\right) J(z(\beta)) \\ \delta \hat{k}(t) &= x_k(t) + a_k \sin\left(\omega_4 t + \frac{\pi}{2}\right) \\ \delta K_j(t) &= \delta \hat{K}_j((I-1)t_f), (I-1)t_f \leq t < It_f, j \in \{1, 2, 3\} \\ \delta k(t) &= \delta \hat{k}((I-1)t_f), (I-1)t_f \leq t < It_f, I = 1, 2, 3, \dots \end{aligned} \quad (6)$$

where $a_{K_1}, a_{K_2}, a_{K_3}$, and a_k are positive and $\omega_p + \omega_q \neq \omega_r$, $p, q, r \in \{1, 2, 3, 4\}$, for $p \neq q \neq r$.

The system of equations (6) may appear complex at first glance; however, it is fundamentally a parallel implementation of the basic ESC equations in (S9), replicated for each of the four gains that require tuning in this application. The final two equations in (6) serve a specific purpose: they discretize the ESC outputs and maintain them constant throughout one learning iteration, which corresponds to a complete armature motion trajectory over a time interval t_f .

We successfully implemented this controller on an elevator brake test bed. As illustrated in Figure 3, during the control-free phase (when the drum is in free fall), the noise monitor displays several impact noise peaks. However, after activating the ESC-based control tuner (either using a laser sensor or an estimator to obtain the displacement measurement x_a), the noise level is reduced to negligible levels. This reduction in the noise level is directly linked to improved control of the moving drum's trajectory, which closely follows a desired smooth landing path, as in Figure 4.

This is an example in which ESC was successfully implemented to enhance the real-time performance of a mechatronics system. However, ESC has also been applied to various other real-world applications, including purely electronic systems, such as DPAs. The following section explores this particular application in more detail.

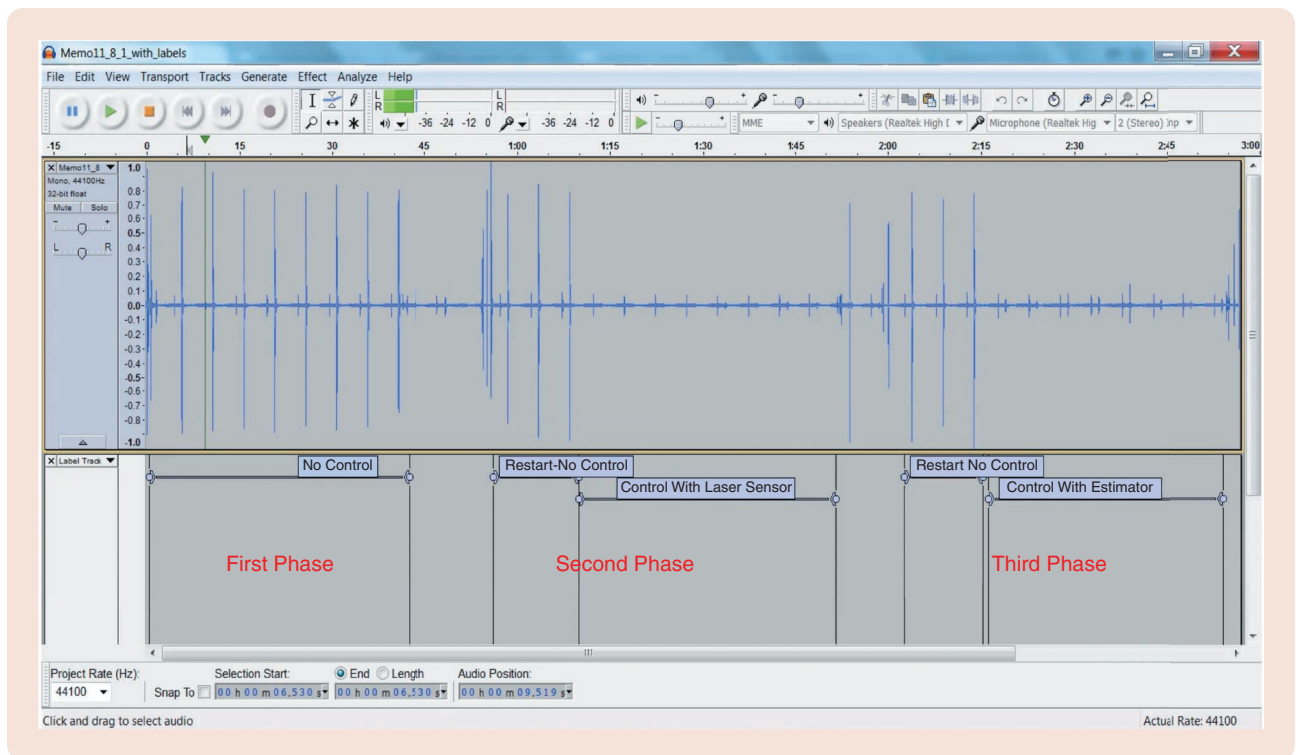


FIGURE 3 An elevator brake: noise level reduction. In the first phase, there is no control. The second phase restarts with no control and then switches to control with armature position measurement. The third phase restarts with no control and then switches to control with armature position estimation from coil current measurement.

DPAS FOR 5G NETWORKS

The rapid evolution of wireless communication systems has been driven by the increasing demand for higher data rates, lower latency, and improved connectivity. 5G networks are at the forefront of this technological advancement, promising to deliver unprecedented performance across a wide range of applications, from enhanced mobile broadband to ultrareliable low-latency communication and massive machine-type communication. Central to the successful deployment of 5G networks is the development of efficient and high-performance PAs, which play a critical role in ensuring the reliability and quality of transmitted signals.

In traditional analog PAs, efficiency and linearity are often at odds, which poses significant challenges in the design of amplifiers that can meet the stringent requirements of 5G. The increasing use of higher frequency bands, such as millimeter waves, and complex modulation schemes like quadrature amplitude modulation further complicate PA design. These factors necessitate a shift toward DPAs, which offer the potential to overcome the limitations of analog designs through digital processing techniques.

DPAs leverage digital signal processing to improve linearity and efficiency while also reducing power consumption (see “DPAs: A Brief Overview”). By converting the input signal into a digital format, DPAs can more precisely control the amplification process, enabling advanced techniques, such as digital predistortion (DPD) and envelope tracking. These techniques are crucial to maintaining signal integrity in the high-frequency high-bandwidth environment of 5G networks. More specifically, DPD is a highly effective linearization technique that is used to counteract nonlinearities in PAs [22]. DPD works by applying a precorrection to the input signal, ensuring that the combined system of the DPD and PA operates as an ideal linear and memoryless amplifier.

The adoption of DPAs in 5G networks also aligns with the broader trend of software-defined radio, where the flexibility of digital components allows for rapid adaptation to changing standards and requirements. This adaptability is particularly important in the context of 5G, where the network must support a wide range of use cases, each with its own unique performance criteria.




FIGURE 4 An elevator brake: the linear trajectory of the moving drum.

In the literature, advanced PA architectures based on dynamic load or supply modulation have been proposed to avoid wasting power resources [23]. Some of the most popular solutions are Doherty [24], envelope tracking [25], Chireix [26], and outphasing [27]. These highly efficient topologies require linearization techniques, such as DPD, to meet the linearity requirement, especially with an increase of the signal bandwidth.

Amplification architectures based on active load modulation, such as the Doherty amplifier, are among the most widely used techniques to enhance PA efficiency. These architectures exploit the nonlinear interaction between the main and auxiliary transistors, particularly for modulated signals with a large dynamic range. Although these designs can be implemented with a single radio-frequency (RF) input for use in transmitters, several studies have highlighted the advantages of maintaining separate inputs [28], [29], [30]. Additionally, the benefits of dual-input PAs over single-input configurations have been explored in [31].

Separation of the RF input introduces additional degrees of freedom, often referred to as “free optimization parameters,” which can be adjusted to optimize performance and enhance PA efficiency [32]. These free parameters encompass critical circuit- and system-level factors, such as PA bias voltages, the power ratio, and phase shift of individual input signals.

When focusing on dual-input Doherty PAs (DIDPAs), identifying the optimal values of these free parameters that ensure high performance typically necessitates experimental cross validation or exhaustive search. These processes can be costly and computationally intensive, particularly when dealing with an unconstrained search space. Optimizing these free parameters within a defined range can be framed as a global optimization problem. Various techniques have been proposed in the literature to determine the optimal set of free parameters within large search intervals [33].

Generally, a DIDPA operates at maximum efficiency, which often results in compromised linearity. To address this limitation, an effective system approach is needed to balance the tradeoff between linearity and efficiency. Designing a DPD model that can accurately determine the inverse characteristics of a DIDPA is challenging, especially when it comes to identifying a model that is both suitable for hardware implementation and possesses favorable numerical properties along with high modeling accuracy.

A practical approach to optimizing a DIDPA involves integrating it into an iterative system-based process that focuses on two key aspects: optimizing the free parameters enabled by the separation of the input signals and compensating for nonlinearities through DPD linearization.

The first significant work on online learning-based optimization of DIDPAs was introduced in [34]. The authors proposed an adaptive technique utilizing a simultaneously

DPAs: A Brief Overview

DPAs represent a significant advancement in the field of electronic amplification, especially in the context of modern communication systems [S2], [S3]. Unlike traditional analog PAs, which amplify continuous signals, DPAs amplify digital signals—discrete-time signals that represent data in the form of bits (zeros and ones). This digital approach allows greater efficiency, flexibility, and integration with digital processing systems, making DPAs a crucial component in modern high-frequency applications, such as wireless communication and broadcasting.

HOW DPAS WORK

A DPA converts a low-power digital signal into a higher-power signal suitable for driving antennas or other loads. The core of a DPA typically involves pulsewidth modulation or delta-sigma modulation, where the digital input signal is used to control the switching of transistors within the amplifier. These transistors rapidly switch between the on and off states, creating a high-power output that mirrors the input signal but with much greater amplitude.

The process involves several key stages.

- **Digital input processing:** The input digital signal is processed to generate the necessary control signals for the transistors.
- **Switching stage:** High-speed transistors switch on and off, controlled by the processed digital input. This stage is responsible for generating the amplified output signal.
- **Filtering:** The output signal, initially a high-frequency high-power signal with switching noise, is passed through filters to remove unwanted noise and smooth the signal into a usable form.
- **Output stage:** The filtered signal is delivered to the load, such as an antenna or a speaker, with the necessary power level for effective operation.

ADVANTAGES OF DPAS

DPAs offer several advantages over traditional analog PAs.

- **Efficiency:** By operating transistors in switching mode (fully on or fully off), DPAs minimize power loss in the form of heat, leading to much higher efficiency.
- **Linearity:** Through advanced digital techniques, such as digital predistortion, DPAs can achieve high linearity, which is essential for minimizing distortion in complex modulation schemes used in modern communications.
- **Integration:** DPAs can be easily integrated with other digital systems, allowing for compact designs and seamless operation with digital signal processors and other digital components.
- **Flexibility:** The digital nature of DPAs allows for dynamic adaptation and real-time tuning, making them suitable for applications requiring variable output characteristics.

APPLICATIONS OF DPAS

DPAs are increasingly being adopted in various high-frequency applications where efficiency, linearity, and integration are paramount. Some of the key applications include

- **Wireless communication:** DPAs are essential in modern wireless communication systems, including 4G and 5G networks. They are used in base stations and mobile devices to amplify signals before transmission, ensuring that data are transmitted reliably over long distances. The high efficiency and linearity of DPAs help to reduce power consumption and minimize signal distortion, which is critical to maintaining the integrity of transmitted data.
- **Broadcasting:** In TV and radio broadcasting, DPAs are used to amplify the signal before it is broadcast over large areas. The ability to efficiently handle high power levels while maintaining signal quality makes DPAs ideal for this purpose.
- **Satellite communication:** DPAs are used in satellite communication systems to amplify signals transmitted to and from satellites. The high efficiency of DPAs is particularly beneficial in satellite applications, where power resources are limited.
- **Radio-frequency identification systems:** Radio-frequency identification (RFID) systems rely on DPAs to power the signals used to communicate with RFID tags. The precision and efficiency of DPAs enhance the performance of RFID systems in various applications, including inventory management and access control.
- **Medical devices:** DPAs are also found in medical devices, such as magnetic resonance imaging machines and other diagnostic equipment that requires high-power RF signals. The precision and efficiency of DPAs contribute to the accuracy and effectiveness of these devices.

CONCLUSION

DPAs represent a significant technological advancement in the amplification of high-frequency signals. Their efficiency, linearity, and ease of integration with digital systems make them indispensable in a wide range of applications, from wireless communication and broadcasting to medical devices. As digital technology continues to evolve, the role of DPAs is expected to expand, driving further innovation in various fields.

REFERENCES

- [S2] V. Pooria, S. Mohammady, B. Mohd Ali, and N. B. Sulaiman, *Power Efficiency in Broadband Wireless Communications*, 1st ed. Boca Raton, FL, USA: CRC Press, 2015.
- [S3] X. Luo, H. J. Qian, Y. Yin, and H. Xu, "Empowering multifunction: Digital power amplifiers, the last RF frontier of the analog and digital kingdoms," *IEEE Microw. Mag.*, vol. 21, no. 12, pp. 47–67, Dec. 2020, doi: [10.1109/MMM.2020.3023221](https://doi.org/10.1109/MMM.2020.3023221).

perturbed stochastic approximation algorithm to tune free parameters, such as the power ratio, phase shift, and bias voltages. A cost function, defined in terms of output gain (G) and power-added efficiency (PAE), was employed to guide the algorithm's convergence.

Building on the work in [34], an extension was presented in [35], where the cost function was expanded to include additional metrics, specifically, the output power (P_{out}) and adjacent channel power ratio (ACPR), to better address linearity considerations. The optimization process in [35] and [36] utilized a global optimization algorithm, such as simulated annealing (SA), as a first step. After achieving an optimal cost function value, a fine-tuning process was implemented using ESC. The results presented below are from [35] and [36].

DIDPA

A dual-input PA, such as Doherty, is depicted in Figure 5.

The DIDPA has two RF inputs, a drain bias V_{dc} and two gate source voltages $V_{GS,1}$ and $V_{GS,2}$ to control the transistor's terminal independently.

The instantaneous amplitude and phase of each baseband input, along with the V_{GS} gate bias voltages, can be independently controlled and adjusted. This provides a significant degree of freedom to optimize these parameters




FIGURE 5 A dual-input PA.

TABLE 1 Parameters controlled by ESC.

Low-Level Signals	Symbol	Unit
Threshold of peak-to-average power ratio reduction	μ	Decibels
Power ratio	α	Percent
Phase shift	ϕ	Degrees
Attenuation difference	ψ	Decibels
Main bias voltage	$V_{GS,m}$	Volts
Peaking bias voltage	$V_{GS,p}$	Volts

for improved power efficiency. These parameters are tuned through ESC, which determines their optimal values in real time, ensuring that the DIDPA operates at a point that offers a better tradeoff between efficiency and linearity.

These free parameters are expressed as functions of low-level signals (refer to [36] for more details), which are summarized in Table 1; such low-level signals are then directly adjusted by the ESC controller.

Learning Cost Function

The cost function is an essential design aspect that ensures a good tradeoff between linearity and efficiency. The linearity requirement is presented in terms of two optimization objectives, namely, the error vector magnitude (EVM) and the ACPR, while the efficiency requirement is presented by the PAE and the output power P_{out} .

The EVM is a metric that measures the in-band distortion level. The ACPR is used to evaluate out-band distortions and is defined for the lower (left) and upper (right) adjacent channels. We propose to design the cost function J , to be maximized with respect to β , according to a linear weighted sum, which is defined as

$$J(\beta) = w_1 J_{EVM} + w_2 J_{ACPR} + w_3 J_{PAE} + w_4 J_{P_{out}}, \quad \sum_{i=1}^4 w_i = 1 \quad (7)$$

with

$$\begin{cases} J_{EVM} = \left| \frac{EVM}{EVM_t} \right| \\ J_{ACPR} = \left| \frac{ACPR}{ACPR_t} \right| \\ J_{PAE} = \left| \frac{PAE}{PAE_t} \right| \\ J_{P_{out}} = \left| \frac{P_{out}}{P_{out,t}} \right| \\ \beta = [\mu, \alpha, \phi, \psi, V_{GS,m}, V_{GS,p}]^T \end{cases} \quad (8)$$

where EVM_t , $ACPR_t$, PAE_t , and $P_{out,t}$ are the EVM target, ACPR target, PAE target, and P_{out} target, respectively. The details of how these optimization cost terms are defined as a function of the control vector β are beyond the scope of this review article and can be found in [36].

IFT Using ESC With Warm Start

The search for an optimal set of the parameters defined in Table 1 is a nonconvex optimization problem due to the nonlinear mapping between these low-level signals and the final objective function. To help the ESC algorithm reach a good local maximum while keeping the ESC simple and deployable under stringent computation power constraints, we choose to warm start the ESC with a black-box optimization method, namely, the SA approach [37]; see the diagram in Figure 6.

The SA algorithm is a stochastic search based on the Metropolis Monte Carlo method [38], the concept of which

is to accept not only the solutions that improve J but also some solutions that worsen it with a probability p , known as the “Metropolis criterion” and defined as

$$p(\Delta E) = e^{-\frac{\Delta E}{k_{\text{bolt}}\mathcal{T}}} \quad (9)$$

where ΔE is the change in the cost function, k_{bolt} is Boltzmann’s constant, and \mathcal{T} is the control parameter analogous to the temperature of the annealing process.

During the search, the temperature is gradually decreased until it reaches zero value in the perfect case.

Once the SA algorithm has reached a neighborhood of $\beta_{\text{opt,SA}}$, the optimization procedure switches to ESC for fine-tuning. The algorithm of ESC with a warm start is described in Algorithm 1.

Experimental Results

The performance of the ESC autotuner was evaluated on the testbed displayed in Figure 7.

The proposed system approach for the DIDPA is tested using a 64-quadrature amplitude modulation-modulated 20-MHz bandwidth long-term evolution signal with a roll-off factor of 0.6 at $f_c = 3$ GHz, with 8 dB of peak-to-average power ratio. The weighting coefficients of the learning cost are chosen as $w_1 = 0.1$, $w_2 = 0.1$, $w_3 = 0.4$, and $w_4 = 0.4$. The target performance levels are set to $\text{EVM}_t = 3\%$, $\text{ACPR}_t = -50$ dB, $\text{PAE}_t = 60\%$, and $P_{\text{out},t} = 40$ dBm. The ESC warm start algorithm is tuned according to the following coefficients: $\mathcal{T}_0 = 1$, $\mathcal{T}_f = 0.01$, $C = 0.96$, $k_{\text{bolt}} = 1$, $\text{SA}_{\text{max}} = 60$, $\omega = 10$, $K = 1$, and $a = 0.01$. The evolution of the learning cost J is reported in Figure 8. The evolution of the optimization parameters appears in Figure 9. The combination of the SA algorithm with ESC fine-tuning led to an optimal performance of $\text{PAE} = 46.58\%$ and $P_{\text{out}} = 35.63$ dBm (refer to [36] for more details).

We conclude this article by presenting another application where ESC played a crucial role in the real-time autotuning of a complex feedback system. In this case, ESC was used to autotune models of airflow and temperature coupling in indoor environments. This coupling is most accurately described by the well-known Navier–Stokes (NS) partial differential equation (PDE). However, solving this PDE in real time is practically infeasible. Therefore, we developed simplified versions of the NS PDE in the form of reduced-order models (ROMs). While ROMs are easier to implement and solve in real time, they require expert tuning to account for the information lost during the simplification of the NS PDE. In the next section, we demonstrate how ESC can automate this tuning process in real time.

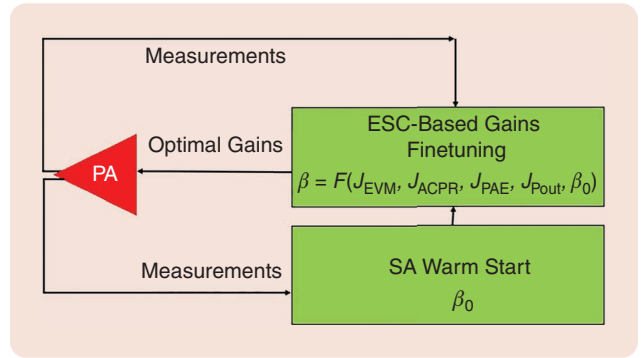


FIGURE 6 A PA: the control loop with SA–ESC gain tuning. Here, F represents a functional relationship among variables, where detailed mathematical expression is omitted for simplicity.

ALGORITHM 1 ESC with a warm start.

Initialization: $\mathcal{T}_0, \mathcal{T}_f, \mathcal{K}, a, \omega$

Run SA algorithm

$\beta_{0,\text{ESC}} = \Theta_{\text{opt,SA}}$

$\beta = \beta_{0,\text{ESC}}$

$J_{\text{old}} = J_{\text{opt,SA}}$

while 1 **do**

 Evaluate $J(\beta)$

if $J(\beta) > J_{\text{old}}$ **then**

$J_{\text{old}} = J(\beta)$

$\xi(\tau) = J(\beta) \times a \sin(\omega\tau)$

$\hat{\beta} = K \times \int_0^t \xi(\tau) d\tau$

$\beta = \hat{\beta} + a \sin(\omega t)$

else

$\beta_{\text{opt}} = \beta$

$J_{\text{opt}} = J(\beta_{\text{opt}})$

end while loop

end

end

Return β_{opt}

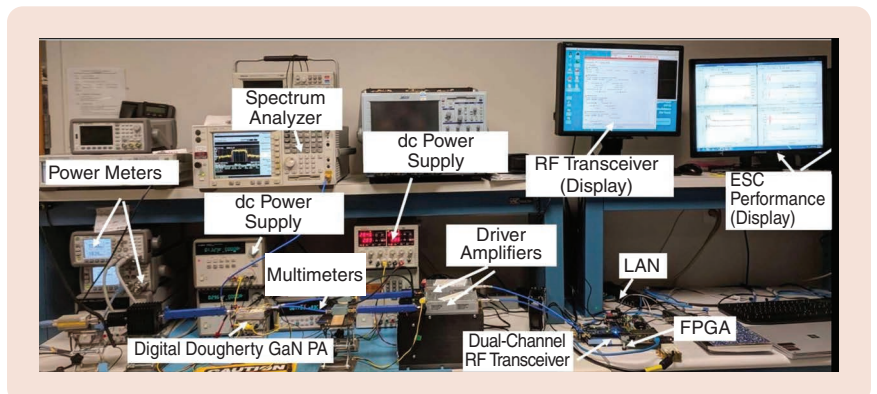


FIGURE 7 A Doherty testbed. LAN: local area network; GaN: gallium nitride; FPGA: field-programmable gate array.

FLUIDIC SPATIOTEMPORAL MODELS FOR HVAC APPLICATIONS

Thermal and fluid systems have been extensively studied across various scientific and engineering disciplines due to their wide range of applications. For example, optimizing energy management in HVAC systems (refer to “HVAC Systems: A Brief Overview”) has become a priority research area in many countries, given the significant impact of these systems on energy consumption [39], particularly in this era of artificial intelligence and large data centers [40].

One of the key challenges in HVAC management is accurately estimating the entire spatial profile of airflow and temperature using a limited number of strategically placed

sensors throughout a building. One approach to addressing this challenge is through the physics-based modeling of airflow dynamics and their interaction with temperature. One of the most accurate models for capturing airflow dynamics is described by NS equations.

State estimation for NS equations has been extensively investigated in the literature over the past decade. In [41], an infinite-dimensional Kalman filter was designed for a linearized NS equation around the velocity field of interest. While the approach demonstrated good performance as an initial contribution to fluid estimation, two major challenges arise in practical implementation. First, the linearization assumption of NS equations can lead to a low-fidelity fluid model. Second, the high computational cost associated with discretizing the infinite-dimensional filter poses significant challenges. To address these issues, the authors of [42] developed a proper orthogonal decomposition (POD)-based model reduction (refer to “POD Model Reduction” for an introduction) for NS equations and designed an extended Kalman filter for state estimation, showing effective results in estimating the velocity profile. Further research on state estimation for the POD-ROM of NS equations can be found in [43] and [44].

Higher-fidelity models of NS equations that account for buoyancy forces, driven by density changes due to temperature dependence, require additional complexity. The dynamics of the temperature profile are also governed by the conservation of energy, resulting in the formulation of Boussinesq equations, which describe coupled thermal and fluid systems. The design and analysis of Boussinesq equations are particularly challenging due to the coupled nonlinearity of the two PDEs, especially in 2D and 3D domains. Moreover, as highlighted in [45], POD-based model reduction can compromise the stability of state variables, often due to the truncation of higher-order modes that serve as stabilizing factors.

A robust and stable model reduction for Boussinesq equations was developed in [46] through the introduction of a new closure model that stabilizes the ROM. The authors demonstrate the robust stability of this closure model in the presence of parameter uncertainties using Lyapunov analysis. Additionally, the gains in the closure terms are autotuned using a learning-based ESC algorithm, which minimizes the errors between the true model solution and the ROM solution. The results presented in [46] show significant improvements in




FIGURE 8 The evolution of cost function J over warm start and ESC iterations.




FIGURE 9 The evolution of control variables over ESC iterations.

HVAC Systems: A Brief Overview

HVAC systems are integral to maintain indoor environmental comfort in residential, commercial, and industrial settings. These systems are designed to regulate temperature, humidity, and air quality within a building, ensuring a comfortable and safe indoor environment. HVAC systems combine three primary functions—heating, ventilation, and air-conditioning—to create and maintain optimal indoor conditions regardless of external weather or environmental factors.

COMPONENTS OF HVAC SYSTEMS

HVAC systems are composed of several key components that work together to achieve the desired indoor climate [S4].

- **Heating system:** The heating component is responsible for maintaining a comfortable indoor temperature during colder weather. Common heating systems include furnaces, boilers, and heat pumps.
- **Ventilation system:** Ventilation is crucial for maintaining indoor air quality by removing contaminants and introducing fresh air into the building. Ventilation systems can be natural or mechanical. Natural ventilation relies on passive airflow through windows, vents, and other openings. Mechanical ventilation involves fans, ducts, and air handling units to actively circulate air within the building. This type of ventilation can be combined with air filters and purifiers to remove pollutants, allergens, and other harmful particles from the air.
- **Air-conditioning system:** The air-conditioning component cools and dehumidifies the indoor air, providing relief during hot weather. The most common air-conditioning systems include central air conditioners that use a compressor, evaporator, and condenser to cool air and distribute it through the building's ductwork and split systems that include an outdoor unit (condenser and compressor) and an indoor unit (evaporator), connected by refrigerant lines. They are often used in residential and small commercial applications.

HOW HVAC SYSTEMS WORK

HVAC systems operate through a combination of heating, cooling, and ventilation processes that are carefully controlled to maintain desired indoor conditions. Here is a simplified overview of how these systems work.

- **Temperature control:** A thermostat monitors the indoor temperature and triggers the heating or cooling system to maintain the set temperature. In heating mode, the furnace or heat pump generates warmth, while in cooling mode, the air conditioner removes heat from the indoor air.
- **Air circulation:** The system's fans and blowers circulate air through ducts, distributing warm or cool air throughout the building. Ventilation fans also introduce fresh air from the outside while exhausting stale air.
- **Humidity regulation:** HVAC systems help regulate indoor humidity levels, particularly through the dehumidifying ef-

fect of air-conditioning. Some systems also include humidifiers to add moisture to the air in dry conditions.

- **Air filtration:** Air filters trap dust, pollen, and other particles, improving indoor air quality. Advanced systems may include air purifiers to remove additional pollutants.

APPLICATIONS OF HVAC SYSTEMS

HVAC systems are ubiquitous across various sectors, providing essential climate control and air quality management in diverse environments. Some of the key applications include

- **Residential buildings:** HVAC systems are standard in homes to maintain comfortable living conditions throughout the year. They ensure warmth in winter, cooling in summer, and proper ventilation year-round.
- **Commercial buildings:** In offices, shopping malls, hotels, and other commercial spaces, HVAC systems are crucial for maintaining a comfortable environment for occupants. They also play a role in ensuring the efficient operation of electronic equipment by preventing overheating.
- **Industrial facilities:** HVAC systems in industrial settings are used not only for worker comfort but also to control the environment in manufacturing processes. For example, clean rooms in semiconductor manufacturing require precise temperature, humidity, and air quality control.
- **Health-care facilities:** Hospitals, clinics, and laboratories rely on advanced HVAC systems to maintain sterile environments, control infection, and ensure the comfort of patients and staff. These systems are critical for air filtration and maintaining specific temperature and humidity conditions.
- **Data centers:** HVAC systems are vital in data centers to prevent servers and other IT equipment from overheating. These systems are designed to maintain a constant temperature and humidity level to ensure the reliability of data processing and storage operations.
- **Transportation:** HVAC systems are used in vehicles, airplanes, trains, and ships to provide passenger comfort and regulate the temperature of sensitive cargo.

CONCLUSION

HVAC systems are an essential part of modern infrastructure, providing critical temperature control, air quality management, and comfort across a wide range of environments. From residential homes to industrial facilities, these systems play a vital role in ensuring the well-being of occupants and the efficient operation of processes. As technology advances, HVAC systems continue to evolve, offering greater efficiency, enhanced control, and improved environmental sustainability.

REFERENCE

[S4] J. W. Mitchell and J. E. Braun, *Principles of Heating, Ventilation, and Air Conditioning in Buildings*, 1st ed. Hoboken, NJ, USA: Wiley, 2012.

One common characteristic of all available feedback controllers is the fact that they all rely on some “well chosen” feedback gains.

solution prediction for laminar flows. However, the method’s effectiveness depends on accurate initial velocity and temperature profiles, a requirement that could be mitigated through state estimation techniques utilizing sparse measurements from a limited number of sensors.

In the following, we summarize the results presented in [47], where the authors proposed a robust ESC-based observer for Boussinesq equations. Consistent with the methodology in [46], a data-driven ESC approach was employed to enable online estimation of model uncertainties.

NS Equations Model

We focus on the dynamics of the velocity field $v(x, t) : \Omega \times \mathbb{R}^+ \rightarrow \mathbb{R}^3$ and the temperature profile $T(x, t) : \Omega \times \mathbb{R}^+ \rightarrow \mathbb{R}$, where x denotes the spatial coordinate $x \in \Omega$ and $t \geq 0$ denotes the time. The spatial domain Ω can be a 2D or 3D space. The governing equations are described by NS equations, with the condition of incompressible flow and the conservation of energy through heat transfer, which leads to the following coupled system:

$$\rho \left(\frac{\partial v}{\partial t} + v \cdot \nabla v \right) = -\nabla p + \nabla \cdot \tau(v) + \rho g \quad (10)$$

$$\nabla \cdot v = 0 \quad (11)$$

$$\rho c_p \left(\frac{\partial T}{\partial t} + v \cdot \nabla T \right) = \nabla \cdot (\kappa \nabla T) \quad (12)$$

where ρ (kg/m³) is the density profile, p [kg/(ms²)] is the pressure field, $\tau(v)$ is the viscous stress, c_p (J/kg^oK) is the constant heat capacity, κ (W/m^oK) is the constant thermal conductivity, and $g = -g e_3$ (m/s²) is the gravitational force. In Boussinesq approximation, the buoyancy force is driven by changes in density $\rho = \rho_0 + \Delta\rho$ from the nominal density ρ_0 , and the density change is modeled as perturbations from the nominal temperature T_0 using the perfect gas law $\Delta\rho g = -\rho_0 \beta (T - T_0) g$, $\beta = 1/T_0$, and the constant term $\rho_0 g$ is absorbed into the pressure. The viscous stress is governed by $\tau(v) = \rho\nu(\nabla v + \nabla v^T)$, with kinematic viscosity ν in square meters per second, whereas the velocity v is in meters per second, and the temperature T is in Kelvin.

By introducing a characteristic length L , characteristic velocity v_0 , and wall temperature T_w , we define the following normalized states:

$$\tilde{x} = \frac{x}{L}, \quad \tilde{t} = \frac{t v_0}{L}, \quad \tilde{v} = \frac{v}{v_0} \quad (13)$$

$$\tilde{p} = \frac{p}{\rho v_0^2}, \quad \tilde{T} = \frac{T - T_0}{T_w - T_0}. \quad (14)$$

Using these variables, PDEs (10)–(12) can be reduced to the following (where we dropped the tilde notation):

$$\frac{\partial v}{\partial t} + v \cdot \nabla v = -\nabla p + \nabla \cdot \tau(v) + \frac{Gr}{Re^2} T e_3 \quad (15)$$

$$\nabla \cdot v = 0 \quad (16)$$

$$\frac{\partial T}{\partial t} + v \cdot \nabla T = \nabla \cdot \left(\frac{1}{RePr} \nabla T \right) \quad (17)$$

where we defined the Reynolds number $Re = v_0 L / \nu$, Grashof number $Gr = (g\beta(T_w - T_0)L^3) / \nu^2$, and Prandtl number $Pr = \nu / (k / \rho_0 c_p)$.

By defining the state vector $q(t) = [q_1(t), q_2(t), \dots, q_{r_v+r_T}(t)]^T$, where q_i is defined in [“POD Model Reduction,” (S12)], and using the Galerkin projection method, we obtain the following ODE (see [46] for the detailed derivation):

$$\dot{q}(t) = \mu D q(t) + [C q(t)] q(t) + b \quad (18)$$

where $\mu > 0$ is the viscosity $\mu = 1/Re$, D is a negative definite diffusion matrix with diagonal blocks corresponding to viscous stress and thermal diffusion, and C is a 3D tensor corresponding to the convection terms in (15) and (17). In the remainder of this article, we call the model (18) “ROM–Galerkin” (“ROM–G”).

To solve the POD–ROM in (18), the values of the viscosity parameter μ and the initial conditions $q(0)$ are required. However, in most applications, these values are uncertain. As an alternative, thermal and velocity sensors can be deployed to measure the partial states of the system, which can then be formulated as a linear map from the POD states.

$$y(t) = \tilde{H} q(t) \quad (19)$$

where $\tilde{H} \in \mathbb{R}^{m \times n}$ is a measurement matrix given by the sensor placement.

To integrate the model (18) with the acquired sensor data, a state estimator is designed to reconstruct the entire state $q(t)$ from the measured data. This reconstruction allows for the estimation of the velocity field $v^{\text{pod}}(x, t)$ and the temperature profile $T^{\text{pod}}(x, t)$ across all x . A well-known design for this purpose is the Luenberger-like observer, which is constructed as a copy of the plant with the addition of measurement error states, and it is expressed as follows:

$$\dot{\hat{q}}(t) = \hat{\mu} D \hat{q}(t) + [C \hat{q}(t)] \hat{q}(t) + b + L(y(t) - \tilde{H} \hat{q}(t)) \quad (20)$$

where $\hat{q}(t)$ is the estimation of states $q(t)$, $\hat{\mu}$ is the estimated value of μ , and $L \in \mathbb{R}^{n \times m}$ is the observer gain to be determined. For analysis purposes, we impose that the lower bound of the viscosity μ is known; that is, $\mu \leq \mu$.

Input State Stability

We introduce some definitions that are used in the sequel, for example, [21]. Consider the system

$$\dot{x} = f(t, x, u) \quad (\text{S10})$$

where $f: [0, \infty) \times \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is piecewise continuous in t and locally Lipschitz in x and u , uniformly in t . The input $u(t)$ is a piecewise continuous bounded function of t for all $t \geq 0$.

A continuous function $\alpha: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is called a “class \mathcal{K} function” if it is strictly increasing and $\alpha(0) = 0$. It is a class \mathcal{K}_∞ function if it is class \mathcal{K} and $\lim_{s \rightarrow \infty} \alpha(s) = \infty$. A continuous function $\beta: \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is a \mathcal{KL} function if, for each fixed $t \geq 0$, the function $\beta(s, t)$ is a \mathcal{K} function and if, for each fixed $s \geq 0$, the function $\beta(s, t)$ is decreasing and such that $\lim_{t \rightarrow \infty} \beta(s, t) = 0$.

Definition 1 ([S5], [21])

The system (S10) is said to be *input-to-state stable (ISS)* if there exist a class \mathcal{KL} function β and a class \mathcal{K} function γ

such that for any initial state $x(t_0)$ and any bounded input $u(t)$, the solution $x(t)$ exists for all $t \geq t_0$ and satisfies

$$\|x(t)\| \leq \beta(\|x(t_0)\|, t - t_0) + \gamma\left(\sup_{t_0 \leq \tau \leq t} \|u(\tau)\|\right).$$

Theorem 1 ([S5], [21])

Let $V: [0, \infty) \times \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable function such that

$$\alpha_1(\|x\|) \leq V(t, x) \leq \alpha_2(\|x\|)$$

$$\frac{\partial V}{\partial t} + \frac{\partial V}{\partial x} f(t, x, u) \leq -W(x), \quad \forall \|x\| \geq \rho(\|u\|) > 0 \quad (\text{S11})$$

for all $(t, x, u) \in [0, \infty) \times \mathbb{R}^n \times \mathbb{R}^m$, where α_1 and α_2 are class \mathcal{K}_∞ functions, ρ is a class \mathcal{K} function, and $W(x)$ is a continuous positive definite function on \mathbb{R}^n . Then, the system (S10) is ISS.

REFERENCE

[S5] M. Malisoff and F. Mazenc, “Further remarks on strict input-to-state stable Lyapunov functions for time-varying systems,” *Automatica*, vol. 41, no. 11, pp. 1973–1978, Nov. 2005, doi: 10.1016/j.automatica.2005.05.015.

Furthermore, our estimated value of the viscosity also satisfies $\mu \leq \hat{\mu}$. Let $\tilde{q}(t)$ be the estimation error state defined by $\tilde{q}(t) := q(t) - \hat{q}(t)$. The ideal performance of the observer is characterized by some kind of stability property of the estimation error $\tilde{q}(t)$. Subtraction of the estimator (20) from the system (18) yields the error dynamics

$$\begin{aligned} \dot{\tilde{q}}(t) &= (\mu D - L\tilde{H})\tilde{q}(t) + \mu_+ D\tilde{q}(t) + \tilde{\mu} Dq(t) \\ &\quad + [C\tilde{q}(t)]q(t) + [Cq(t)]\tilde{q}(t) - [C\tilde{q}(t)]\tilde{q}(t) \end{aligned} \quad (21)$$

where $\mu_+ := \mu - \hat{\mu} > 0$ and $\tilde{\mu} = \mu - \hat{\mu}$. The stability of the observer was analyzed in [47], where it was shown that under the proper choice of the observer matrix L , this observer leads to an input-to-state stable result (see “Input State Stability”) between the system parameters’ estimation error and the state estimation error.

Parametric Uncertainty Estimation by ESC

To enhance the performance of the state observer (20) based on the measured value, the uncertain parameters estimate $\hat{\mu} \in \mathbb{R}$ in (20) for the ROM-G used in the observer should be corrected through learning iterations; see the diagram in Figure 10.

The learning cost function to be minimized with respect to $\hat{\mu}$ at every iteration is an error between the measured value and its estimate, which is formulated as

$$Q(\hat{\mu}) = \int_0^{t_f} (y(t) - \tilde{H}\hat{q}(t))^T R (y(t) - \tilde{H}\hat{q}(t)) dt, \quad R > 0. \quad (22)$$

Since the estimate of the parameter $\hat{\mu}$ affects the observer state \hat{q} , the cost function Q is an implicit function of $\hat{\mu}$.

Following [48], we impose that the cost function $Q(\cdot)$ in (22) has a local minimum at $\hat{\mu} = \mu$ and is analytic and that its variation with respect to $\hat{\mu}$ is bounded in the neighborhood of μ ; that is, $\|\nabla_{\hat{\mu}} Q(\delta\mu)\| \leq \xi_2$, $\xi_2 > 0$, for all $\delta\mu \in \mathcal{N}(\mu)$, where $\mathcal{N}(\mu)$ denotes a compact neighborhood of μ .

We apply an extremum seeking algorithm as a learning method for parameter identification. Let $\hat{\mu}^{(i)}$ be the estimate of the parameter μ at the i th iteration, and we introduce an internal variable $z^{(i)}$ with initial value $z^{(1)} = \hat{\mu}^{(1)}$. The parameter update through the ESC algorithm is given by

$$z_1^{(i+1)} = z_1^{(i)} + a_1 \delta \sin\left(\omega_1 i \delta + \frac{\pi}{2}\right) Q(\hat{\mu}^{(i)}) \quad (23)$$

$$\hat{\mu}^{(i+1)} = z_1^{(i+1)} + a_1 \sin\left(\omega_1 i \delta - \frac{\pi}{2}\right) \quad (24)$$

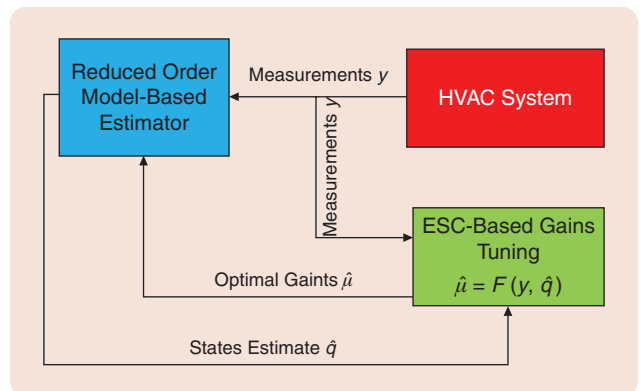


FIGURE 10 HVAC: the estimation loop with ESC gain tuning. Here, F represents a functional relationship among variables, where detailed mathematical expression is omitted for simplicity.

where the tuning parameters are (a_1, ω_1, δ) , which are the amplitude, frequency, and iteration increment $\delta > 0$, respectively. Owing to the convergence analysis of the extremum seeking algorithm as presented, for example, in [2] and [48], the performance of the learning-based observer has been quantified in [47].

Numerical Simulation for 2D Boussinesq Equations

Input Parameters

We implement the numerical simulation by setting the spatial domain as a 2D rectangular shape; namely, $\Omega = (0, 8) \times (0, 1)$. The nondimensional parameters are chosen as $\text{Re} = 10^4$, $\text{Pr} = 1$, and $\text{Gr} = 4 \times 10^8$. The experiment is set up with two fluids at different temperatures, separated by a vertical barrier located at $x = 4$. On the right side of the barrier, the fluid has a lower temperature of one, while on the left side, the fluid has a higher temperature of 1.5. Once the barrier is removed, we anticipate that the warmer and less dense fluid will rise, while the cooler and denser fluid will sink. Initially, the flow is stationary, with the vorticity and stream function set to zero at $t = 0$, and the temperatures on either side of the barrier remain constant.

The 2D equations were simulated using a vorticity stream function formulation with no-slip boundary conditions for vorticity and adiabatic boundary conditions for temperature. We use 10 POD basis functions for vorticity and 10 POD basis functions for temperature variables

computed from snapshots taken every 2×10^{-2} s. The simulation time is set as $t_f = 10$ s. In Figure 11, we present the direct numerical simulation of (15)–(17). The snapshot in Figure 11 corresponds to the spatial view of the flow at $t = 8$ s. A mixing process becomes visible after the barrier between the two fluids is removed, which presents a challenging and complex example of fluid dynamics. The interaction between the fluids during mixing introduces additional complexity, particularly when predicting and analyzing the resulting flow patterns and behaviors.

For data collection, we deploy five sensors to measure temperature and five sensors to measure velocity, with their positions strategically selected using the Q -discrete empirical interpolation method, as described in [50] and examined in Figure 12. The measurement matrix \tilde{H} in (19) is then constructed based on the POD basis functions (see “POD Model Reduction”).

ESC-Based Observer Autotuning

After discretizing the time, the dynamics of the ROM-G model (18) and its observer (19) are represented by the difference equation

$$\hat{q}_{k+1} = f(\hat{q}_k, y_k, \hat{\mu})$$

where $k = 0, 1, \dots, T_f$, $\hat{q}_k = \hat{q}(k\Delta t)$, and $T_f = t_f/\Delta t$. In the numerical experiment, measurements $\{y_k\}_{k=0}^{T_f}$ are generated through forward simulation of the model. These measurements are then used to calculate the observer state $\{\hat{q}_k\}_{k=0}^{T_f}$, starting with an initial guess for $\hat{\mu}$.

Once the observer state $\{\hat{q}_k\}_{k=0}^{T_f}$ is obtained, the estimated parameter is updated iteratively using the ESC algorithm. This process, alternating between running the observer and updating the parameter, continues for I iterations, where I is the predefined number of iterations. The final estimate of the uncertain parameter is calculated by averaging the parameter values over the last 10% of the steps. Finally, the observer is rerun with the updated learned parameter to refine the results. This entire procedure is outlined in Algorithm 2.

Performance of the Observer

In these tests, the initial observer state is set as $\hat{q}(0) = (1 + \epsilon)q(0)$, where $\epsilon = 0.1$, indicating a 10% error in the initial conditions of the observer state. The true viscosity is normalized to $\mu = 1$, while the initial estimate of the viscosity is set to $\hat{\mu}^{(1)} = 0.85$, representing a 15% error. The parameters for the dither signals are configured as $a_1 = 0.02$, $\omega_1 = 1,000$, and $\delta = 0.001$. The learned parameter is updated progressively over the course of 1,000 iterations.

The simulation of the learning-based observer was carried out, and Figure 13 presents the time evolution of the first two modes for both velocity and temperature. It shows the true values (red), the estimates before learning (green), and the estimates after learning (blue). For all four states,

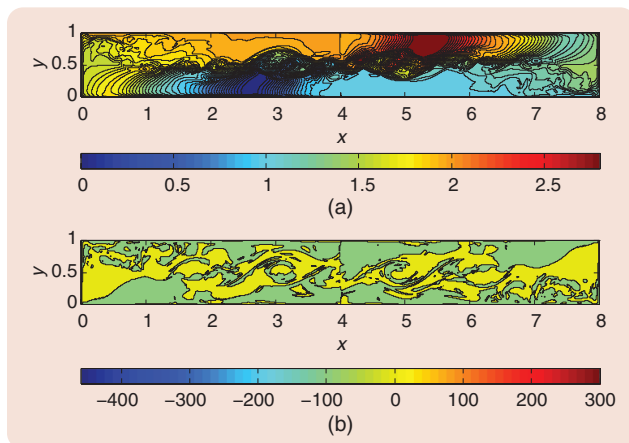


FIGURE 11 A true solution profile snapshot at 8 s. (a) The temperature contour (the exact solution at $t = 8$ s). (b) The vorticity contour (the exact solution at $t = 8$ s).




FIGURE 12 The sensor placement for both temperature (red) and velocity (blue).

POD Model Reduction

Following Galerkin projection onto the subspace spanned by the POD basis functions, we have

$$\begin{aligned} \mathbf{v}^{\text{pod}}(\mathbf{x}, t) &= \mathbf{v}_0(\mathbf{x}) + \sum_{i=1}^{r_v} q_i(t) \phi_i^v(\mathbf{x}) \\ T^{\text{pod}}(\mathbf{x}, t) &= T_0(\mathbf{x}) + \sum_{i=r_v+1}^{r_v+r_T} q_i(t) \phi_i^T(\mathbf{x}) \end{aligned} \quad (\text{S12})$$

where $\mathbf{v}_0(\mathbf{x})$ and $T_0(\mathbf{x})$ are the steady-state solution to (10)–(12) and $\phi_i^v(\mathbf{x})$ and $\phi_i^T(\mathbf{x})$ are POD basis functions given by

$$\phi_i^v(\mathbf{x}) = \frac{1}{t} \int_0^{t_i} (\mathbf{v}_{\text{sim}}(\mathbf{x}, t) - \mathbf{v}_0(\mathbf{x})) w_i(t) dt \quad (\text{S13})$$

for $i = 1, \dots, r_v$, and

$$\phi_i^T(\mathbf{x}) = \frac{1}{t} \int_0^{t_i} (T_{\text{sim}}(\mathbf{x}, t) - T_0(\mathbf{x})) w_i(t) dt \quad (\text{S14})$$

for $i = r_v + 1, \dots, r_v + r_T$, with the orthogonal weight functions $w_i(t)$ satisfying $\int_0^{t_i} w_i(t) w_j(t) dt = 0$ if $i \neq j$. The coefficients $q_i(t)$ in (S12) are the dynamical states representing the POD-ROM of Boussinesq equations.

Considering the POD projection error, for example, [S6], let $\mathcal{T} : \mathcal{H} \rightarrow \hat{\mathcal{H}}$ be the orthogonal projector from \mathcal{H} (infinite-dimension Hilbert space) to the finite-dimension space $\hat{\mathcal{H}}$ (hence, $\|\mathcal{T}\|_{\mathcal{H}} = 1$ and \mathcal{T}^\dagger are the injection of $\hat{\mathcal{H}}$ into \mathcal{H} : $\mathcal{T}^\dagger \hat{z} = z$ for all $\hat{z} \in \hat{\mathcal{H}} \subset \mathcal{H}$). Then we define the *reduced estimation error* as

$$\mathbf{e}(t) = \hat{z}(t) - \mathcal{T}z(t) \in \hat{\mathcal{H}}. \quad (\text{S15})$$

This can be used as a proxy for the *state estimation error*

$$\mathbf{e}_{\text{se}} \equiv \mathcal{T}^\dagger \hat{z} - z \in \mathcal{H} \quad (\text{S16})$$

when \mathcal{T} produces a small projection error ($z - \mathcal{T}^\dagger \mathcal{T}z$) since

$$\mathbf{e}_{\text{se}}(t) = \mathcal{T}^\dagger \mathbf{e}(t) - (z(t) - \mathcal{T}^\dagger \mathcal{T}z(t)). \quad (\text{S17})$$

When $\hat{\mathcal{H}}$ is the span of r dominant POD basis functions and \mathcal{T}_{POD} is the corresponding projection for a specific trajectory z , that is, $\mathcal{T} \equiv \mathcal{T}_{\text{POD}} : \mathcal{H} \rightarrow \hat{\mathcal{H}}$,

$$[\mathcal{T}_{\text{POD}}z](\cdot) = \sum_{i=1}^r \phi_i(\cdot) \langle \phi_i, z \rangle_{\mathcal{H}} \quad (\text{S18})$$

then \mathcal{T}_{POD} minimizes the projection error

$$\mathcal{P}(\mathcal{T}, z) = \left(\int_0^{t_i} \|z(t) - \mathcal{T}^\dagger \mathcal{T}z(t)\|_{\hat{\mathcal{H}}}^2 dt \right)^{1/2} \quad (\text{S19})$$

over all projections \mathcal{T} into subspaces of \mathcal{H} with dimension r and where t_i denotes the finite time support over which the projection error is evaluated [S7].

REFERENCES

- [S6] M. Benosman and J. Borggaard, "Robust nonlinear state estimation for a class of infinite-dimensional systems using reduced-order models," *Int. J. Control*, vol. 94, no. 5, pp. 1309–1320, 2019.
- [S7] P. Holmes, J. L. Lumley, and G. Berkooz, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge, U.K.: Cambridge Univ. Press, 1998.

the estimates after the ESC-based learning of the uncertain parameter show a marked improvement, aligning significantly closer to the true values. As evident in Figure 14, the norm of the estimation error vector containing all modes converges to zero after learning (blue), while the error norm before learning demonstrates a diverging behavior.

The evolution of the learned parameter $\hat{\mu}^{(i)}$ throughout the learning iterations is depicted in Figure 15(a). Despite an initial estimate with a 15% error, the ESC approach shows that after approximately 200 iterations, the learned parameter $\hat{\mu}^{(i)}$ converges near the true value $\hat{\mu}$. The amplitude of fluctuations around the averaged learned parameter is directly influenced by the dither signal amplitude a , which can be adjusted by the user. While reducing a decreases the fluctuations, it typically slows down the convergence. This tradeoff can be optimized through performance improvements to the ESC algorithm, such as those proposed in [51], which guarantee convergence to the global optimum even in the presence of local extrema. Refer also to [49] for more recent techniques tackling this issue. Finally, the iterative update of the cost function is given in Figure 15(b), where convergence to the minimum value of Q is observed after

ALGORITHM 2 Iterative learning for online parameter estimation in the state observer.

Input: $\{y_k\}_{k=0}^{T_r}$, $\hat{\mu}^{(1)}$, \hat{q}_0 ;
 $z \leftarrow \hat{\mu}^{(1)}$;
for $i = 1, 2, \dots, I$, **do**
 for $k = 0, \dots, t_i - 1$, **do**
 $\hat{q}_{k+1} \leftarrow f(\hat{q}_k, y_k, \hat{\mu}^{(i)})$;
 $\hat{y}_k \leftarrow y_k - H\hat{q}_k$;
 end for
 $Q \leftarrow \text{Trapz}\{\hat{y}_k^T \hat{y}_k\}_{k=0}^{t_i}$;
 $z \leftarrow z + a\delta \sin(\omega i\delta + \frac{\pi}{2})Q$;
 $\hat{\mu}^{(i+1)} \leftarrow z + a \sin(\omega i\delta - \frac{\pi}{2})$;
end for
 $\hat{\mu}^{\text{ave}} \leftarrow \text{Mean}(\{\hat{\mu}^{(i)}\}_{i=0.9..I})$;
for $k = 0, \dots, T_r - 1$, **do**
 $\hat{q}_{k+1} \leftarrow f(\hat{q}_k, y_k, \hat{\mu}^{\text{ave}})$;
end for
Output: $\hat{\mu}^{\text{ave}}$, $\{\hat{q}_k\}_{k=0}^{T_r}$

100 steps. Overall, this learning-based state estimation method exhibits strong performance in the numerical study of challenging 2D Boussinesq equations.

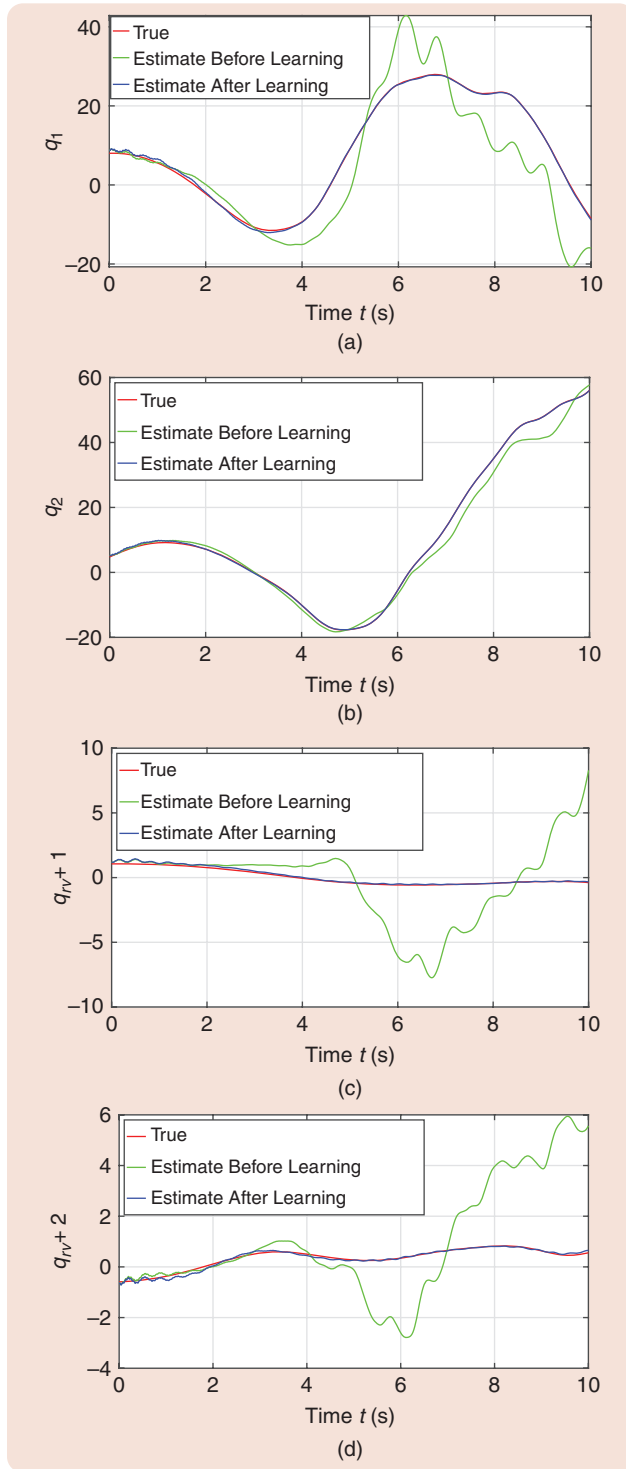


FIGURE 13 The evolution of the first and second modes in both the velocity and temperature states of the ROM-G, showing the true value (red), estimate before learning (green), and estimate after learning (blue). For all the modes given here, we observe that the modes' estimate is highly improved after learning the viscosity uncertain parameter $\hat{\mu}$.

For a spatial 2D view of the ESC performance, **Figure 11** shows the true temperature and velocity at 8 s, while **Figures 16** and **17** provide snapshots of the estimated spatial temperature and velocity at 8 s, along with the associated




FIGURE 14 A comparison of the estimation error's norm before learning (green) and after learning (blue), which illustrates the high improvement of the estimator performance through extremum seeking-based real-time tuning of the parameter $\hat{\mu}$.

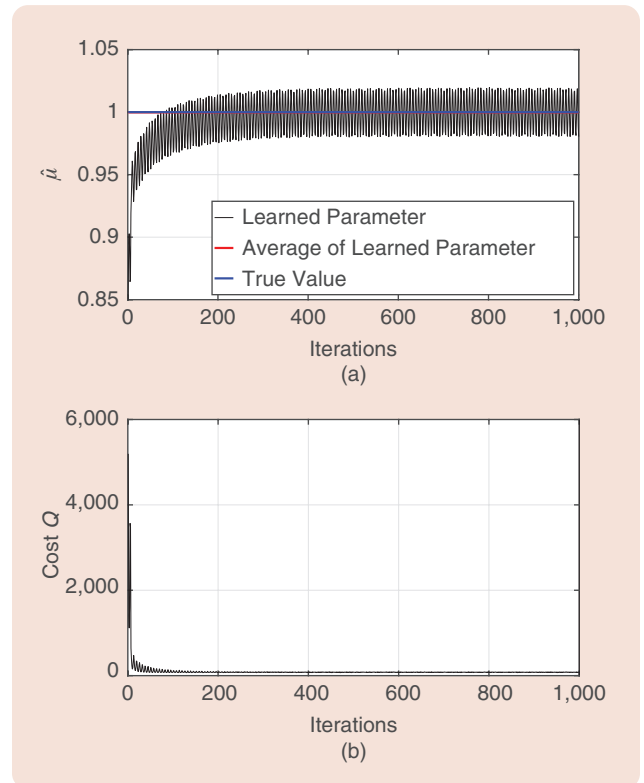


FIGURE 15 Updates of (a) the learned parameter and (b) the cost value Q by extremum seeking through 1,000 iterations. After 200 steps, the learned parameter $\hat{\mu}^{(l)}$ stays in the neighborhood of the true value $\mu = 1$. The averaged value of the learned parameter over the last 100 steps becomes $\hat{\mu}^{ave} = 1$, the same value as the true value. The value of the cost is largely decreased during the first 100 steps, and after that, it maintains small oscillations around the minimum.

The application of control theory in industry has evolved from simple linear PID controllers to more advanced nonlinear feedback systems.

spatial estimation error, when the ESC tuning is turned off. These are in contrast with the results obtained when the online ESC tuning is applied, as apparent in Figures 18 and 19. It is evident that ESC tuning reduced the estimation error by an order of magnitude.

CONCLUSION

The application of control theory in industry has evolved from simple linear PID controllers to more advanced nonlinear feedback systems. This progression has been driven by ongoing advancements in control theory, particularly in the

development of feedback laws that ensure stability and robustness—key factors in any industrial implementation.

In particular, data-driven nonlinear controllers have recently started gaining traction in industrial settings. In this article, we focused on several key applications where a data-driven controller has been used, namely, ESC, to continuously optimize and enhance the performance of feedback systems.

For example, we demonstrated how ESC has been applied to electromagnetic brakes for smooth elevator control in high-rise buildings. In this case, ESC enables continuous

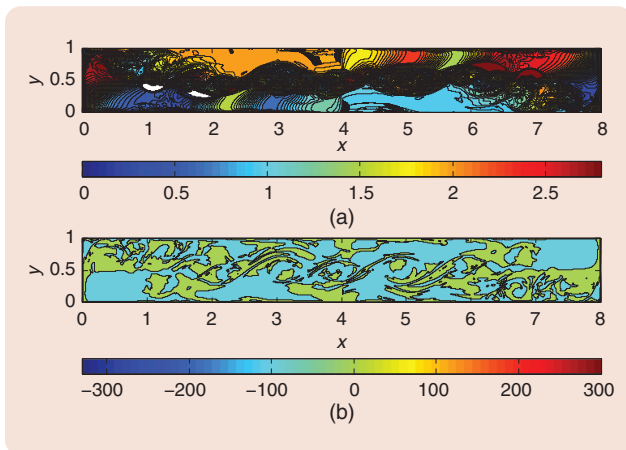


FIGURE 16 A ROM-based (without ESC tuning) estimated solution profile snapshot at 8 s. (a) The temperature contour (ROM-G). (b) The vorticity contour (ROM-G).




FIGURE 18 A ROM-based (with ESC tuning) estimated solution profile snapshot at 8 s. (a) The temperature contour (ROM-G). (b) The vorticity contour (ROM-G).

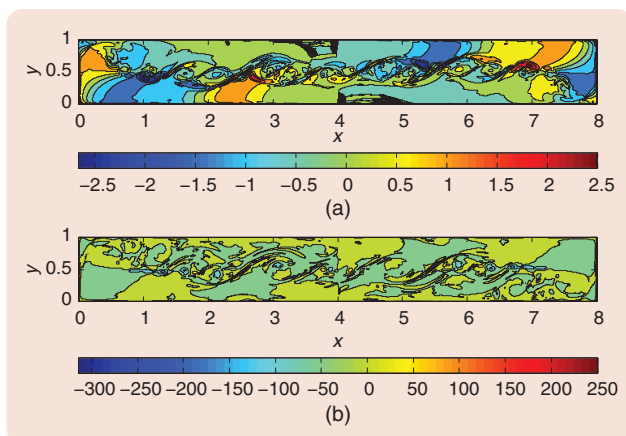


FIGURE 17 The ROM-based (without ESC tuning) estimated solution error at 8 s. (a) The temperature error contour (ROM-G). (b) The vorticity error contour (ROM-G).

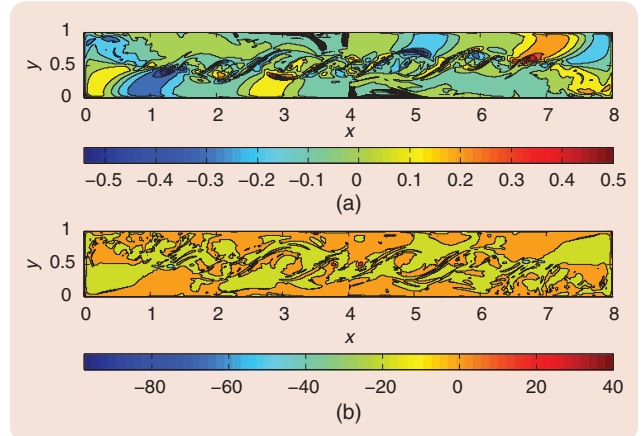


FIGURE 19 The ROM-based (with ESC tuning) estimated solution error at 8 s. (a) The temperature error contour (ROM-G). (b) The vorticity error contour (ROM-G).

ESC has been effectively applied in various industrial settings to enhance the estimation and control of highly complex and nonlinear dynamics.

tuning of a nonlinear feedback controller, effectively compensating for the aging of the electromagnet actuator. This performance guarantee is crucial, given the critical safety role that the electromagnetic actuator brake plays in the elevator system.

We also highlighted the application of ESC in the field of 5G wireless communications, where we employed ESC to automatically tune the performance of a DPA. Traditionally, this tuning process relied on manual expert-driven adjustments conducted in the factory. We transformed this offline expert-based tuning into an online adaptation control problem, one that ESC is ideally suited to address. In addition to performance enhancement, this application required the design of an online autotuner that is both simple to implement and efficient in its use of computational resources. As demonstrated in this article, dither-based multiparametric ESC proved to be the ideal solution for meeting these deployment constraints.

Finally, we highlighted work in the area of spatiotemporal fluidic systems modeling and estimation for HVAC applications. Traditional HVAC systems primarily rely on temperature measurements to control the unit and maintain room temperature at a setpoint. A significant shift from this conventional approach is obtained by considering the coupling between temperature and airflow within the room. This interaction, often overlooked in HVAC controller design, is crucial for effectively modeling, visualizing, estimating, and controlling the temperature distribution in a space. In this article, we demonstrated how we model this interaction and how ESC enables us to autotune these models and estimators in real time.

In conclusion, we have demonstrated how ESC has been effectively applied in various industrial settings to enhance the estimation and control of highly complex and nonlinear dynamics. Despite significant advances, several research opportunities remain for the academic community to enhance ESC's applicability in industrial settings. Particularly in factory automation and warehouse applications, where discrete-event and hybrid systems predominate, the system dynamics often experience sudden event-based transitions. This characteristic necessitates the development of data-driven discrete-event ESC methods. Another promising application lies in artificial intelligence, where models are initially trained on large datasets and subsequently deployed for real-time inference. In this context, novel ESC algorithms for large-scale systems could enable adaptive fine-tuning, allowing models to better respond to changing real-world conditions.

ACKNOWLEDGMENT

This work was conducted prior to Dr. Benosman's affiliation with Amazon Robotics. The opening artwork was generated using OpenAI ChatGPT. Prompt by Mouhacine Benosman (summarized): "Futuristic scientific illustration of extremum-seeking optimization represented as a glowing three-dimensional optimization surface with a moving point climbing toward the maximum. Surrounding semi-transparent overlays illustrate application domains including aerospace, industrial automation, renewable energy, medical devices, and communications. The style combines wireframe geometry with glowing blue, orange, and gold gradients, resulting in a clean, editorial, research-cover aesthetic."

AUTHOR INFORMATION

Mouhacine Benosman (m_benosman@ieee.org) received the Ph.D. degree in applied mathematics in 2002 from Ecole Centrale de Nantes, Nantes, France, and the M.A. degree in psychology from Harvard University, Cambridge, MA, USA, in 2025. He is a principal research scientist at Amazon Robotics (AR), North Reading, MA 01864 USA. Before joining AR, he worked at Reims University, France; Strathclyde University, U.K.; National University of Singapore; and Mitsubishi Electric Research Labs, Cambridge. His research interests include multiagent distributed control, with applications to robotics, physics-informed machine learning for modeling, estimation, and control of infinite-dimensional systems, with applications to fluid dynamics, data-driven robust reinforcement learning control, and applying psychology theory to artificial intelligence systems design. He served as associate editor of *IEEE Control Systems Letters* and *Journal of Optimization Theory and Applications*, and he served as senior editor of *International Journal of Adaptive Control and Signal Processing*. He is a Senior Member of IEEE.

REFERENCES

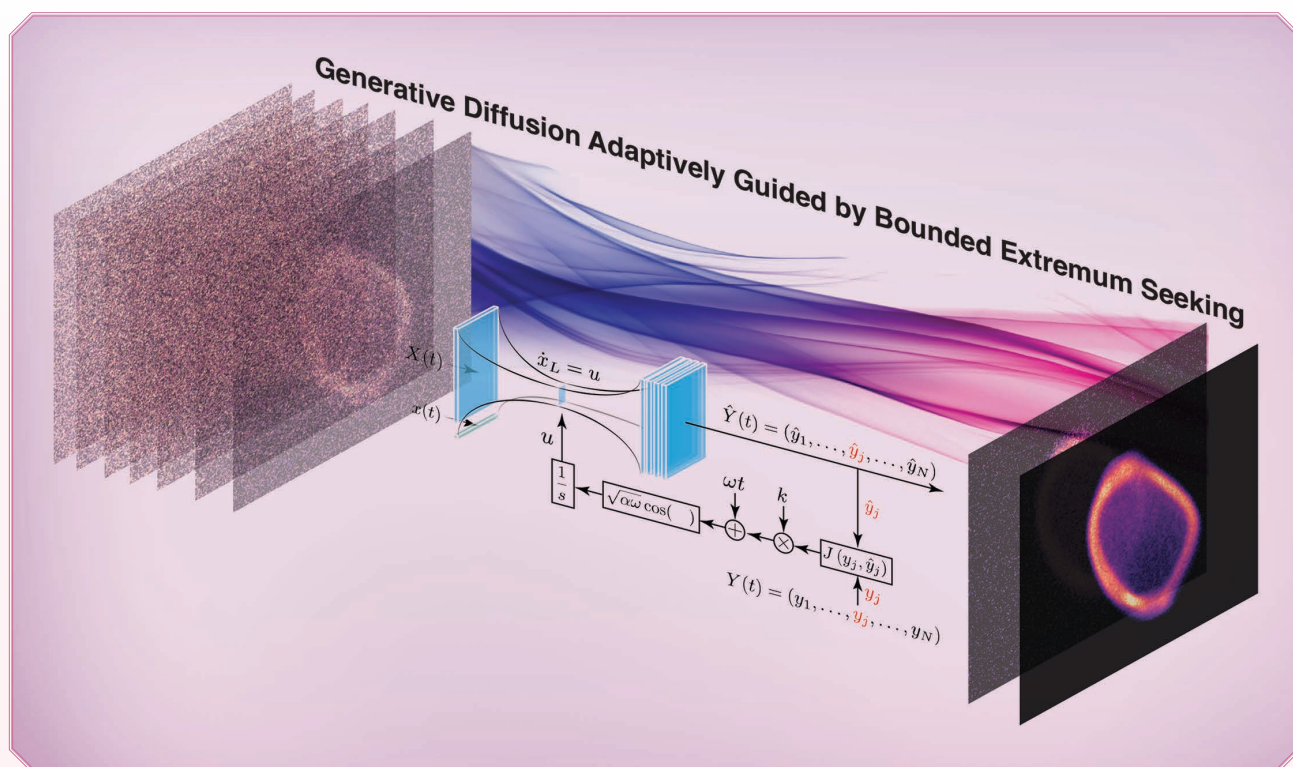
- [1] M. Leblanc, "Sur l'électrification des chemins de fer au moyen de courants alternatifs de fréquence élevée," *Revue Générale de L'Electricité*, vol. 12, no. 8, pp. 275–277, 1922.
- [2] M. Krstić and H.-H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, no. 4, pp. 595–601, 2000, doi: [10.1016/S0005-1098\(99\)00183-1](https://doi.org/10.1016/S0005-1098(99)00183-1).
- [3] S.-J. Liu and M. Krstić, *Stochastic Averaging and Stochastic Extremum Seeking*. London, U.K.: Springer-Verlag, 2012.
- [4] P. Frihauf, M. Krstić, and T. Basar, "Nash equilibrium seeking in noncooperative games," *IEEE Trans. Autom. Control*, vol. 57, no. 5, pp. 1192–1207, May 2012, doi: [10.1109/TAC.2011.2173412](https://doi.org/10.1109/TAC.2011.2173412).
- [5] J. Poveda and A. Teel, "A framework for a class of hybrid extremum seeking controllers with dynamic inclusions," *Automatica*, vol. 76, pp. 113–126, Feb. 2017, doi: [10.1016/j.automatica.2016.10.029](https://doi.org/10.1016/j.automatica.2016.10.029).
- [6] K. Ariyur and M. Krstić, *Real Time Optimization by Extremum Seeking Control*. Hoboken, NJ, USA: Wiley, 2003.

- [7] C. Zhang and R. Ordonez, *Extremum-Seeking Control and Applications: A Numerical Optimization-Based Approach*. London, U.K.: Springer-Verlag, 2012, doi: [10.1007/978-1-4471-2224-1_2](https://doi.org/10.1007/978-1-4471-2224-1_2).
- [8] A. Scheinker and M. Krstić, *Extremum-Seeking Control and Applications: A Numerical Optimization-Based Approach*. London, U.K.: Springer-Verlag, 2016.
- [9] J. Ziegler and N. Nichols, "Optimum settings for automatic controllers," *Trans. ASME*, vol. 64, no. 8, pp. 759–765, 1942, doi: [10.1115/1.4019264](https://doi.org/10.1115/1.4019264).
- [10] H. Hjalmarsson, S. Gunnarsson, and M. Gevers, "A convergent iterative restricted complexity control design scheme," in *Proc. IEEE Conf. Decis. Control*, 1994, pp. 1735–1740, doi: [10.1109/CDC.1994.411185](https://doi.org/10.1109/CDC.1994.411185).
- [11] O. Lequin, M. Gevers, M. Mossberg, E. Bosmans, and L. Triest, "Iterative feedback tuning of PID parameters: Comparison with classical tuning rules," *Control Eng. Pract.*, vol. 11, no. 9, pp. 1023–1033, 2003, doi: [10.1016/S0967-0661\(02\)00303-9](https://doi.org/10.1016/S0967-0661(02)00303-9).
- [12] H. Hjalmarsson, "Iterative feedback tuning—an overview," *Int. J. Adapt. Control Signal Process.*, vol. 16, no. 5, pp. 373–395, 2002, doi: [10.1002/acs.714](https://doi.org/10.1002/acs.714).
- [13] N. J. Killingsworth and M. Krstić, "PID tuning using extremum seeking: Online, model-free performance optimization," *IEEE Control Syst. Mag.*, vol. 26, no. 1, pp. 70–79, Feb. 2006, doi: [10.1109/MCS.2006.1580155](https://doi.org/10.1109/MCS.2006.1580155).
- [14] L. Koszalka, R. Rudek, and I. Pozniak-Koszalka, "An idea of using reinforcement learning in adaptive control systems," in *Proc. Int. Conf. Netw. Int. Conf. Syst. Int. Conf. Mobile Commun. Learn. Technol. (ICN/ICONS/MCL)*, 2006, p. 190, doi: [10.1109/ICN/ICONS/MCL.2006.52](https://doi.org/10.1109/ICN/ICONS/MCL.2006.52).
- [15] H. Hjalmarsson, M. Gevers, S. Gunnarsson, and O. Lequin, "Iterative feedback tuning: Theory and applications," *IEEE Control Syst.*, vol. 18, no. 4, pp. 26–41, Aug. 1998, doi: [10.1109/37.710876](https://doi.org/10.1109/37.710876).
- [16] Y. Wang, F. Gao, and F. J. D. Iii, "Survey on iterative learning control, repetitive control, and run-to-run control," *J. Process Control*, vol. 19, no. 10, pp. 1589–1600, 2009, doi: [10.1016/j.jprocont.2009.09.006](https://doi.org/10.1016/j.jprocont.2009.09.006).
- [17] H. Hjalmarsson, "Control of nonlinear systems using iterative feedback tuning," in *Proc. Amer. Control Conf. (ACC) (IEEE Cat. No. 98CH36207)*, 1998, pp. 2083–2087, doi: [10.1109/ACC.1998.702994](https://doi.org/10.1109/ACC.1998.702994).
- [18] J. Sjöberg and M. Agarwal, "Model-free repetitive control design for nonlinear systems," in *Proc. IEEE Conf. Decis. Control*, 1996, pp. 2824–2829, doi: [10.1109/CDC.1996.573544](https://doi.org/10.1109/CDC.1996.573544).
- [19] F. DeBruyne, B. D. O. Anderson, M. Gevers, and N. Linard, "Iterative controller optimization for nonlinear systems," in *Proc. 36th IEEE Conf. Decis. Control*, 1997, pp. 3749–3754, doi: [10.1109/CDC.1997.652439](https://doi.org/10.1109/CDC.1997.652439).
- [20] J. Sjöberg et al., "Iterative controller optimization for nonlinear systems," *Control Eng. Pract.*, vol. 11, no. 9, pp. 1079–1086, 2003, doi: [10.1016/S0967-0661\(02\)00231-9](https://doi.org/10.1016/S0967-0661(02)00231-9).
- [21] H. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2002.
- [22] A. Katz, J. Wood, and D. Chokola, "The evolution of pa linearization: From classic feedforward and feedback through analog and digital predistortion," *IEEE Microw. Mag.*, vol. 17, no. 2, pp. 32–40, Feb. 2016, doi: [10.1109/MMM.2015.2498079](https://doi.org/10.1109/MMM.2015.2498079).
- [23] D. J. Sheppard, J. Powell, and S. C. Cripps, "An efficient broadband reconfigurable power amplifier using active load modulation," *IEEE Microw. Wireless Compon. Lett.*, vol. 26, no. 6, pp. 443–445, Jun. 2016, doi: [10.1109/LMWC.2016.2559503](https://doi.org/10.1109/LMWC.2016.2559503).
- [24] W. H. Doherty, "A new high efficiency power amplifier for modulated waves," *Proc. Inst. Radio Engineers*, vol. 24, no. 9, pp. 1163–1182, Sep. 1936, doi: [10.1109/JRPROC.1936.228468](https://doi.org/10.1109/JRPROC.1936.228468).
- [25] F. Wang, A. H. Yang, D. F. Kimball, L. E. Larson, and P. M. Asbeck, "Design of wide-bandwidth envelope-tracking power amplifiers for OFDM applications," *IEEE Trans. Microw. Theory Techn.*, vol. 53, no. 4, pp. 1244–1255, Apr. 2005, doi: [10.1109/TMTT.2005.845716](https://doi.org/10.1109/TMTT.2005.845716).
- [26] T. Cappello, C. Florian, T. W. Barton, M. Litchfield, and Z. Popovic, "Multi-level supply-modulated Chireix outphasing for LTE signals," in *Proc. IEEE MTT-S Int. Microw. Symp. (IMS)*, Jun. 2017, pp. 1846–1849, doi: [10.1109/MWSYM.2017.8059012](https://doi.org/10.1109/MWSYM.2017.8059012).
- [27] H. Chireix, "High power outphasing modulation," *Proc. Inst. Radio Engineers*, vol. 23, no. 11, pp. 1370–1392, Nov. 1935, doi: [10.1109/JRPROC.1935.227299](https://doi.org/10.1109/JRPROC.1935.227299).
- [28] R. Darraji, F. M. Ghannouchi, and O. Hammi, "A dual-input digitally driven Doherty amplifier architecture for performance enhancement of Doherty transmitters," *IEEE Trans. Microw. Theory Techn.*, vol. 59, no. 5, pp. 1284–1293, May 2011, doi: [10.1109/TMTT.2011.2106137](https://doi.org/10.1109/TMTT.2011.2106137).
- [29] R. Darraji and F. M. Ghannouchi, "Digital Doherty amplifier with enhanced efficiency and extended range," *IEEE Trans. Microw. Theory Techn.*, vol. 59, no. 11, pp. 2898–2909, Nov. 2011, doi: [10.1109/TMTT.2011.2166122](https://doi.org/10.1109/TMTT.2011.2166122).
- [30] C. M. Andersson, D. Gustafsson, J. C. Cahuana, R. Hellberg, and C. Fager, "A 1–3-GHz digitally controlled dual-RF input power-amplifier design based on a Doherty-outphasing continuum analysis," *IEEE Trans. Microw. Theory Techn.*, vol. 61, no. 10, pp. 3743–3752, Oct. 2013, doi: [10.1109/TMTT.2013.2280562](https://doi.org/10.1109/TMTT.2013.2280562).
- [31] A. Piacibello et al., "Comparison of S-band analog and dual-input digital Doherty power amplifiers," in *Proc. 13th Eur. Microw. Integr. Circuits Conf. (EuMIC)*, Madrid, Spain, 2018, pp. 269–272, doi: [10.23919/EuMIC.2018.8539954](https://doi.org/10.23919/EuMIC.2018.8539954).
- [32] R. Darraji, P. Mousavi, and F. M. Ghannouchi, "Doherty goes digital: Digitally enhanced Doherty power amplifiers," *IEEE Microw. Mag.*, vol. 17, no. 8, pp. 41–51, Aug. 2016, doi: [10.1109/MMM.2016.2561478](https://doi.org/10.1109/MMM.2016.2561478).
- [33] R. Horst, P. M. Pardalos, and N. V. Thoai, *Introduction to Global Optimization: Non-convex Optimization and Its Applications*. Berlin, Heidelberg: Springer-Verlag, 2002.
- [34] S. Niu, A. M. Koushik, R. Ma, K. H. Teo, S. Shinjo, and Y. Komatsuzaki, "Stochastically approximated multiobjective optimization of dual input digital Doherty power amplifier," in *Proc. IEEE 10th Int. Workshop Comput. Intell. IEEE Int. Symp. Spread Spectr. Tech. Appl. (IWCIA)*, Hiroshima, Japan, 2017, pp. 147–152, doi: [10.1109/IWCIA.2017.8203576](https://doi.org/10.1109/IWCIA.2017.8203576).
- [35] R. Ma et al., "Machine-learning based digital Doherty power amplifier," in *Proc. IEEE Int. Symp. Radio-Freq. Integration Technol. (RFIT)*, Melbourne, VIC, Australia, 2018, pp. 1–3, doi: [10.1109/RFIT.2018.8524126](https://doi.org/10.1109/RFIT.2018.8524126).
- [36] C. Kantana, M. Benosman, R. Ma, and Y. Komatsuzaki, "A system approach for efficiency enhancement and linearization technique of dual-input Doherty power amplifier," *IEEE J. Microwaves*, vol. 3, no. 1, pp. 115–133, Jan. 2023, doi: [10.1109/JMW.2022.3226838](https://doi.org/10.1109/JMW.2022.3226838).
- [37] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, 1983, doi: [10.1126/science.220.4598.671](https://doi.org/10.1126/science.220.4598.671).
- [38] N. Metropolis, A. W. Rosenbluth, N. M. Rosenbluth, H. Augusta, and E. Teller, "Equation of state calculations by fast computing machines," *J. Chem. Phys.*, vol. 21, no. 6, pp. 1087–1092, 1953, doi: [10.1063/1.1699114](https://doi.org/10.1063/1.1699114).
- [39] J. Borggaard, J. A. Burns, A. Surana, and L. Zietsman, "Control, estimation and optimization of energy efficient buildings," in *Proc. Am. Control Conf. (ACC)*, 2009, pp. 837–841, doi: [10.1109/ACC.2009.5160552](https://doi.org/10.1109/ACC.2009.5160552).
- [40] R. Ciesla, *The Book of Chatbots: From ELIZA to ChatGPT*. New York, NY, USA: Springer-Verlag, 2024.
- [41] M. Chevalier, J. Hoepffner, T. R. Bewley, and D. S. Henningson, "State estimation in wall-bounded flow systems. Part 2. Turbulent flows," *J. Fluid Mech.*, vol. 552, pp. 167–187, Apr. 2006, doi: [10.1017/S0022112005008578](https://doi.org/10.1017/S0022112005008578).
- [42] M. Guay and N. Harihatan, "Airflow velocity estimation in building systems," in *Proc. Am. Control Conf. (ACC)*, 2009, pp. 908–913, doi: [10.1109/ACC.2008.4586608](https://doi.org/10.1109/ACC.2008.4586608).
- [43] T. John, M. Guay, N. Hariharan, and S. Narayan, "POD-based observer for estimation in Navier–Stokes flow," *Comput. Chem. Eng.*, vol. 34, no. 6, pp. 965–975, 2010, doi: [10.1016/j.compchemeng.2009.12.001](https://doi.org/10.1016/j.compchemeng.2009.12.001).
- [44] W. MacKunis, S. V. Drakunov, M. Reyhanoglu, and L. Ukeiley, "Nonlinear estimation of fluid flow velocity fields," in *Proc. 50th Conf. Decis. Control (CDC)*, 2011, pp. 6931–6935, doi: [10.1109/CDC.2011.6161193](https://doi.org/10.1109/CDC.2011.6161193).
- [45] V. L. Kalb and A. E. Deane, "An intrinsic stabilization scheme for proper orthogonal decomposition based low-dimensional models," *Phys. Fluids*, vol. 19, no. 5, 2007, Art. no. 054106, doi: [10.1063/1.2723149](https://doi.org/10.1063/1.2723149).
- [46] M. Benosman, J. Borggaard, O. San, and B. Kramer, "Learning-based robust stabilization for reduced-order models of 2D and 3D Boussinesq equations," *Appl. Math. Modell.*, vol. 49, pp. 162–181, Sep. 2017, doi: [10.1016/j.apm.2017.04.032](https://doi.org/10.1016/j.apm.2017.04.032).
- [47] S. Koga, M. Benosman, and J. Borggaard, "Extremum seeking-based robust observer design for coupled thermal and fluid systems," *Int. J. Adapt. Control Signal Process.*, vol. 35, no. 7, pp. 1316–1335, 2021, doi: [10.1002/acs.3252](https://doi.org/10.1002/acs.3252).
- [48] M. Benosman, *Learning-Based Adaptive Control: An Extremum Seeking Approach—Theory and Applications*. Oxford, U.K.: Butterworth-Heinemann, 2016.
- [49] A. Scheinker, "100 years of extremum seeking: A survey," *Automatica*, vol. 161, Mar. 2024, Art. no. 111481, doi: [10.1016/j.automatica.2023.111481](https://doi.org/10.1016/j.automatica.2023.111481).
- [50] Z. Drmac and S. Gugercin, "A new selection operator for the discrete empirical interpolation method—Improved a priori error bound and extensions," *Methods Algorithms Sci. Comput.*, vol. 32, no. 8, pp. A631–A648, 2016.
- [51] Y. Tan, D. R. Nesic, I. M. Mareels, and A. Astolfi, "On global extremum seeking in the presence of local extrema," *Automatica*, vol. 45, no. 1, pp. 245–251, 2009, doi: [10.1016/j.automatica.2008.06.010](https://doi.org/10.1016/j.automatica.2008.06.010).

Adaptive Control Meets Deep Learning

EXTREMUM SEEKING IN LATENT SPACE WITH PARTICLE ACCELERATOR APPLICATIONS

ALEXANDER SCHEINKER 



Generative deep learning has recently emerged as a transformational technology for a wide range of tasks, including image generation from user-defined captions, prediction of 3D protein structures directly from DNA sequences, and large language models (LLMs) that can write custom computer code based on user prompts. One major challenge

faced by machine learning (ML)- and artificial intelligence (AI)-based tools is that of time-varying systems or systems with distribution shift. ML typically relies on brute-force retraining to readjust learned models when systems change. Extremum seeking (ES) is a model-independent adaptive feedback technique that can be used to stabilize unknown and open-loop unstable time-varying dynamic systems and to optimize their analytically unknown time-varying output functions. ES is model independent and robust to both noise and time variations. This article

Digital Object Identifier 10.1109/MCS.2025.3646850
Date of current version: 26 March 2026

provides a brief review of several generative deep learning techniques and how they have been combined with ES so that while the generative models act as highly detailed virtual diagnostics of otherwise inaccessible system states, incorporating adaptive feedback within their latent embeddings increases their robustness for use with uncertain and time-varying systems. The approach is demonstrated with applications for time-varying charged particle beams in particle accelerators.

INTRODUCTION

Although ML is technically a subset of AI, and deep learning is a subset of ML, in this article these terms will be used interchangeably. Data-driven AI and ML algorithms have been demonstrated to have powerful abilities to learn useful representations of complex systems directly from vast amounts of data. Two clear examples of this are LLMs and diffusion-based image-generating models. For language, the transformer architecture and its use of the attention mechanism was a major breakthrough in enabling deep neural networks to capture long-range dependencies between tokens (words) in long sequences of text and to learn the contextual embeddings in which the meaning of words changes depending on the surrounding context [1]. As a result, transformer-based LLMs have successfully learned effective statistical models of language. These LLMs are able to predict the next most likely word given a sequence of text so well that they can carry on conversations, answer questions, summarize large documents, and write custom computer code [2], [3].

For images, a statistical thermodynamics-inspired diffusion process for generative deep neural networks of gradually building up complex distributions from noise was a major breakthrough [4]. The diffusion approach has been utilized for the generation of high-resolution images [5], [6], [7]. Diffusion-based generative models are the state of the art for generating high-resolution images, especially when the images have a wide variety. The generative ability of diffusion-based models has made them powerful tools for a wide range of scientific applications. Diffusion-based models are being used for complex generative tasks in biology [8], including Google's AlphaFold 3, which is a diffusion-based model capable of predicting 3D structures of proteins, nucleic acids, and small molecules [9]. Diffusion-based models have also been used to study designs of new superconductors [10] and to generate megapixel-resolution virtual views of charged particle beams in large particle accelerators [11].

Although they are incredibly successful at learning useful representations over large data training sets, a major difficulty faced by AI and ML techniques is their use in dynamic, time-varying environments. This difficulty comes from the fact that once AI or ML models are trained for a given dataset, they are unable to adapt when applied to a time-varying system that continuously undergoes

large distribution shifts. In such cases, AI- and ML-based techniques typically require brute-force retraining. While retraining is possible for many tasks, such as fine-tuning pretrained LLMs for specialized applications, it is a slow process that cannot be used for short-time-scale applications, such as real-time control of quickly changing dynamic systems.

In contrast to deep learning-based approaches, which are in many ways analogous to feedforward systems, the control theory community has always worked under the assumption that we cannot fully trust models no matter how great the amount of data used to fit them and that we need active feedback to robustly handle uncertainty and time-varying situations. That is why control theory has provided the state-of-the-art methods for real-time control of everything from a building's temperature based on a thermostat setpoint, cruise control maintaining a set speed on dry, wet, or slippery roads, to rockets bringing humans to the Moon.

Although they are much more robust than ML-based methods, a similar challenge faced by control theory-based approaches is that they are designed for specific subsystems or components. For example, a kilometer-long particle accelerator may be composed of thousands of magnets and resonant accelerating cavities, and while an experienced human with a global view of the entire system understands how to adjust any of these components to improve beam quality, each component requires its own local feedback loop.

The control theory community has been developing hybrid control-ML approaches that take advantage of the robustness of adaptive feedback controllers and the learning ability of ML-based models, such as deep neural networks. A promising approach to combine deep learning methods and adaptive feedback control is to utilize generative deep learning models as virtual diagnostics or high-dimensional state space observers that provide views of otherwise inaccessible states of complex systems. These virtual diagnostics can be made to be robust and to track time-varying systems by incorporating adaptive feedback tuning of their learned internal representations, as shown in [Figure 1](#). These virtual diagnostics can then be used by robust adaptive feedback controllers to control the system's state in a robust way.

In the first demonstration of adaptive deep learning for particle accelerators, it was shown that, for controlling the shapes of intense relativistic charged particle beam distributions, deep neural networks can be used to give initial estimates of parameter settings, which are imprecise but do escape local minima based on their broad high-level estimate of the system. Those initial estimates can then be adaptively refined by robust methods, such as ES, to zoom in on and robustly track the optimal accelerator parameter settings despite the time variation of the entire complex system [12]. This adaptive deep learning approach relied on an available online beam diagnostic that provided

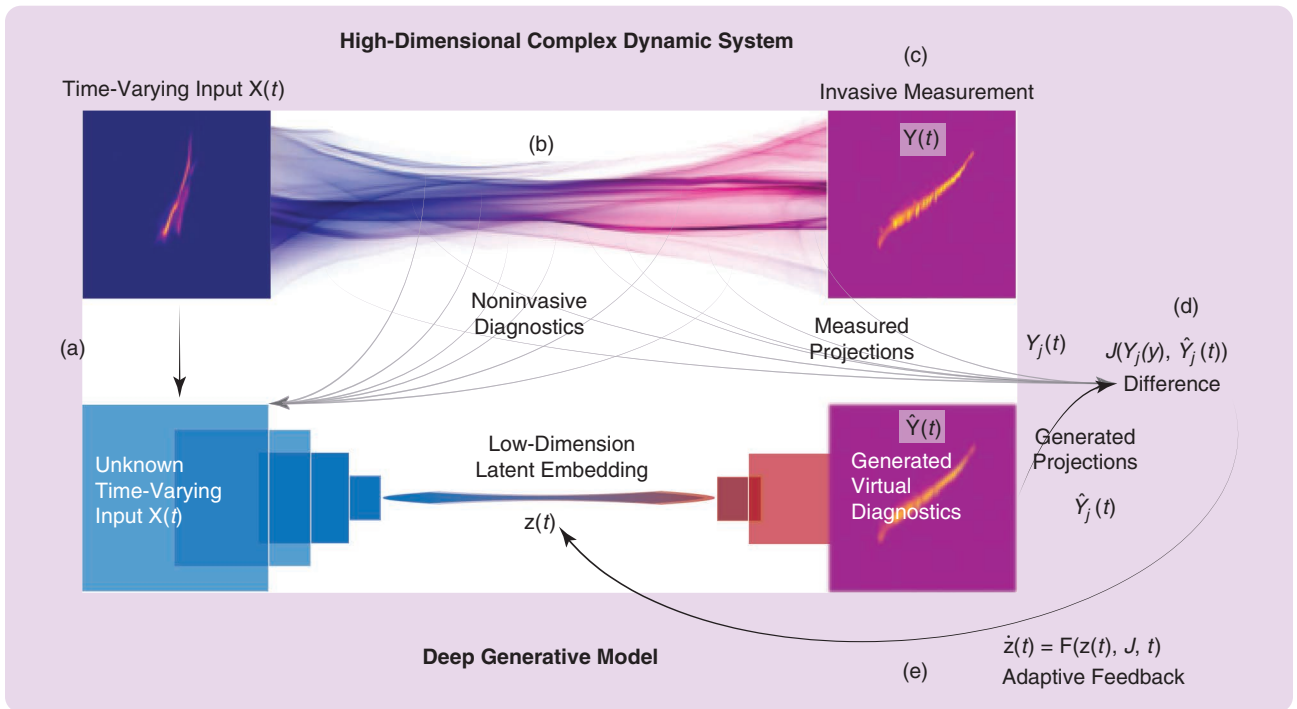


FIGURE 1 A high-level view is shown of adaptive generative deep learning for noninvasive virtual diagnostics of complex dynamics systems. (a) During the training phase, the deep generative models may have access to time-varying input beam conditions, $X(t)$, but we assume that later on, when we apply the model, the initial conditions may change, and we may lose access to them. (b) The initial conditions evolve according to some complex process, such as a 6D distribution of charged particles accelerating through a kilometer-long particle accelerator. (c) Our goal is to create an estimate of the system's time-varying state $Y(t)$ by using a generative model without having to rely on a destructive measurement, such as placing a screen into the beam pipe, which measures the beam's profile but destroys the beam during the measurement. (d) Our generative model creates an estimate $\hat{Y}(t)$, and from that we estimate projections (such as 1D current measurements) $\hat{Y}_j(t)$ that can be compared to measurements $Y_j(t)$ with the difference defining a cost J . (e) The cost function is minimized by adaptive feedback acting on the low-dimensional latent representation of the generative model $z(t)$, allowing us to track $Y(t)$ despite unknown and time-varying initial conditions.

images as online measurements of the charged particle beam's 2D energy-versus-time projection. These measurement images could be compared to a target 2D distribution, and the difference was used as a cost function to

Summary

Generative deep learning tools, such as variational auto-encoders, diffusion models, and transformers, attempt to learn the underlying probability distributions of large complex systems by utilizing deep neural networks. Once trained, the representations encoded by these models can be sampled to generate realistic new examples from the underlying distributions. This article explores the use of bounded ES-based feedback directly within the learned latent embeddings of such generative models to adaptively track the high-dimension states of complex time-varying systems without access to the systems' time-varying initial conditions and without brute-force retraining, which is the standard ML approach for time-varying systems with large distribution shifts.

guide the feedback loop. However, in many cases, such detailed measurements are not available online in a non-invasive way that does not interrupt system operation. Generative deep neural network-based approaches, such as conditional guided generative diffusion models, can generate complex virtual representations based only on noninvasive measurements, as recently demonstrated in [11], where vectors of noninvasive accelerator measurements guided a generative diffusion process to generate detailed virtual views of accelerator beams.

This article provides a brief overview of generative deep learning methods and some recent efforts on combining the ability of deep learning methods with adaptive feedback to create virtual high-level global views of large complex systems, including the use of adaptive feedback directly within the learned latent embeddings of these generative models to enable them to track time-varying systems without retraining. (See "Summary.")

To motivate the rest of this article, we start with an illustrative example from [13], shown in Figure 2, which is a real-world application of the abstract setup shown in Figure 1. We consider a particle accelerator in which it is

difficult (slow and destructive) to measure the input beam distribution $X(t)$, but it is much easier to measure the output beam $Y(t)$. We therefore train a generative model that solves the inverse problem of mapping measured output beams to estimates of the input beam $\hat{X}(t)$. The generated input beam estimates are used as initial conditions in a physics model that simulates the complex beam dynamics of the accelerator and provides its own estimate of the output beam $\hat{Y}(t)$. A cost function is then defined by comparing the simulated and measured output beam distributions:

$$E(t) = \|Y(t) - \hat{Y}(\mathbf{p}(t), \mathbf{z}(t))\| \quad (1)$$

where $\mathbf{p}(t)$ represents accelerator parameters in the physics model, and $\mathbf{z}(t)$ is a low-dimensional intermediate latent representation created by the generative model from which the estimated input beam is created (latent variables and generative models are described in more detail next).

After training, during operations, when we no longer have access to the input beam, and it is changing with time, as are the actual parameters of the accelerator, we track these changes by adaptively adjusting both $\mathbf{p}(t)$ and $\mathbf{z}(t)$ by using a recently developed novel form of bounded ES [14]. Bounded ES is incredibly simple and elegant compared to the standard ES algorithms that have large collections of

input and output filters that must be designed. The bounded ES feedback is simply defined as

$$\dot{z}_i(t) = \sqrt{\alpha_i \omega_i} \cos(\omega_i t + kE(t)) \quad (2)$$

$$\dot{p}_j(t) = \sqrt{\alpha_j \omega_j} \cos(\omega_j t + kE(t)) \quad (3)$$

which results in the minimization of $E(t)$ by finding the correct input beam distribution by optimizing $\mathbf{z}(t)$ and the correct accelerator settings by simultaneously optimizing $\mathbf{p}(t)$.

There are two major benefits of adjusting the latent representation \mathbf{z} . The first benefit of this method is that \mathbf{z} is a much lower dimensional object, $\mathcal{O}(10)$, as opposed to attempting to modify the unknown input beam image estimate itself, which even for a low-resolution 52×52 pixel image is a 2,704D object, which makes fast adaptive control feasible. The second benefit is that this becomes an unsupervised task, where we only use the network's predicted outputs $\hat{X}(t)$. We do not need access to the true difficult-to-measure $X(t)$ with this approach.

CONTROL AND AI

While reinforcement learning (RL)-based methods have combined aspects of optimal control theory with deep learning techniques, deep RL is still fundamentally a learning process in which deep neural network-based models of the controller, the value function, or both are learned for a specific dynamic

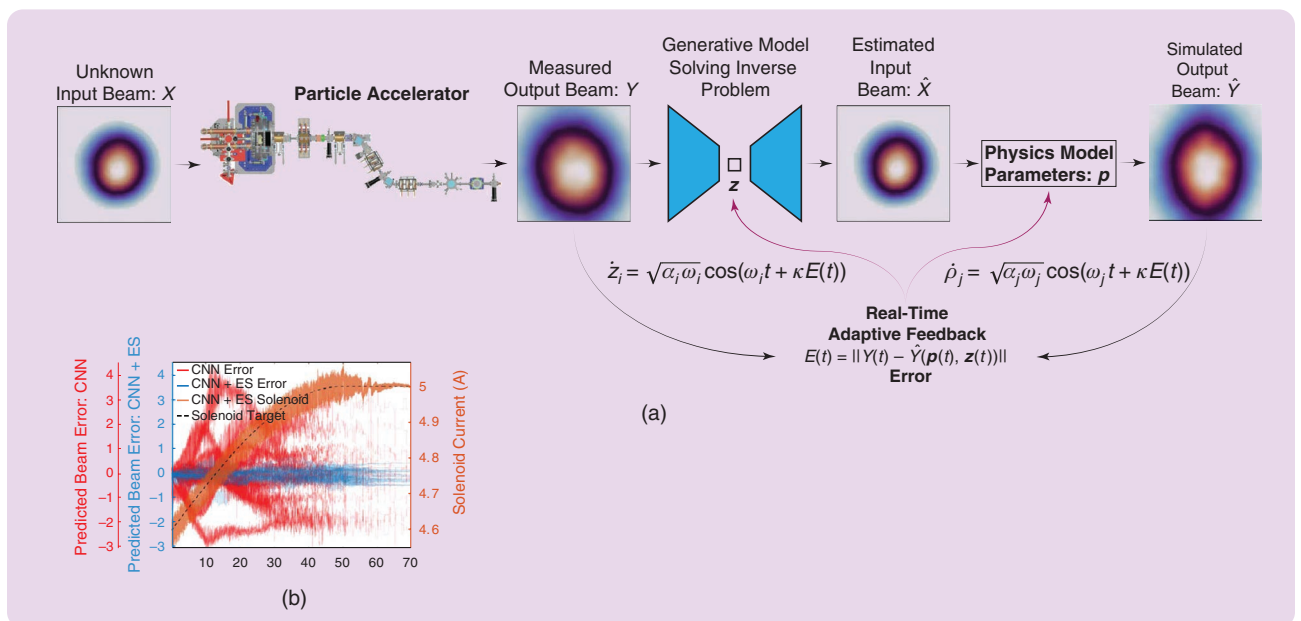


FIGURE 2 (a) We consider a particle accelerator for which measuring the time-varying input beam distribution $X(t)$ is difficult and would interrupt operations, while measurements of the output beam $Y(t)$ are available in real time. A generative model is developed to solve the inverse problem of mapping $Y(t)$ to estimates of the input beam $\hat{X}(t)$. Input beam estimates are used to initialize a physics model that attempts to simulate the beam evolution and predicts its own version of the output beam $\hat{Y}(t)$. Comparing the measured and predicted output beams, we define a cost function that is minimized by simultaneously adaptively tuning a low-dimensional latent representation of the generative model $\mathbf{z}(t)$ as well as parameters of the physics model $\mathbf{p}(t)$. (b) Minimizing this error in real time by using a recently developed bounded ES algorithm allows us to accurately track the time-varying input beam as well as the time-varying accelerator parameters, for which a pretrained generative model alone catastrophically fails. CNN: convolutional neural network.

system and need to be retrained if the system changes. For example, while an RL-based model might be able to master the game of chess, if any of the rules of chess were changed, such as allowing pawns to eat only the pieces directly in front of them rather than acting diagonally, any reasonable chess player could quickly adapt to the new set of rules, but the pre-trained RL agent would completely fail until it was again retrained on millions of virtual chess games that incorporate the new rule. This well-known issue, that RL methods struggle with changes to the environment for which they were designed, is an active area of research [15].

Approaches to this problem include efforts to speed up learning from fewer observations [16], [17], [18], and also to model RL problems as partially observable, in an attempt to design agents that are more robust [19]. Another approach is contextual RL, in which the underlying system is assumed to be parameterized by a set of contexts C , and a function is learned to map any given context $c \in C$ to the appropriate Markov decision process (MDP) [20], resulting in a collection of models that share the same state and action space. This approach then relies on first identifying which context the current state is in via clustering, based on previous state examples labeled by the known list of contexts, then switching to the appropriate MDP. More recent versions of contextual RL focus on more efficient training by using Bayesian methods to focus on specific scenarios in various contexts [21]. Despite ongoing efforts, the main limitations of RL-based methods remain; they still rely on retraining for any new context not already contained in C , mostly relying on assumptions of knowing C a priori or being able to somehow parameterize it in a continuous way. An interesting recently developed approach for improving the robustness of RL-based methods is the use of a stochastic model predictive control framework for nonlinear systems subject to unbounded process noise with closed-loop guarantees [22].

This is also one of the reasons why, despite years of success in other problems and so many billions of invested dollars in research and compute costs for training the world's biggest neural networks, and despite all of the hype, ML-based methods still struggle with dynamic problems, such as self-driving cars, especially when they face new and uncertain conditions (bad weather). In terms of real-time control applications, RL-based methods have focused on specific subsets of large complex systems, performing iterative optimization while relying on the fact that the rest of the system is properly operating, thanks to large numbers of feedback control algorithms. For example, the highly celebrated RL-based approach for control of tokamak plasmas focused on a very specific aspect of the tokamak's operation while relying on operators and many active adaptive feedback loops to first bring the machine to a proper operating condition and maintain it there. Only then could iterative refinement of the plasma profile take place, and only that specific problem within the incredibly diverse and complex tokamak could be approached at all as

the RL agent was specifically trained for that one task via large numbers of simulations [23].

We recommend that interested readers take a look at a great survey on RL, its connections to control theory, and existing challenges, which makes the following statement, which is very relevant to the preceding discussion: "In the adaptive setting, we want to learn the policy online. We only get one trajectory. The goal is, after a few steps, to have a model whose reward from here to eternity will be large. This is very different and much harder than what people are doing in RL. In episodic RL, you get endless access to a simulator. In adaptive control, you get one go" [24].

Ideas are diffusing at an increasing rate between the control theory and ML communities, with many researchers now working within the intersection of these fields. The following is a nonexhaustive list of examples illustrating some of these efforts. Researchers have incorporated neural networks into hybrid multiagent control barrier function-based controllers combining offline and online learning with radial basis function neural networks, with network weights updated online with new observations [25]. This approach allows the agents to ensure safety by taking advantage of the neural network-based predictions of other agents' dynamics. A distributional RL approach has been developed that enhances the understanding of the effects of the randomness in the environment by letting agents learn the distribution of a random return including a model-free approach for estimating the return distribution for unknown models [26]. A model-free RL algorithm has also been developed that enables the use of linear temporal logic (LTL) to formulate a goal for unknown continuous-state/action MDPs, and for synthesizing a corresponding optimal policy maximizing the probability of satisfying the given LTL property [27]. A method has been developed that incorporates robust adaptive latent space feedback control within the architectures of generative variational autoencoder (VAE) and diffusion-based models to make them more robust for extrapolating beyond the training data for complex time-varying dynamic systems [28]. A robust kernelized data-enabled predictive control algorithm has been developed to perform model-free optimal control for nonlinear systems using only input and output data by combining kernel methods with data-enabled predictive control [29]. A method has also been developed using a neural state-space model to approximate nonlinear systems, where a deep encoder network learns the nonlinearity from data and a state-space component captures the temporal relationship, to enable reference tracking in uncertain nonlinear systems [30]. A method of safety verification and robustness analysis has been developed for neural networks against norm-bounded perturbations in their input based on semidefinite programming and quadratic constraints [31]. A method has also been developed extending

aggregation graph neural networks to time-varying signals and time-varying network support to learn a common local controller that exploits information from distant teammates using only local communication interchanges for control of large networks of mobile robots with interacting dynamics and sparsely available communications [32]. A general safety framework based on Hamilton–Jacobi reachability methods that can work in conjunction with an arbitrary learning algorithm has been developed that exploits approximate knowledge of the system dynamics to guarantee constraint satisfaction while minimally interfering with the learning process [33]. DeepReach is a recently developed method that leverages new developments in sinusoidal networks to construct a neural partial differential equation solver for high-dimensional reachability problems [34].

GENERATIVE DEEP LEARNING

Generative deep learning is a subset of ML in which deep neural networks are used to generate new objects that are different from what was included in their training data. Technically, what these generative models are doing is high-dimensional interpolation and extrapolation, usually with a stochastic component built in to increase the diversity of the generated data.

LLMs are a great example of this. Given a set of text, LLMs are trained to statistically predict the next most probable token (for this discussion, one can think of tokens as words). Once trained, LLMs are capable of creating new sentences and stories that are outside of the training dataset.

Uncertainty Quantification

Before moving on to more complex image-based generative models, we consider regression, as described in the great work of Nix and Weigand from 1994 for uncertainty quantification in ML [35]. Given a noisy dataset $D = \{(x_i, y_i) \in \mathbb{R} \times \mathbb{R}\}_{i=1}^N$, we fit a function by assuming the data are coming from a normal distribution whose mean value is a function of x . A simple approach is to first assume a uniform standard deviation σ for the data. One then defines a function $f_\mu(x)$ parameterized by weights \mathbf{w} , such as a polynomial or a dense neural network (see “Dense Neural Networks”)

$$f_\mu(x, \mathbf{w}) = \sum_{n=0}^m w_n x^n$$

and a probabilistic model

$$p(y|x, \mathbf{w}) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{y - f_\mu(x, \mathbf{w})}{\sigma}\right)^2\right]. \quad (4)$$

The model is fit by adjusting the weights \mathbf{w} to maximize the likelihood of the observed data, assuming that the data are independent and identically distributed:

$$\begin{aligned} \mathbf{w}^* &= \operatorname{argmax}_{\mathbf{w}} P(D|\mathbf{w}) \\ &= \operatorname{argmax}_{\mathbf{w}} \prod_{i=1}^N p(y_i|x_i, \mathbf{w}) \\ &= \operatorname{argmax}_{\mathbf{w}} \prod_{i=1}^N \exp\left[-\frac{1}{2}\left(\frac{y_i - f_\mu(x_i, \mathbf{w})}{\sigma}\right)^2\right] \end{aligned} \quad (5)$$

where the constant term $1/\sigma\sqrt{2\pi}$, which does not influence the optimization, has been removed.

Because probabilities lie within the range $[0, 1]$, a large product of this type quickly becomes incredibly small, causing numerical problems. One solution is to take the negative log of (5) to also convert it to a minimization problem, which is now the negative log-likelihood (NLL)

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} -\log\left[\prod_{i=1}^N \exp\left[-\frac{1}{2}\left(\frac{y_i - f_\mu(x_i, \mathbf{w})}{\sigma}\right)^2\right]\right].$$

Expanding the logarithm of a product as a sum gives

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \frac{1}{2\sigma^2} \sum_{i=1}^N (y_i - f_\mu(x_i, \mathbf{w}))^2.$$

Again, removing the constant term $1/2\sigma^2$ and dividing by the number of data points N to make the procedure less sensitive to the size of the dataset results in the typical mean squared error loss

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \frac{1}{N} \sum_{i=1}^N (y_i - f_\mu(x_i, \mathbf{w}))^2. \quad (6)$$

In this approach, the variance has been completely eliminated from the optimization process; it can be chosen as an arbitrary constant resulting in a probabilistic generative model. For a given data point x , the value $\hat{y}(x)$ is then randomly sampled from the normal distribution defined by plugging $f_\mu(x, \mathbf{w}^*)$ into (4). This approach can also be used as a deterministic generative model in which, associated with any point x , we simply choose the most likely estimate as $\hat{y}_j = f_\mu(x, \mathbf{w}^*)$.

This simple derivation is provided because most deep neural networks are, in fact, trained by minimizing a mean squared error, which is equivalent to interpreting them as generative models whose predictions are those of the mean of a normal distribution. A more powerful approach is, of course, to also model $\sigma(x)$ as a second function $f_\sigma(x)$ without assuming that it is uniform over the data. The NLL derivation follows as shown previously, but we can no longer simply remove $\sigma(x)$ from the optimization; thus, the probabilistic model becomes

$$p(y|x, \mathbf{w}) = \frac{1}{f_\sigma(x, \mathbf{w}_2)\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{y - f_\mu(x, \mathbf{w}_1)}{f_\sigma(x, \mathbf{w}_2)}\right)^2\right] \quad (7)$$

Dense Neural Networks

Fully connected or dense neural networks act directly on n -dimensional vectors $\mathbf{x} \in \mathbb{R}^n$. Typically, each of the k layers of a dense neural network has $N_k \gg 1$ nodes $\{n_i^k\}_{i=1}^{N_k}$ each with its own bias and vector of weights \mathbf{w}_i^k , where N_k is referred to as the width of the layer. Considering a single input node n_i^1 of a dense neural network, it typically performs an inner product of the input vector \mathbf{x} with a vector of input weights $\mathbf{w}_i^1 \in \mathbb{R}^n$, adds the bias term $b_i^1 \in \mathbb{R}$, and then applies a nonlinear activation function a_i^k to the result, producing the output

$$n_i^1(\mathbf{x}) = a_i^1(\langle \mathbf{x}, \mathbf{w}_i^1 \rangle + b_i^1). \quad (\text{S1})$$

Applying all N_1 input nodes in parallel to the single input vector \mathbf{x} results in the N_1 -dimensional vector

$$\mathbf{v}^1(\mathbf{x}) = [n_1^1(\mathbf{x}), n_2^1(\mathbf{x}), \dots, n_{N_1}^1(\mathbf{x})].$$

This vector is then passed to the next layer, whose nodes perform similar operations, producing outputs of the form

$$n_i^2(\mathbf{v}^1(\mathbf{x})) = a_i^2(\langle \mathbf{v}^1(\mathbf{x}), \mathbf{w}_i^2 \rangle + b_i^2).$$

Note that, at this step, having used just one single layer, if we use just one single node n_i^2 and the identity function $a_i^2(x) = x$ as our activation, the inner product gives us

$$n_i^2(\mathbf{v}^1(\mathbf{x})) = \sum_{i=1}^{N_1} a_i^1(\langle \mathbf{x}, \mathbf{w}_i^1 \rangle + b_i^1) + b_i^2 \quad (\text{S2})$$

which is a linear combination of functions applied to \mathbf{x} . If we choose the a_i^1 as sines and cosines, this is equivalent to a standard Fourier series. If we instead choose the a_i^1 as x^i , this is a polynomial expansion. Basically, we are creating functions out of arbitrary linear combinations of whatever basis functions we desire. Therefore, when \mathbf{x} is confined to some compact set $K \subset \mathbb{R}^n$, it follows from standard functional analysis results that we can approximate any measurable function arbitrarily accurately over the compact set K by using a sufficiently large number of functions. This is in essence the well-known universal approximation theorem of neural networks and demonstrates why they are useful for such a wide range of systems. For example, compared to a classical Fourier/Taylor series performing a trigonometric/polynomial function approximation, a neural network using the popular rectified linear unit $\text{ReLU}(x) = \max\{0, x\}$ activation functions will perform a piecewise affine function approximation. Other activation functions can be understood as smoothing the ReLU.

In practice, what works much better than trying to use a single incredibly wide layer is to use a deep neural network that has many layers that can be interpreted as compositions of many functions. Considering a dense network's second layer,

applying all N_2 nodes of the second layer in parallel to the vector $\mathbf{v}^1(\mathbf{x})$ results in an N_2 -dimensional vector:

$$\mathbf{v}^2(\mathbf{v}^1(\mathbf{x})) = [n_1^2(\mathbf{v}^1(\mathbf{x})), n_2^2(\mathbf{v}^1(\mathbf{x})), \dots, n_{N_2}^2(\mathbf{v}^1(\mathbf{x}))].$$

This process can be continued indefinitely, depending on how many layers deep the dense neural network is, where the dimension of each representation \mathbf{v}^i is defined by the width of the layer. A high-level summary of this architecture is shown in Figure S1.

It is found that, by using deep networks (tens of layers), useful intermediate latent representations are discovered by the networks, automatically resulting in their state-of-the-art performance in regression and classification.

Given a dataset of inputs and outputs $D = \{\mathbf{x}_i, \mathbf{y}_i\}$, a dense neural network is trained by stochastic gradient descent applied to the values of the weights and biases to minimize a cost function that compares predicted outputs $\hat{\mathbf{y}}_i$ to their true values over batches of N_b data examples. Typically, the cost function is something like the mean squared error

$$C = \frac{1}{N_b} \sum_{i=1}^{N_b} (\hat{\mathbf{y}}_i(\mathbf{x}_i) - \mathbf{y}_i)^2. \quad (\text{S3})$$

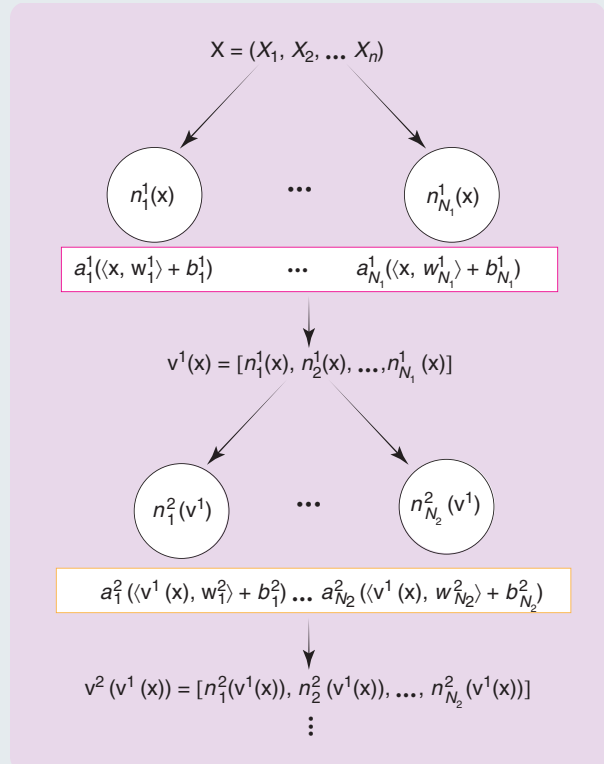


FIGURE S1 A high-level overview of the first two layers of a dense neural network is shown. Each subsequent layer can be thought of as a general function of the preceding layer's output.

where $\mathbf{w} = \{\mathbf{w}_1, \mathbf{w}_2\}$ is now two sets of weights that parameterize the two unknown functions. The overall NLL calculation results in

$$\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmin}} \left\{ -\frac{1}{N} \sum_{i=1}^N \log \left(\frac{1}{f_\sigma(x, \mathbf{w}_2) \sqrt{2\pi}} \right) + \frac{1}{N} \sum_{i=1}^N \frac{(y_i - f_\mu(x_i, \mathbf{w}_1))^2}{2f_\sigma^2(x, \mathbf{w}_2)} \right\}. \quad (8)$$

The preceding example is a generative model developed by statistical regression, which is referred to as a deep generative model if the functions $f_\mu(x, \mathbf{w}_1)$ and $f_\sigma(x, \mathbf{w}_2)$ are deep neural networks. One neural network-based example of this approach for a noisy dataset is shown in Figure 3.




FIGURE 3 Example of a statistical generative model fit to a noisy dataset.

CNNs

When working with images instead of vectors, it is also possible to utilize deep dense neural networks. A digital image is usually represented by an object of the form $X \in \mathbb{R}^{n_x \times n_y \times n_c}$ with n_x and n_y representing pixel numbers, where n_c is the number of channels. Typically, images have three color channels of red, green, and blue (RGB). It is possible to unwrap such an image into a vector \mathbf{x} of length $n_x n_y n_c$; however, this is usually not done. Even a relatively low-resolution 100×100 pixel black-and-white image would then become a 10,000D vector for which just a single input node would require an equally large 10,000D vector of weights for the dot product shown in (S1). Since the layers of dense neural networks typically contain hundreds of nodes, the size of just a single-layer neural network would quickly explode.

Despite the high dimensionality of the problem mentioned previously, the real reason that dense neural networks are not used for images is that they are suboptimal for this task as unwrapping an image into a vector destroys all of the spatial relationships that neighboring pixels share. CNNs are a different form of deep neural networks, ones that can incorporate spatial correlations. CNNs were first introduced by K. Fukushima in 1980 for the classification of images of letters and numbers, with their architecture inspired by the hierarchy model of the visual nervous system [S1]. CNNs then grew in popularity in the 1990s with the work of Y. LeCun, who utilized backpropagation to train a deep CNN to identify handwritten zip codes for the U.S. Postal Service [S2]. In 2012, there was a major breakthrough as the deep CNN AlexNet was created by A. Krizhevsky, I. Sutskever, and G. Hinton and won the ImageNet competition of classifying large ($224 \times 224 \times 3$ pixels) images into 1,000 classes [S3]. CNNs have dominated the ImageNet competition ever since and it is now considered solved.

The main idea of a CNN is for a single node to utilize an $n_f \times n_f$ matrix W as a filter that is traversed over an image with zeros padded on at the edges. For example, if using a 3×3 filter, at each position in the image at pixel (i, j) , the following convolution takes place over the entire 3×3 neighborhood containing that pixel to generate its new representation by adding a bias and applying an activation function:

$$Y(i, j) = a \left(\sum_{k:1:3} \sum_{l:1:3} X(i-2+k, j-2+l) W(k, l) + b \right). \quad (S4)$$

As with the dense neural networks described previously, a large number of filters can be used in parallel with networks that are many layers deep, with each subsequent image represented by the CNN having as many channels as there are filters.

Each subsequent layer of the CNN learns how to extract useful features directly from the intermediate images. Furthermore, by jumping the filters more than one pixel at a time, subsequent images decrease in size. For example, if a stride of 2 is used, then the width and height of each subsequent image is decreased by a factor of 2. For example, AlexNet uses such a procedure with strides of varying lengths at different layers as well as MaxPool layers that replace a neighborhood of pixel values with the value of the maximum pixel intensity to eventually decrease the initial 227×227 pixel images with three color channels down to 6×6 pixel images with 256 channels before being reshaped and passed through a dense neural network to final representation as a 1,000D vector.

While AlexNet solves a classification problem with position i of the 1,000D output vector representing the probability of the input image being of class i , it can be thought of as a compression technique in which the original $227 \times 227 \times 3 = 154,587$ D

(Continued)

CNNs (Continued)

objects are compressed by a factor of approximately 155, down to 1,000D vectors. AEs, which are described next, utilize such an approach to compress and then regenerate high-dimension objects.

As an illustrative example, Figure S2 shows 10 input images from the Modified National Institute of Standards and Technology (MNIST) dataset (28×28 pixel handwritten digits) together with their embedded representations within a small CNN that is performing classification to convert them into one-hot encoded vectors of length 10, where after training, position i of the vector should be 1 and all others 0 when acting on an input image of the digit i . It can be seen that initially there is

some edge detection going on before the network builds more abstract representations in subsequent layers.

REFERENCES

- [S1] F. Kuniyiko, "Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position," *Biol. Cybern.*, vol. 36, no. 4, pp. 193–202, 1980, doi: [10.1007/BF00344251](https://doi.org/10.1007/BF00344251).
 [S2] L. Yann, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998, doi: [10.1109/5.726791](https://doi.org/10.1109/5.726791).
 [S3] K. Alex, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," 2012. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf




FIGURE S2 One example of each digit from the MNIST dataset is used as input to a classification CNN, showing the digits as well as their internal representations within the network as a CNN converts the images into 10D vectors that predict their numerical values.

AEs

AEs are neural networks that compress high-dimensional data into a low-dimensional latent embedding from which the data are then recreated. When working with images, AEs typically utilize 2D CNNs to first compress the images down to a low-dimensional latent representation and then build them back up. One of the first uses of 2D CNNs as AEs was published in 2007 [S4].

Consider a CNN-based encoder whose input is an $n_x \times n_y$ -pixel RGB image. This is a high-dimensional object even for low-resolution images because $X \in \mathbb{R}^{n_x \times n_y \times n_c}$. For example, a 100×100 -pixel black-and-white image is a 10,000D object. AEs are designed to greatly reduce the dimensionality of input images, to encode them down into a low-dimensional

latent space representation $\mathbf{z}_\mu = \text{CNN}_{\text{en}}(X) \in \mathbb{R}^{n_z}$, where $n_z \ll n_x \times n_y \times n_c$. The latent embedding is followed by a CNN-based generative decoder that takes inputs from the low-dimensional latent space and is trained to reproduce estimates $\hat{X} = \text{CNN}_{\text{gen}}(\mathbf{z}_\mu)$ of the original high-dimensional inputs. The AE process is

$$X \rightarrow \mathbf{z}_\mu = \text{CNN}_{\text{en}}(X) \rightarrow \hat{X} = \text{CNN}_{\text{gen}}(\mathbf{z}_\mu). \quad (\text{S5})$$

Training takes place by adjusting the AE's weights and biases via stochastic gradient descent to minimize a loss function

(Continued)

AEs (Continued)

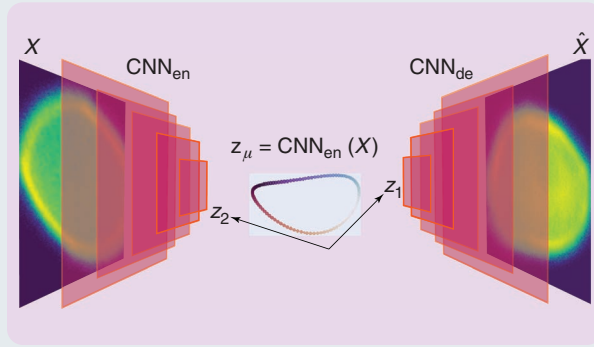


FIGURE S3 An AE compresses a high-dimensional object, such as an image X , down into a low-dimensional latent embedding \mathbf{z}_μ that the decoder neural network then transforms into an estimate of the original image \hat{X} . This is a deterministic process; for a given input, the latent space location and the output are always the same, as determined by the trained neural network. Therefore, if an input image is perturbed in a periodic manner, its latent representation will also perform a periodic motion, as shown here.

EMBEDDING INTO LATENT SPACE

In practice, the strength of deep generative models comes from their incredible complexity as they are composed of millions of nonlinear activation functions, weights, and biases. This article will focus on two abilities of these large generative models that emerge as their scale is increased and that bifurcate their abilities, clearly separating them from traditional statistical generative approaches. The first feature that sets large generative deep learning models apart from standard statistics approaches is their ability to extend the preceding simple generative process, which works for numbers and vectors, to work directly with high-dimensional numerical representations of text and images. The second major strength that this article focuses on is the powerful data compression ability of these models. To explore these capabilities, first consider a convolutional neural network (CNN)-based (see “CNNs”) autoencoder (AE).

The AE

An AE can map a large image into a low-dimensional latent embedding, which is typically a vector $\mathbf{z} \in \mathbb{R}^{n_z}$ whose dimension n_z is much smaller than the original input image, as described in “AEs.” For example, a 100×100 -pixel image, which is a 10,000D object when considering every pixel, can typically be compressed and regenerated with high accuracy from a latent embedding of dimension $n_z \sim 10$. This compression is possible because AEs are composed of complex nonlinear deep neural networks, and so they can compress to much lower dimensions than linear methods, such as principal component analysis (PCA). In

that is a metric of the difference between the generated and true images, such as the mean squared error averaged over a batch of images of size N_b :

$$C(\mathbf{w}, \mathbf{b}) = \frac{1}{N_b} \sum_{j=1}^{N_b} \left[\frac{1}{n_x n_y n_c} \sum_{i_x, i_y, i_c} (X_{j, i_x, i_y, i_c} - \hat{X}_{j, i_x, i_y, i_c})^2 \right] \quad (\text{S6})$$

where \mathbf{w} are the neural network weights, \mathbf{b} are the neural network biases, the (i_x, i_y) are pixel positions, and i_c is the channel number.

Figure S3 shows a high-level overview of an AE that is encoding images into a 2D latent space and then attempting to reconstruct them. Because the AE is a collection of deterministic functions, if the input image is varied periodically, then the latent embedding will also follow a periodic trajectory.

REFERENCES

[S4] M. Ranzato, F. J. Huang, Y.-L. Boureau, and Y. LeCun, “Unsupervised learning of invariant feature hierarchies with applications to object recognition,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Minneapolis, MN, USA, 2007, pp. 1–8, doi: [10.1109/CVPR.2007.383157](https://doi.org/10.1109/CVPR.2007.383157).

fact, for a given embedding dimension n_z , a purely linear AE without any nonlinear activation functions can be shown to recover the largest n_z principal components that would be generated by PCA [36], [37], [38].

As previously described relative to statistical regression, the encoder neural network of the AE maps an image X to the mean vector of a statistical multivariate normal distribution model according to

$$\text{NN}_{\text{en}}(X, \mathbf{w}_{\text{en}}) = \mathbf{z}_\mu(X, \mathbf{w}_{\text{en}})$$

and the generative decoder neural network then attempts to recreate the image from this vector according to

$$\hat{X} = \text{NN}_{\text{de}}(\mathbf{z}_\mu, \mathbf{w}_{\text{de}})$$

where \mathbf{w}_{en} and \mathbf{w}_{de} are the weights and biases of the encoder and decoder neural networks, respectively.

The VAE

As discussed previously, it is also possible to include a non-zero variance in the latent statistical model. In this case, the encoder neural network splits into two branches to generate

$$\text{NN}_{\text{en}}(X) = [\mathbf{z}_\mu, \mathbf{z}_\sigma]$$

and models the probabilistic latent representation as

$$p(\mathbf{z}|X, \mathbf{w}) = \frac{(2\pi)^{-n_z/2}}{\sqrt{\det(\Sigma)}} \exp\left[-\frac{1}{2}((\mathbf{z} - \mathbf{z}_\mu)^T \Sigma (\mathbf{z} - \mathbf{z}_\mu))\right] \quad (9)$$

where Σ is a diagonal matrix with entries $\mathbf{z}_\sigma(X, \mathbf{w}_2)$, as described in “VAEs.”

If an AE is trained without regularization, the network may place latent embeddings over a very wide range of the latent space, which can make interpolation within the latent space difficult as large regions must be traversed that have never been seen by the decoder neural network. In such cases, the decoder may generate poorly behaved outputs, which would make any kind of continuous adaptive movement in the latent space more difficult. To alleviate this, a simple solution is to perform

regularization in the form of an additional term in the AE optimization cost function that penalizes the magnitudes of the created latent vectors according to $\epsilon \sum_{X_i \in D_b} |\mathbf{z}_\mu|^2$, where D_b is a batch of the training data and $\epsilon > 0$ must be chosen small enough so as to not force the network’s representation to collapse so much that it is unable to recreate images. The same approach of weighing various parts of the cost function must be considered for VAEs as the default weight of 1 on the KL part of the cost function may cause the network’s latent representation to collapse.

VAEs

Another compression approach is to train an encoder–decoder pair as a VAE. This approach typically doubles the size of the encoder’s output so that it produces the mean value of the latent embedding location \mathbf{z}_μ as in the preceding case, and another vector of the same size \mathbf{z}_σ , which represents the diagonal covariance matrix of an n_z -dimensional Gaussian distribution [S5]. (See Figure S4.)

There is then an additional step to the training process. Instead of simply generating samples from their latent representations \mathbf{z}_μ , random noise is added to each component of \mathbf{z}_μ by sampling from a random n_z -dimensional Gaussian that is centered at \mathbf{z}_μ , whose diagonal covariance matrix is defined by \mathbf{z}_σ . The VAE process is

$$\begin{aligned} X &\rightarrow (\mathbf{z}_\mu, \mathbf{z}_\sigma) = \text{CNN}_{\text{en}}(X) \\ &\rightarrow \hat{\mathbf{z}}_\mu = \mathbf{z}_\mu + \mathbf{z}_\sigma \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \\ &\rightarrow \hat{X} = \text{CNN}_{\text{gen}}(\hat{\mathbf{z}}_\mu). \end{aligned} \quad (\text{S7})$$

The cost function used to optimize the weights and biases of the VAE now has two parts. The first part is a measure of reconstruction accuracy, as defined in (S6). The

second part of the cost function is the Kullback–Leibler (KL) divergence between a mean zero unit variance multivariate Gaussian probability distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and the multivariate Gaussian distribution parameterized by the two branches of the encoder $\rho(\mathbf{z}|X)$:

$$D_{\text{KL}} = \frac{1}{2} [\text{Tr}[\Sigma] + \mathbf{z}_\mu^T \mathbf{z}_\mu - \log[\det(\Sigma)]] \quad (\text{S8})$$

where Σ is a diagonal covariance matrix with entries \mathbf{z}_σ .

From (S7), it is apparent that an AE is just a VAE with $\mathbf{z}_\sigma = \mathbf{0}$. Because the latent space points of a VAE are perturbed by random noise before being passed to the decoder, the decoder becomes more robust. Furthermore, the KL divergence pushes the learned probabilistic latent space distribution toward $\mathcal{N}(\mathbf{0}, \mathbf{I})$, which results in inputs being placed closer together in the latent space, which removes large gaps and helps the decoder to interpolate over the latent space when moving between points that were seen in the training data.

Rather than using a probabilistic loss term, such as the KL divergence on the latent space, it is possible to simply introduce regularization on the network weights and an explicit L2 norm on the latent space vector, which forces the encoder to place points into a smaller, more compact set than a vanilla AE. This helps to create a latent space in which we can move around without large gaps between latent embedding points to help the generator create realistic objects as we move.

Whether we randomly perturb the latent points as in a VAE approach or choose them by moving along a parameterized trajectory as in our AE approach, in both cases, we sample new points of the latent space, and the generative decoder CNN_{gen} generates new outputs. A more thorough study of deterministic generative models and these ideas is provided in [S6].

REFERENCES

- [S5] D. P. Kingma, “Auto-encoding variational Bayes,” 2013, *arXiv:1312.6114*.
- [S6] P. Ghosh, M. S. Sajjadi, A. Vergari, M. Black, and B. Schölkopf, “From variational to deterministic autoencoders,” 2019, *arXiv:1903.12436*.




FIGURE S4 A VAE is an AE in which the low-dimensional latent embedding is treated as the parameterized probability density function: a multivariate Gaussian whose mean vector and diagonal covariance matrix are determined by two branches of the encoder neural network. During both training and inference, samples are randomly sampled from the learned distribution before the decoder neural network uses them as inputs.

AE Versus VAE

As an illustrative example, both an AE and a VAE, with various regularization and Kullback–Leibler (KL) weights were trained on a toy problem in which 128×128 -pixel images of rectangles are embedded into a 2D latent space and reconstructed. Intuitively, the networks could arrange the latent space coordinates as rectangle width and height (w, h) . The networks are trained with 1,024 randomly generated images. Examples of four random reconstructions by the AE with and without regularization are shown in Figure 4.

Locations of the 2D latent embeddings z_μ are shown for 1,024 input images of various heights and widths for an AE and a VAE in which rectangles are embedded into the latent space and then regenerated by the decoder in Figure 5. For the AE, results are shown with and without L2 regularization on the latent vector length. In both cases, the latent embedding is nicely arranged relative to both the width and height of the input rectangle, but the embeddings of the AE without regularization are very widely spread. When an L2 penalty with a $1e-2$ weight is used, the network learns a compact latent representation for which the decoder is better able to handle movement in the latent space. For the VAE, results are shown with a default 1 weight on the KL divergence as well as $1e-2$ and $1e-3$ weights. The lowest weight of $1e-3$ is the only one that gives the VAE enough freedom to arrange the latent space in a reasonable way. The black lines show the path along which input rectangle images are embedded with their width and height (w, h) interpolated from $(1, 1)$ to $(11, 8)$ over 1,000 steps. The path for the VAE with KL weight 1 is not shown as it is randomly scattered over the entire space. The reconstruction mean percent error for points along the black latent paths of Figure 5 is shown in Figure 6 on a log scale, where the mean percent error for reconstruction \hat{X} of image X is calculated as

$$100 \times \frac{\sum_{i_x, i_y, i_c} |X_{i_x, i_y, i_c} - \hat{X}_{i_x, i_y, i_c}|}{\sum_{i_x, i_y, i_c} |X_{i_x, i_y, i_c}|}.$$

Robustness of Image Compression

From Figures 5 and 6, it is clear that the default approach to AEs, without latent vector length regularization, results in latent representations that are very widely distributed. Including a small regularization term contracts the latent embedding, resulting in a latent space that is easier to traverse with higher accuracy reconstructions. The default approach to VAEs has the opposite problem, as a default gain of 1 or a gain as low as $1e-2$ on the KL divergence term in the cost function leads the VAE to overly compress the latent space, making it very difficult for the decoder to distinguish between different inputs. In this example, once

the KL gain is decreased to $1e-3$, the performance of the VAE matches that of the regularized AE, and their latent embeddings take on almost identical shapes that are only rotated and scaled relative to each other.

It is not surprising that AEs and VAEs can compress images of rectangles down into a 2D representation, as each rectangle is defined by its 2D coordinates of width and height (w, h) . To clearly demonstrate the power of AE- and VAE-based nonlinear compression, consider the Modified National Institute of Standards and Technology (MNIST) dataset, a large collection of 28×28 -pixel images of handwritten digits 0–9. First, PCA is performed on the set of 60,000 MNIST training images by unwrapping each image as a $1 \times 28^2 = 784$ D vector and then stacking the images into a $60,000 \times 784$ data matrix D . PCA is performed by first factorizing the data matrix D via singular value decomposition:

$$D = U\Sigma W^T$$

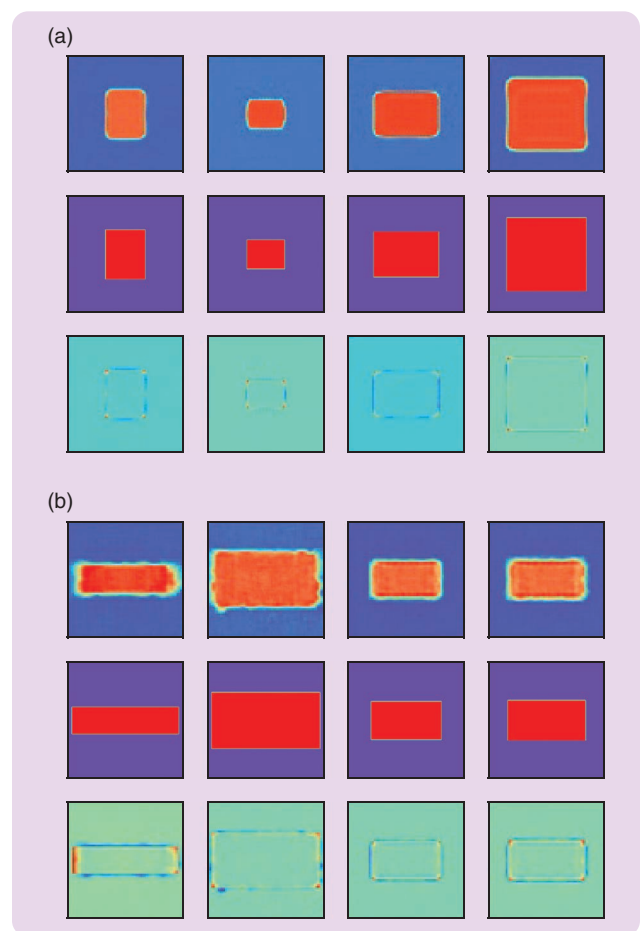


FIGURE 4 (a) Top row shows AE reconstructions with $1e-2$ latent regularization, middle row shows the true images, and the bottom row is the difference. (b) AE reconstructions without regularization. With a more compact latent embedding to work with, the decoder of the regularized AE is having an easier time reconstructing the rectangles more accurately.

where Σ is a diagonal matrix of singular values of D , and U and W are the left and right singular vectors.

A score matrix T is then calculated according to

$$T = DW = U\Sigma W^T W = U\Sigma$$

where T is now a 784×784 matrix whose components can be reshaped into 28×28 basis images whose linear combinations can be used to generate the original images. An image X in D can be decomposed as

$$X = \sum_{i=1}^{28^2} c_i T_i, \quad \mathbf{c} = XW$$

where T_i rows of T are reshaped into 28×28 -pixel images.

PCA can then be used for dimensionality reduction by keeping only the right singular vectors associated with the N largest singular values of D in a matrix W_N and performing the transformation

$$\mathbf{x}^{\text{PCA},N} = XW_N. \quad (10)$$

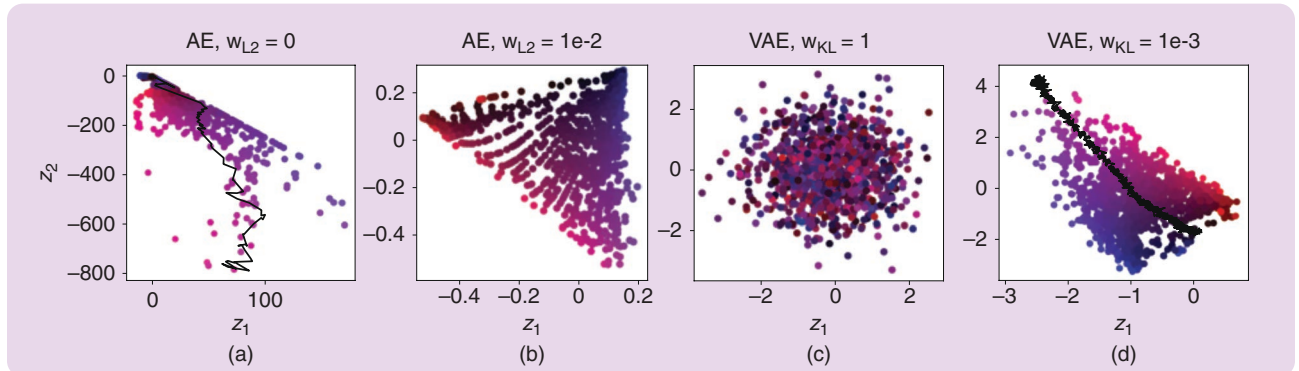


FIGURE 5 Locations of the 2D latent embeddings \mathbf{z}_μ are shown for 1,024 input images of various heights and widths for an AE and a VAE in which rectangles are embedded into the latent space and then regenerated by the decoder. The latent space points are colored by RGB settings with the green channel set to zero, red proportional to rectangle width, and blue proportional to rectangle height. For the AE, results are shown with and without L2 regularization on the latent vector length. In both cases, the latent embedding is nicely arranged relative to both the width and height of the input rectangle, but the embeddings of the AE without regularization are very widely spread. When an L2 penalty with a $1e-2$ weight is used, the network learns a compact latent representation for which the decoder is better able to handle movement in the latent space. For the VAE, results are shown with a default 1 weight on the KL divergence as well as $1e-2$ and $1e-3$ weights. The lowest weight of $1e-3$ is the only one that gives the VAE enough freedom to arrange the latent space in a reasonable way. The black lines show the path along which input rectangle images are embedded as their width and height (w, h) are interpolated from (1, 1) to (11, 8) over 100 steps. The path for the VAE with KL weight 1 is not shown as it is randomly scattered over the entire space.




FIGURE 6 The reconstruction error is shown for input images whose latent embeddings move along the black curves shown in Figure 5 for an AE and a VAE with different levels of regularization.




FIGURE 7 Cumulative sum of the explained variance over the MNIST training dataset as a function of number of PCA components used to represent the data. One hundred fifty-four PCA components are required for 95% explained variance.

The original 28×28 image X can then be approximately reconstructed from the N -dimensional vector x_N by again using the components of the original score matrix T as basis images and constructing

$$\hat{X} = \sum_{i=1}^{N \leq 28^2} x_i^{\text{PCA}, N} T_i.$$

For the MNIST training data, it is found that at least 154 PCA components are required to achieve 95% explained variance of the data, as shown in Figure 7.

If we instead use a nonlinear CNN-based AE to embed the images into a low-dimensional latent space from which they are then regenerated by the decoder, it is found that they can be accurately represented with just a 2D latent embedding, where the dimension reduction is performed by the embedding neural network to create

$$\mathbf{z}(X) = (z_1(X), z_2(X)) \in \mathbb{R}^2 = \text{NN}_{\text{en}}(X).$$

Figure 8 shows the 2D locations of the 60,000 latent embeddings of a trained AE acting on the MNIST training dataset images. It can be seen that the AE has organized the 2D latent space into separate sections for each of the digits.

The reconstructed approximation of an image is then created by the generative decoder neural network acting on the latent embedding

$$\hat{X} = \text{NN}_{\text{de}}(\mathbf{z}(X)) = \text{NN}_{\text{de}}(\text{NN}_{\text{en}}(X)) \in \mathbb{R}^{28 \times 28}.$$

In Figure 9, the reconstruction abilities of the trained AE with a 2D latent space are compared to PCA-based reconstructions that use 10 and 154 PCA components for test images that were not part of the original training data for either method.

It is important to consider the advantages and disadvantages of both methods. The AE is able to compress to a much lower dimension latent space while still reconstructing realistic test images that were unseen during training, but this is only true if the new images are coming from the same distribution as the original data. If the AE is given a completely new image, it will compress it to some point in the latent space, and the trained decoder will then generate whatever is closest to that point in the latent space. PCA, on the other hand, is more general in this sense; the projection of (10), while limited to using the basis vectors created from the original dataset, will try to reproduce the input image as best as possible. This can be seen by taking both the trained AE and the PCA and presenting them with a new unique input image that is a Gaussian, as shown in Figure 10. While neither method can reproduce

it accurately as it is very different from images of numbers, PCA comes closer. The AE produces a perfect number 0 as the Gaussian is most similar in shape to this number, and so it is embedded into a latent position that coincides with zeros, which is then passed through the generative decoder.




FIGURE 8 Latent embeddings of the AE acting on 60,000 training images of the MNIST dataset colored by the value of each embedded digit.




FIGURE 9 Reconstructions of (a) 10 test images that were not part of the training data are shown when using (b) 10 PCA components, (c) 154 PCA components, and (d) an AE with a 2D latent space.




FIGURE 10 Comparison of the images produced by both PCA and the AE when acting on a Gaussian input image that is very different from anything seen in the training data. While using all available PCA components starts to approximate the Gaussian image, the AE simply sends it to a representation of 0.

What these generative models are doing is high-dimensional interpolation and extrapolation, usually with a stochastic component built in to increase the diversity of the generated data.

ADAPTIVE LATENT SPACE TUNING

The previous section was meant to provide an intuition for how to use AEs and VAEs to embed high-dimensional objects into useful low-dimensional embeddings that can then be easily and quickly traversed to enable the use of adaptive feedback mechanisms that work directly in the latent space.

In this section, we explore how such embeddings are enabling adaptive deep learning approaches for tracking the states of time-varying systems via virtual diagnostics, by which we mean estimates of the internal states of systems that may be inaccessible or difficult to measure, similar to what are referred to as state observers in the control theory community [41]. Unlike traditional state space observers, which are based on differential equations, ours are based on adaptively tuned deep neural networks. Specifically, we focus on particle accelerator-related applications.

Particle Accelerators

Particle accelerators are complex systems composed of thousands of coupled components, including resonant structures for beam acceleration and magnets for beam focusing and steering. The performance of these components drifts with time because of external disturbances, aging, and misalignment and shift of components during maintenance. Charged particle beams themselves are also complex objects evolving in a 6D position–momentum phase space (x, x', y, y', z, E) ; see “The 6D Phase Space of Charged Particle Beams.” These beams themselves have time-varying initial conditions that are due to time-varying beam source properties. Changing the initial beam conditions strongly influences the detailed 6D phase space distributions of beams along the entire accelerator, especially for intense beams whose evolution is governed by complex collective effects, such as space charge forces and coherent synchrotron radiation (CSR).

The 6D Phase Space of Charged Particle Beams

Many of the applications that we show in this article are focused on using adaptive generative AI to give us virtual views or estimates of the states of 2D projections of the 6D phase space of charged particle beams in large particle accelerators. Therefore, we give a very brief review of what this 6D phase space is, in which beam dynamics evolve, and which we hope to estimate to better control beam properties.

Particle accelerator beams are composed of many charged particles that are held together in precisely shaped bunches that are accelerated to precise design energies. This is a complex process, as even a low-charge 1-nC electron bunch contains more than 6 billion particles, whose mutually repulsive electromagnetic fields are constantly trying to rip the bunch apart. These beams evolve in a 6D phase space (x, y, z, p_x, p_y, p_z) , where (x, y, z) are particle positions and (p_x, p_y, p_z) are momentum components [S7]. Particle dynamics satisfy the relativistic Vlasov equation

$$\frac{\partial \rho}{\partial t} + \mathbf{v} \cdot \nabla_x \rho + \frac{\partial \mathbf{p}}{\partial t} \cdot \nabla_p \rho = 0, \quad \frac{\partial \mathbf{p}}{\partial t} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \quad (\text{S9})$$

where

$$\mathbf{v} = (\dot{x}, \dot{y}, \dot{z}), \quad \mathbf{p} = (p_x, p_y, p_z) = \gamma m \mathbf{v}$$

$$\gamma = 1 / \sqrt{1 - \frac{v^2}{c^2}}, \quad v = |\mathbf{v}|.$$

The dynamics of the electric and magnetic fields \mathbf{E} and \mathbf{B} , respectively, are coupled and depend on the current and charge densities \mathbf{J} and ρ , as described by Maxwell's equations

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0}, \quad \nabla \cdot \mathbf{B} = 0 \quad (\text{S10})$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad \nabla \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right). \quad (\text{S11})$$

The electric and magnetic fields include contributions from charges within the beam as well the external electromagnetic fields of accelerator components, such as radio-frequency resonant accelerating cavities, solenoids, dipoles, and quadrupole magnets [S8].

In high-energy particle accelerators, the z direction is usually defined parallel to the accelerator axis, and the p_z momentum component is many orders of magnitude larger than p_x and p_y . Therefore, it is common to consider the density function $\rho(x, y, z, x', y', E)$, where $(x', y') = (p_x/p_z, p_y/p_z)$ are the angles at which the particles are traveling relative to the accelerator axis.

REFERENCES

- [S7] M. Reiser, *Theory and Design of Charged Particle Beams*. New York, NY, USA: Wiley, 2008.
- [S8] L. D. Landua, *The Classical Theory of Fields*. Amsterdam, The Netherlands: Elsevier, 2013.

Particle accelerators are complex systems composed of thousands of coupled components, including resonant structures for beam acceleration and magnets for beam focusing and steering.

Space charge collective effects refer to the fact that, besides being influenced by external electromagnetic field sources, such as resonant cavities for acceleration and magnetic fields for shaping and turning, intense charged particle beam dynamics are also influenced by their own self-fields. The electric fields of the individual particles in a beam all repel each other, trying to rip the beams apart and introducing an energy spread. The magnetic fields of the beams have the opposite effect as the magnetic fields of beams traveling in parallel have a tendency to pull the particles back together. Space charge forces dominate at lower energies before the particles become highly relativistic. In highly relativistic beams, the forces from their own self-electric and magnetic fields start to cancel out, with highly relativistic beams sometimes referred to as being “frozen.”

For electron beams, this cancellation effect is very helpful as electron beams with as little as 5 MeV of kinetic energy are already highly relativistic. This energy of 5 MeV is the typical energy of the beams these days as they exit the beam source at the very beginning of the accelerator. For beams with higher mass, space charge is much more of an issue, as this field cancellation effect does not kick in until much later. When compared to electron beams, even much higher energy 800-MeV proton beams are only slightly relativistic because of their factor of 2,000 greater mass, and so they experience significant distortion and filamentation of their 6D phase space.

Furthermore, while space charge forces start to cancel for highly relativistic electron beams, this cancellation effect only happens in the direction transverse to beam travel, and incredibly short and intense electron beams can experience significant longitudinal space charge forces that result in an energy spread. At high energies, when space charge forces are not of concern for electron beams, the second collective effect of CSR starts to kick in.

Because of their low mass, highly relativistic electron beams asymptotically approach the speed of light. Therefore, the only way to compress the beams longitudinally, to make the bunches shorter, is to use a path difference trick, in which first a resonant cavity is used to introduce an energy chirp along the electron bunch so that particles near the head of the bunch are at slightly lower energy than particles at the end. The beam is then passed through a series of four dipole magnet chicanes. The first magnet typically curves the beam slightly to the left if we are flying along with the beam looking down the z-axis of the accelerator.

The second dipole magnet has the opposite effect, which straightens out the beam to again head straight down the z-axis. The third dipole magnet turns the beam back to the right, back toward the axis, and the final fourth magnet restores the beam trajectory to the center of the accelerator beam pipe. At each of these turns, the slightly higher energy particles at the end of the bunch are turned less by the same magnetic field because of their higher momentum, and so at the end of the chicane, the particles that were originally at the end of the beam have gone through a shorter path length than the slightly lower energy particles at the head of the bunch. If the beam energy and magnetic field strengths are set properly, the net result is that the beam is compressed into a shorter bunch by this path length difference trick.

Because of their incredibly low mass, as they accelerate through each of the turns in the chicane, the highly energetic electron beams will emit high-energy light, and as the electron bunch length itself is incredibly short, the light from different parts of the beam will overlap coherently and be amplified. This is the CSR effect. This light will then catch up to and impart forces on electrons within the bunch as they travel along the chicane, creating a new source of phase space distortion and energy spread in the beams.

Virtual Beam Diagnostics (State Observers)

While the beams and accelerators are both complex and time varying, as described previously, there is a lack of detailed noninvasive beam diagnostics beyond simple 1D measurements, such as beam position monitors, beam current monitors, or beam loss monitors. The result is that accelerator beams are difficult to tune up after a maintenance period or an outage (typically requiring weeks of effort by large teams), and they are difficult to retune between different experiments even when the machines are running. Besides the large amount of time wasted in tuning the machines, they are also typically running at suboptimal states, with operators chasing time-varying drifts to decrease beam loss by tuning tens to hundreds of parameters by hand, based on experience.

The most detailed single-shot beam measurements that are available are typically the (x, y) or (z, E) 2D projections of a beam's 6D phase space distribution. The (x, y) 2D projection can be collected by inserting a scintillating screen in the beam's path. The intensity and repetition rate of the beams that can be measured with such diagnostics are

limited to prevent damage of these screens. The state-of-the-art method for measuring a beam's (z, E) distribution utilizes an X-band transverse deflecting cavity (XTCAV) to measure the beam. The XTCAV streaks a charged particle bunch, translating longitudinal position to transverse position. The sheared bunch is then passed through a vertical dipole, causing an energy-dependent curvature of the electron trajectory. The beam, which has been significantly spread in terms of longitudinal position (z) and energy (E) , impacts a scintillating screen, and the generated light intensity is recorded, providing a measurement of the (z, E) longitudinal phase space density that captures both the longitudinal bunch current profile and the energy distribution [39].

Other 2D projections of charged particle beams are also regularly measured, such as the (x, x') and (y, y') distributions, but these are very time-consuming processes that require dedicated measurement times and setups; they are not single-shot measurements. To precisely control charged particle beams, it would be very useful to have a full view of all possible 2D projections of the beam's 6D phase space. If we consider all of the unique 2D projections, there are 15 of them $(x, y), (x, x'), \dots, (z, E)$. One example of all 15 unique projections of a beam's 6D phase space is shown in Figure 11.

The problem of creating a virtual view of this partially observable system, where we would like to generate all 15 2D projections of the beam's 6D phase space based on observations of only one or two of the projections, is ideally suited for a combination of adaptive feedback and generative ML. The approach is to create an AE-type model that is not trained to recreate its inputs, but rather is trained to

map a combination of the phase space projections of the beam's time-varying initial conditions at the accelerator entrance, $X(t)$ as well as a vector of scalar accelerator and beam measurements, such as magnet and resonant cavity set points, $x(t)$ to all 15 projections $(y_1(t), y_2(t), \dots, y_{15}(t))$ of the beam's 6D phase space at a downstream location. Note that, although the AE and VAE frameworks described previously were always designed to recreate their own inputs from a low-dimensional latent embedding, the main point is that deep neural networks can create useful low-dimensional representations of complex objects, which is also known as *representation learning*.

In this approach, the encoder half of the network creates the latent embedding according to

$$\mathbf{z}(t) = \text{NN}_{\text{en}}(X(t), \mathbf{x}(t))$$

and the decoder half then maps this low-dimensional representation to estimates of the downstream phase space projections

$$\hat{Y}(t) = [\hat{y}_1(t), \dots, \hat{y}_{15}(t)] = \text{NN}_{\text{de}}(\mathbf{z}(t)).$$

Adaptive Virtual Diagnostics

The model is first trained in a typical supervised learning approach, but it is then applied in an unsupervised adaptive framework, where it is assumed that we no longer have access to a projection of the time-varying input beam $X(t)$ and we also do not have access to the time-varying vector of beam and accelerator parameters $\mathbf{x}(t)$. The idea is that

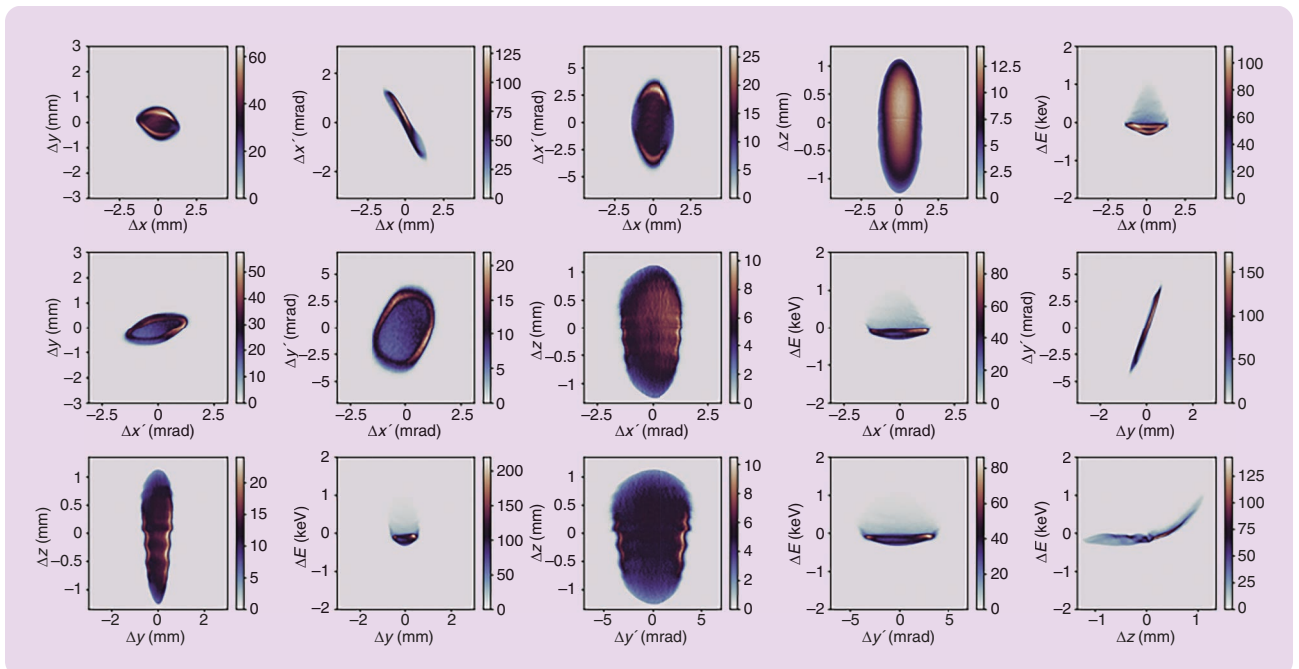


FIGURE 11 The 15 unique 2D projections of the simulated 6D phase space distribution of an electron beam in the HiRES compact accelerator. HiRES: High Repetition-Rate Electron Scattering apparatus.

we still want to track all 15 of the time-varying 2D downstream projections while only having access to one of them (such as the (z, E) or the (x, y) projection, depending on the accelerator).

The approach is adaptive tuning of $\mathbf{z}(t)$ by bounded ES [14]; see “Bounded ES for Nonlinear Systems With Unknown Control Directions.” It is worth noting that many forms of ES have been developed over the years. The major strength of this method is that it was shown to be able to stabilize and control unknown systems with unknown and time-varying control directions [42]. This method has been applied for finite-time control of static maps [43], for building observers or reduced-order models of coupled thermal and fluid systems [44], for high-voltage converter modulator rise-time optimization [45], for distributed control [46], for hybrid systems [47], for aircraft control [48], for piezoelectric control [49], for vibrational control [50], for event-triggered adaptive control [51], for heat exchanger network optimization [52], for spark-plug timing calibration [53], and for type 1 diabetes personalized zone adaptation [54].

The relatively recently developed bounded form of ES [14] is especially useful for real-world applications beyond academic studies because of its simplicity and analytically guaranteed bounds on the parameter update rates. Having known guaranteed bounds makes this method useful for in-hardware real-time control of powerful physical systems where incorrect or large abrupt parameter changes can result in system damage, such as in high-energy particle accelerators, where this method has been experimentally

demonstrated for beam loss minimization via magnet and radio-frequency (RF) tuning [55], [56], and has been studied for the development of optimal feedback control for unknown time-varying systems [57]. For a broad overview of recent developments of the ES approach, the reader is referred to the recent survey [58].

The latent vector \mathbf{z} is first given a fixed initial condition, which may come from the encoder acting on the last known measurement of X and \mathbf{x} . The components of \mathbf{z} are then adaptively tuned according to

$$\dot{z}_i(t) = \sqrt{\alpha\omega_i} \cos(\omega_i t + kC(y_j(t), \hat{y}_j(t))) \quad (11)$$

where $i = 1, \dots, n$ for the components of the latent vector $\mathbf{z} = (z_1, \dots, z_n)$, $j = 1, \dots, 15$ for the 15 unique 2D images y_1, \dots, y_{15} which represent the (x, y) , (x, p_x) , \dots , (z, p_z) 2D projections of a beam’s 6D position momentum (x, y, z, p_x, p_y, p_z) phase space distribution, and the cost function

$$C(y_j(t), \hat{y}_j(t)) = \iint (y_j(t) - \hat{y}_j(t))^2 \quad (12)$$

only compares the single measured projection $y_j(t)$ to its generated version $\hat{y}_j(t)$. It is worth pointing out that, even for typical low-resolution beam measurements, this is a very high-dimensional problem. A measured input beam image of just 52×52 pixels is a $\sim 2,700$ D object, and if the generative CNN creates all 15 projections of the beam’s 6D phase space at 256×256 pixels, the overall output is a ~ 1 million-dimension object that is generated as a single image with 15 different channels by the CNN. While

Bounded ES for Nonlinear Systems With Unknown Control Directions

ES is a method for optimizing output functions of stable systems by tuning control parameters [S9], for systems with delays [S10], with prescribed fixed-time convergence [S11], and with safety performance optimization while simultaneously respecting nonconvex analytically unknown safety barrier functions [S12].

A feedback-stabilizing form of ES was developed that is applicable not just for the optimization of stable systems but also for the stabilization and optimization of time-varying unstable systems [42]. Bounded ES is an incredibly simple (no need for families of various external filters) new form of stabilizing ES in which the control efforts have analytically guaranteed bounds [14]. Bounded ES has been studied for various control problems [S13], [S14], and its convergence properties have been proven for a wide range of dithers, including nondifferentiable and discontinuous controllers [S15]. The most general form of bounded ES has been studied for systems of the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t, u),$$

where $\mathbf{x} \in \mathbb{R}^n$ and \mathbf{f} is unknown [S16]. For clarity, this discussion focuses on the simplest form of bounded ES for control affine nonlinear systems of the form

$$\dot{\mathbf{x}}_i = f_i(\mathbf{x}, t) + g_i(\mathbf{x}, t)u_i \quad (S12)$$

$$u_i = \sqrt{\alpha\omega_i} \cos(\omega_i t + kV(\mathbf{x}, t)) \quad (S13)$$

where the functions $f_i(\mathbf{x}, t), g_i(\mathbf{x}, t): \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$ are unknown, and where $g_i(\mathbf{x}, t)$ may pass through zero, changing sign, such as $g_i(t) = \sin(\omega_i t)$, making this a difficult problem of stabilization under unknown and time-varying control direction.

$V(\mathbf{x}, t)$ depends on the situation. If the goal is stabilization of the unknown dynamic system, where one is free to choose V , it can be chosen as a Lyapunov-type function, such as $V = \mathbf{x}^T \mathbf{x}$. V may also be predefined as a noise-corrupted measurement

$$V(\mathbf{x}, t) = h(\mathbf{x}, t) + n(t) \quad (S14)$$

(Continued)

Bounded ES for Nonlinear Systems With Unknown Control Directions (Continued)

where $h(\mathbf{x}, t)$ is an analytically unknown output function of the system, and $n(t)$ is random noise. In either of the previous two scenarios, given a compact set $K \subset \mathbb{R}^n$ and any $\delta > 0$, there exists sufficiently large ω^* such that for all $\omega_i > \omega^*$ where the ω_i are distinct, the averaged and actual closed-loop dynamics satisfy

$$\dot{\bar{\mathbf{x}}}_i = f_i(\bar{\mathbf{x}}, t) - \frac{k\alpha}{2} g_i^2(\bar{\mathbf{x}}, t) \frac{\partial V(\bar{\mathbf{x}}, t)}{\partial \bar{\mathbf{x}}_i}, \quad \|\mathbf{x}(t) - \bar{\mathbf{x}}(t)\| < \delta. \quad (\text{S15})$$

Stabilization of the averaged system (S15) ensures practical stabilization within δ of the original system (S12).

A strength of this approach is that the unknown and time-varying control direction problem has been eliminated as $g_i^2 \geq 0$, so the system can be stabilized by a sufficiently large

choice of $k\alpha > 0$. Also, although acting on an unknown and noise-corrupted function, the control effort (S13) has analytically guaranteed bounds

$$|u_i| = |\sqrt{\alpha\omega_i} \cos(\omega_i t + kV(\mathbf{x}, t))| \leq \sqrt{\alpha\omega_i}. \quad (\text{S16})$$

The simplicity, robustness, model independence, and guaranteed control bounds of this approach make it a good choice for in-hardware implementation on systems where safety is important, such as high-power particle accelerators, where large, sudden jumps in parameter values can result in damage or irradiation of components. Figure S5 shows the experimental application of bounded ES for automatic adjustment of 125 magnets simultaneously based on a noisy cost function that experienced sudden jumps when the beam was temporarily shut off, resulting in minor parameter phase shifts.




FIGURE S5 Experimental demonstration of bounded ES tuning 125 magnets simultaneously in the European X-ray Free-Electron Laser Facility. A sudden beam outage did not cause large parameter jumps.

REFERENCES

- [S9] M. Krstic and H. Wang, “Stability of extremum seeking feedback for general nonlinear dynamic systems,” *Automatica*, vol. 36, no. 4, pp. 595–601, 2000, doi: [10.1016/S0005-1098\(99\)00183-1](https://doi.org/10.1016/S0005-1098(99)00183-1).
- [S10] T. R. Oliveira, M. Krstic, and D. Tsubakino, “Extremum seeking for static maps with delays,” *IEEE Trans. Autom. Control*, vol. 62, no. 4, pp. 1911–1926, Apr. 2017, doi: [10.1109/TAC.2016.2564958](https://doi.org/10.1109/TAC.2016.2564958).
- [S11] J. I. Poveda and M. Krstic, “Nonsmooth extremum seeking control with user-prescribed fixed-time convergence,” *IEEE Trans. Autom. Control*, vol. 66, no. 12, pp. 6156–6163, Dec. 2021, doi: [10.1109/TAC.2021.3063700](https://doi.org/10.1109/TAC.2021.3063700).
- [S12] A. Williams, M. Krstic, and A. Scheinker, “Semiglobal safety-filtered extremum seeking with unknown CBFs,” *IEEE Trans. Autom. Control*, vol. 70, no. 3, pp. 1698–1713, Mar. 2025, doi: [10.1109/TAC.2024.3469785](https://doi.org/10.1109/TAC.2024.3469785).
- [S13] A. Scheinker, “Extremum seeking for force and torque actuated systems,” in *Proc. IEEE Conf. Decis. Control (CDC)*, 2018, pp. 7107–7111, doi: [10.1109/CDC.2018.8619556](https://doi.org/10.1109/CDC.2018.8619556).
- [S14] A. Scheinker and M. Krstic, “Extremum seeking with bounded update rates,” *Syst. Control Lett.*, vol. 63, pp. 25–31, 2014, doi: [10.1016/j.sysconle.2013.10.004](https://doi.org/10.1016/j.sysconle.2013.10.004).
- [S15] A. Scheinker and D. Scheinker, “Bounded extremum seeking with discontinuous dithers,” *Automatica*, vol. 69, pp. 250–257, Jul. 2016, doi: [10.1016/j.automatica.2016.02.023](https://doi.org/10.1016/j.automatica.2016.02.023).
- [S16] A. Scheinker and D. Scheinker, “Constrained extremum seeking stabilization of systems not affine in control,” *Int. J. Robust Nonlinear Control*, vol. 28, no. 2, pp. 568–581, 2018, doi: [10.1002/rnc.3886](https://doi.org/10.1002/rnc.3886).
- [S17] A. Scheinker et al., “Model-independent tuning for maximizing free electron laser pulse energy,” *Phys. Rev. Accelerators Beams*, vol. 22, no. 8, 2019, Art. no. 082802, doi: [10.1103/PhysRevAccelBeams.22.082802](https://doi.org/10.1103/PhysRevAccelBeams.22.082802).

directly adaptively attempting to modify even the small input image would be computationally infeasible, working directly in the low-dimensional latent greatly reduces the dimension of the problem and enables fast automatic adaptive tuning. By low dimensional, we typically refer to $\mathbf{z} \in \mathbb{R}^n$ with $n < 10$, so that adaptive tuning in this low-dimensional space is very fast. Although ES can handle relatively high-dimensional systems, its convergence is

typically faster for lower dimensions and also depends on how noisy the system is, how big we allow the extremum-seeking controller dither amplitude to be, and whether the system is time varying or not. For the example shown in Figure S5, based on work from [S17] (see “Bounded ES for Nonlinear Systems With Unknown Control Directions”), ES is tuning 125 magnets and converges in approximately 200 steps. In [56], where we were adjusting eight particle

accelerator parameters for a high-power high-current beam, we had to be careful to make very small iterative changes, and the ES method converged within 50 steps.

An overview of such a setup, which was first developed using experimentally measured input beam distributions at the compact particle accelerator-based ultrafast electron diffraction microscope known as the High Repetition-Rate Electron Scattering apparatus (HiRES) at Lawrence Berkeley National Laboratory [59], is shown in Figure 12 with a $n = 2$ 2D latent embedding $\mathbf{z} \in \mathbb{R}^2$. The bounded ES feedback (11) results in average latent dynamics

$$\begin{aligned} \dot{\mathbf{z}}(t) &= -\frac{k\alpha}{2} \nabla_{\mathbf{z}} C(y_j(t), \hat{y}_j(t)) \\ &= -\frac{k\alpha}{2} \nabla_{\mathbf{z}} C(y_j(t), \text{NN}_{\text{de}}(\hat{\mathbf{z}}(t))) \end{aligned} \quad (13)$$

and therefore

$$\dot{C}(y_j(t), \hat{y}_j(t)) = -\frac{k\alpha}{2} [\nabla_{\mathbf{z}} C(y_j(t), \hat{y}_j(t))]^T \nabla_{\mathbf{z}} C(y_j(t), \hat{y}_j(t)).$$

The generative strength of the CNN is being utilized to directly track the latent representation that tracks the measured time-varying signal without relying on access to the time-varying input beam conditions. It is, of course, possible for this approach to find the wrong answer since the adaptive feedback only has access to one of the 15 projections, and the solution to this partially observed problem is not unique. However, in practice, a lengthy beam tune-up first takes place, during which slow, invasive measurements are used to characterize the beam and properly set accelerator parameters. Following such a tune-up procedure, based on slow and invasive measurements that would otherwise interrupt operations, we can assume that we temporarily do know an accurate representation of both the beam's initial conditions at the beam source and the state of the accelerator's many parameters.

Once an accelerator goes into production, the beam's initial conditions, as well as the accelerator components start to drift. If our virtual diagnostic has been initialized with the correct beam state that it must now track, the local

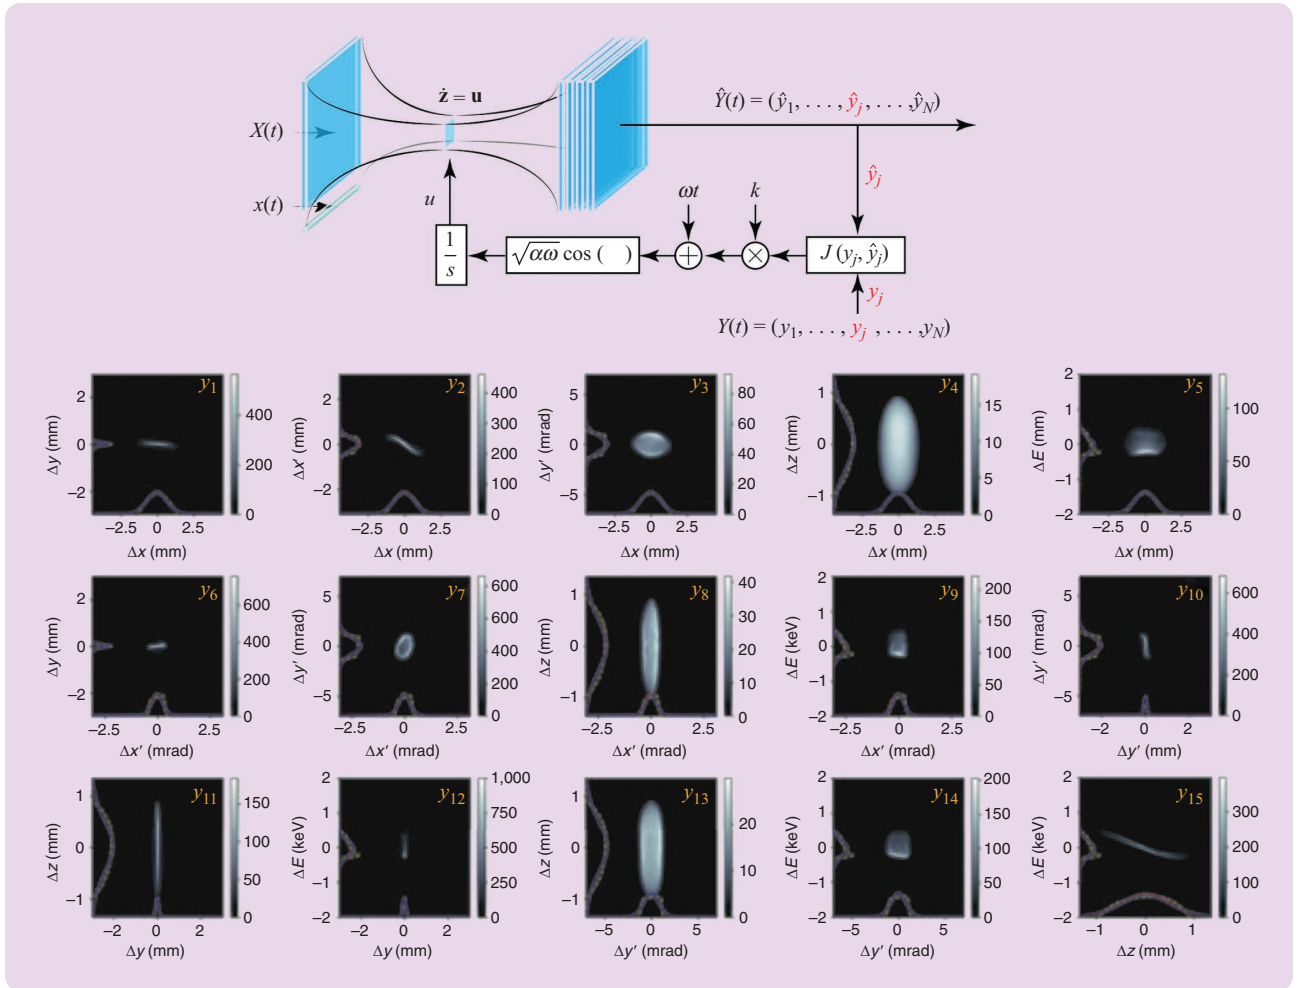


FIGURE 12 Although all 15 of the unique 2D projections of the simulated 6D phase space distribution of the electron beam in the HiRES compact accelerator are generated simultaneously, only a single projection is available for online measurement and is compared to its generated version to guide the adaptive feedback.

Another major benefit of diffusion-based models is their ability to smoothly interpolate between very different images by moving along the image manifold.

minimum of the cost function (12) coincides with the global minimum, which uniquely reconstructs the full 6D phase space of the beam. This is in contrast to starting with a random accelerator state and trying to minimize (12), which could lead to finding a local minimum that does not coincide with the correct match of the other unmeasured beam projections. Furthermore, as the beam continues to evolve, if we couple this adaptive approach with a calibrated online physics model, then we are able to uniquely track the true beam based only on a measurement of a single projection further than otherwise possible [13], but even with such an approach there is no guarantee that we can always uniquely track the correct entire 6D phase space when using just a single 2D projection. The problem is analogous to starting with a 3D sphere and attempting to predict its shape by tracking only a 2D projection of it, which is a circle, without knowing or being able to detect whether the sphere is transforming over time into a cone or a cylinder. Again, the physics constraints that can be incorporated into the scheme, such as what was done in [13], do help by constraining us to predictions that are physically consistent with respect to the setup, but even that is limited to the accuracy of the physics model and how well it can interpolate and extrapolate as the accelerator and beam change with time.

In practice, this adaptive tracking approach to virtual 6D phase space diagnostics has been demonstrated not only to uniquely track a time-varying beam, but also to further increase the robustness of the generative model beyond the span of the training data, as compared to a purely generative approach without latent space feedback [60]. In a completely different application, one of dynamic 3D imaging for material science, such an adaptive latent space tuning-based method has also been successful at tracking the 3D electron densities of time-varying samples [61], [62].

GENERATIVE DIFFUSION MODELS

While CNN-based encoder–decoder neural networks are powerful tools for dimension reduction, diffusion-based models have been proven to be the state of the art for generating higher resolution images, especially when the images are highly varying, as is the case for complex accelerator beams over a wide range of accelerator parameters, for class-conditional image synthesis [63], [64], and super-resolution methods [65]. Diffusion-based methods utilize advanced U-Net-based recursive architectures that share data at various image scales [66] as well as attention layers

to learn relationships in the lower dimensional latent representations of the U-Nets [1].

Another major benefit of diffusion-based models is their ability to smoothly interpolate between very different images by moving along the image manifold, which is again especially useful for exploring large configuration spaces of complex particle accelerators [67]. See “[Conditional Generative Diffusion](#)” for a technical introduction to the generative diffusion process.

The conditional guided diffusion process for virtual accelerator beam diagnostics was first demonstrated in [11], where a vector \mathbf{c} was used to guide the generative process to create virtual views of the (z, E) distribution of the electron beam in the European X-ray Free-Electron Laser Facility (FEL). The vector \mathbf{c} consisted of noninvasive measurements of accelerator parameters and noninvasive beam measurements, such as beam position monitors, which estimate the beam’s root mean square (x, y) location based on the image charges on conducting plates placed inside the beam pipe. The overall setup is shown in Figure 13.

Although this guided diffusion process is generating ~ 1 million-dimensional megapixel resolution images of the beam’s 2D (z, E) distribution, the predictions can now be adaptively tuned in a similar way as described previously, with the use of bounded ES for adaptive adjustment of just the 15D conditional vector \mathbf{c} , which is a greatly reduced dimensionality problem. In this setup, where only the (z, E) distribution is generated, adaptive feedback can be guided by noninvasive measurements of the beam’s average energy, such as energy spread spectra measurements. Adaptive tuning would then guide the generative diffusion process to generate projections (z, E) that match the measured energy spectra when integrated along z to give generated estimates of the distribution of E . Similarly, by integrating the generated images along E and fitting the resulting current profile to a Gaussian, σ_z of the generated distribution can be compared to noninvasive estimates of the bunch length for adaptive tuning. When working with multiple parameters, the ES algorithm is utilizing unique values of dithering frequencies ω_i in (11) so that each component is adjusted independently. A good way to choose the dithering frequencies is $\omega_i = \omega r_i$, where the values r_i are chosen uniformly within the range [1, 1.75] so that no two frequencies are equal and also are not integer multiples of each other. The value of ω is then increased until stable convergence of the overall scheme is achieved, as explained in more detail in references such as [14].

Conditional Generative Diffusion

Generative models based on diffusion were first created to utilize a gradual denoising approach inspired by statistical thermodynamics for modeling complex distributions [4]. The method was then modified to enable generation of high-resolution images [5]. The generative diffusion process adds noise to an image X over a large number of steps T until the pixel values of the image are transformed into a mean zero unit variance Gaussian distribution. The first step of the diffusion process creates the noise-corrupted image \mathbf{z}_1 defined by

$$\mathbf{z}_1 = \sqrt{1 - \beta_1} X + \sqrt{\beta_1} \epsilon_1, \quad \epsilon_1 \sim \mathcal{N}(\epsilon_1 | \mathbf{0}, \mathbf{I}). \quad (\text{S17})$$

Subsequent diffusion steps are iteratively defined as

$$\mathbf{z}_t = \sqrt{1 - \beta_t} \mathbf{z}_{t-1} + \sqrt{\beta_t} \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(\epsilon_t | \mathbf{0}, \mathbf{I}), \quad t \in \{2, \dots, T\}. \quad (\text{S18})$$

The noise schedule $\beta_t \in [0, 1]$ with $\beta_1 < \beta_2 < \dots < \beta_T$ prescribes the variance for the additive unit variance Gaussian noise ϵ_t at each step t , which defines how quickly images are converted to pure noise. This diffusion sequence forms a Markov chain with conditional distributions of the form

$$q(\mathbf{z}_t | \mathbf{z}_{t-1}) = \mathcal{N}(\mathbf{z}_t | \sqrt{1 - \beta_t} \mathbf{z}_{t-1}, \beta_t \mathbf{I}) \quad (\text{S19})$$

which is convenient for sampling random diffusion steps t without having to rerun the entire chain as \mathbf{z}_t can be rewritten as

$$\mathbf{z}_t = \sqrt{\alpha_t} X + \sqrt{1 - \alpha_t} \epsilon_t \quad (\text{S20})$$

where $\sqrt{1 - \alpha_t} \epsilon_t \sim \mathcal{N}(\epsilon_t | \mathbf{0}, \sqrt{1 - \alpha_t} \mathbf{I})$ represents the total noise added to the image in moving from step 1 to step t , with α_t given by

$$\alpha_t = \prod_{\tau=1}^t (1 - \beta_\tau). \quad (\text{S21})$$

Equation (S20) implies that

$$q(\mathbf{z}_t | X) = \mathcal{N}(\mathbf{z}_t | \sqrt{\alpha_t} X, (1 - \alpha_t) \mathbf{I}). \quad (\text{S22})$$

Because $(1 - \beta_t) < 1$, as $T \rightarrow \infty$ the terms α_t and $1 - \alpha_t$ approach 0 and 1, respectively, and

$$\lim_{T \rightarrow \infty} q(\mathbf{z}_T | X) = \mathcal{N}(\mathbf{z}_T | \mathbf{0}, \mathbf{I}) \quad (\text{S23})$$

so that any image is converted to a signal indistinguishable from mean zero unit variance Gaussian noise.

To generate images, the model learns to undo the noising diffusion process, and it is guided by a condition vector \mathbf{c} so that it creates particular images rather than purely random ones. For an image X , using Bayes' rule and the Gaussian change of variables, the conditional probability $q(\mathbf{z}_{t-1} | \mathbf{z}_t, X)$

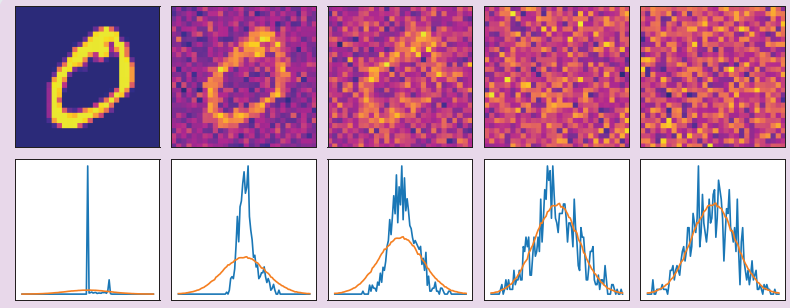


FIGURE S6 Several steps of the noising process for an MNIST image are shown in the top row, with the bottom row showing a histogram of the image pixel values compared to a mean zero unit Gaussian distribution.

describing one reverse step of the diffusion process can be written as

$$\mathcal{N}\left(\mathbf{z}_{t-1} \mid \frac{1 - \alpha_{t-1}}{1 - \alpha_t} \sqrt{1 - \beta_t} \mathbf{z}_t + \frac{\sqrt{\alpha_{t-1}} \beta_t}{1 - \alpha_t} X, \frac{\beta_t (1 - \alpha_{t-1})}{1 - \alpha_t} \mathbf{I}\right). \quad (\text{S24})$$

By rewriting (S20) as

$$X = \frac{1}{\sqrt{1 - \beta_t}} \mathbf{z}_t - \frac{\sqrt{\beta_t}}{\sqrt{1 - \beta_t}} \epsilon_t \quad (\text{S25})$$

the neural network $D_\theta(\mathbf{z}_t, t, \mathbf{c})$ can be used to predict how to remove noise between iterative steps to restore the original image according to

$$\mathbf{z}_{t-1} = \frac{1}{\sqrt{1 - \beta_t}} \left[\mathbf{z}_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} D_\theta(\mathbf{z}_t, t, \mathbf{c}) \right] + \sqrt{\beta_t} \epsilon \quad (\text{S26})$$

where θ represents the neural network's weights. At the final step, the image is created according to

$$X = \frac{1}{\sqrt{1 - \beta_1}} \left[\mathbf{z}_1 - \frac{\beta_1}{\sqrt{1 - \alpha_1}} D_\theta(\mathbf{z}_1, t, \mathbf{c}) \right]. \quad (\text{S27})$$

The diffusion model is trained to predict the noise at each diffusion step t by sampling a random $\epsilon \sim \mathcal{N}(\epsilon | \mathbf{0}, \mathbf{I})$ and updating the neural network's weights according to a descent of the gradient

$$\nabla_\theta \|\epsilon - D_\theta(\mathbf{z}_t, t, \mathbf{c})\|^2 = \nabla_\theta \|\epsilon - D_\theta(\sqrt{\alpha_t} X + \sqrt{1 - \alpha_t} \epsilon, t, \mathbf{c})\|^2. \quad (\text{S28})$$

Several steps from an example of the diffusion process applied to an MNIST image are shown in Figure S6.

Instead of working directly in the image space, as described here, it is also possible for a powerful VAE to first embed images into lower dimensional latent representations and then perform diffusion on those embeddings directly in a process known as latent diffusion. While this speeds up the diffusion process, there is a tradeoff between final image quality and the amount of compression. All of the results discussed in this article utilize diffusion directly in the image space.

The adaptive conditionally guided generative diffusion approach has also been extended for predicting all of the 15 unique 2D projections of a beam's 6D phase space distribution, as was done previously with an AE [40]. In this approach, the method was extended to include not only vector-based accelerator and beam inputs but also 2D projections of the beam's initial conditions. This second approach first utilizes a VAE to embed input beam images into a low-dimensional latent vector \mathbf{z} , which is then concatenated with a vector of noninvasive beam and accelerator

measurements \mathbf{p} and also with a vector \mathbf{d} that conditions the diffusion process on which of the 15 projections is to be generated for a conditional vector $\mathbf{c} = [\mathbf{z}, \mathbf{p}, \mathbf{d}]$. In this setup, a single model is used to generate one single projection at a time. The model is no longer attempting to generate all 15 projections simultaneously as the 15 channels of a single image. This method is referred to as conditional diffusion guided by variational autoencoder (cDVAE).

A high-level overview of the cDVAE setup is shown in Figure 14 for generating all of the projections of the HiRES




FIGURE 13 The diffusion process is guided by a condition vector \mathbf{c} composed of noninvasive accelerator component and beam measurements to generate a virtual view of the beam's (z, E) distribution.)

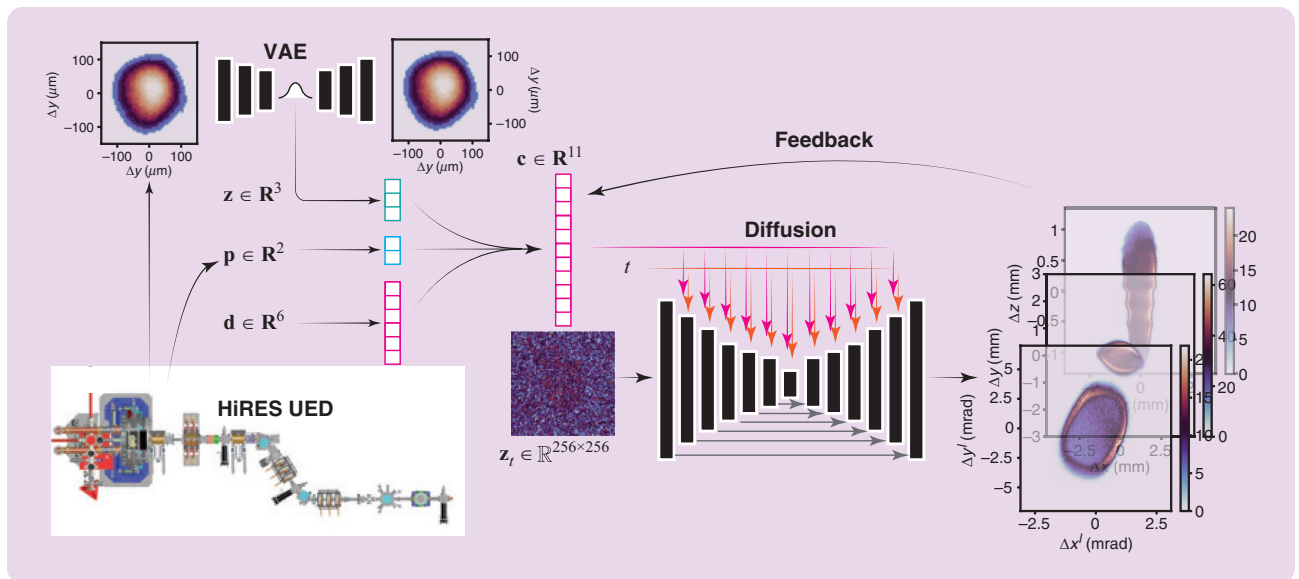


FIGURE 14 The diffusion process is guided by a condition vector $\mathbf{c} = [\mathbf{z}, \mathbf{p}, \mathbf{d}]$, where a VAE is used to embed input beam images into a low-dimensional latent vector \mathbf{z} , the vector \mathbf{p} is based on noninvasive beam and accelerator measurements, and the vector \mathbf{d} specifies which of the 15 projections of the 6D phase space should be generated. When applied in an adaptive unsupervised approach, the vector \mathbf{d} is first fixed so that the model generates only a single 2D projection to be compared to the available 2D measurement \mathbf{z} . The difference between the measured and predicted projection is then minimized by adaptively adjusting \mathbf{z} and \mathbf{p} via bounded ES. Once a good match is found, the vectors \mathbf{z} and \mathbf{p} are fixed, and \mathbf{d} is varied to generate all of the associated 2D projections. UED: ultrafast electron diffraction.

electron beam's phase space. The generation process is shown for several 2D phase space projections in **Figure 15**. When applied in an adaptive un-supervised approach, we assume that we no longer have access to the input images which are converted to the \mathbf{z} component of the conditional vector \mathbf{c} or to accurate measurements of the accelerator parameters \mathbf{p} . We also assume that both the input beam and accelerator components begin to change and drift with time, so we must adaptively tune our estimates of them to accurately track them to generate and predict the correct output beam distribution. To do this, the vector \mathbf{d} is first fixed so that the model generates only a single 2D projection to be compared to the available 2D measurement \mathbf{z} . The difference between the measured and predicted projection is then minimized by adaptively adjusting \mathbf{z} and \mathbf{p} via bounded ES. Once a good match is found, the vectors \mathbf{z} and \mathbf{p} are fixed, and \mathbf{d} is varied to generate all of the associated 2D projections.

An adaptive superresolution conditional generative diffusion method, based on the approach described previously but with hard physics constraints built into the generative process, has been developed for generating 6D tensor-based representations of the 6D phase space density of charged particle beams [28].

ADAPTIVE LATENT DIFFUSION

While diffusion-based generative models are the state-of-the-art method for generating a widely diverse set of accurate representations of complex objects, such as high-resolution images, their main drawback is that they can be computationally

expensive, especially when operating in image space and being used to generate high-resolution images over many iterative steps. Fortunately, a method known as latent diffusion, which combines the ideas from AEs and from diffusion models, as discussed previously, has been developed, which greatly speeds up the generative diffusion process [68], [69], [70]. In the latent diffusion process, the first step is to train an AE or a VAE to compress high-resolution images X into much smaller latent images \mathbf{z} , from which the decoder then reconstructs an estimate \hat{X} of X .

The conditional latent diffusion model is then used to directly generate latent representations based on a conditional input \mathbf{c} , so that the diffusion model learns to create estimates $\hat{\mathbf{z}}$ of the true latent representations \mathbf{z} that have been created by the encoder network. As diffusion is happening in a much lower dimensional space, it can use a smaller network and also fewer diffusion steps. Finally, the generated latent image $\hat{\mathbf{z}}$ is mapped directly to the estimate of an image \hat{X} by the trained decoder from the AE model. Besides speeding up the diffusion process, this approach is also a great way to enable conditional generation for applications such as solving inverse problems.

For example, an adaptive latent diffusion-based method for solving the extreme inverse problem of mapping scalar beam current and beam loss measurements to detailed images of the beam's phase space distribution is being developed for the Los Alamos Neutron Science Center (LANSCE) linear accelerator [67]. In this approach, a VAE is first used to compress 256×256 -pixel images by a factor of $\sim 340 \times$ down to $8 \times 8 \times 4$ latent images. The latent




FIGURE 15 Steps along the generative diffusion process are shown for seven of the 15 projections as labeled on the left edge of the figure. The generated images are compared to the true figures, and their differences are also shown.

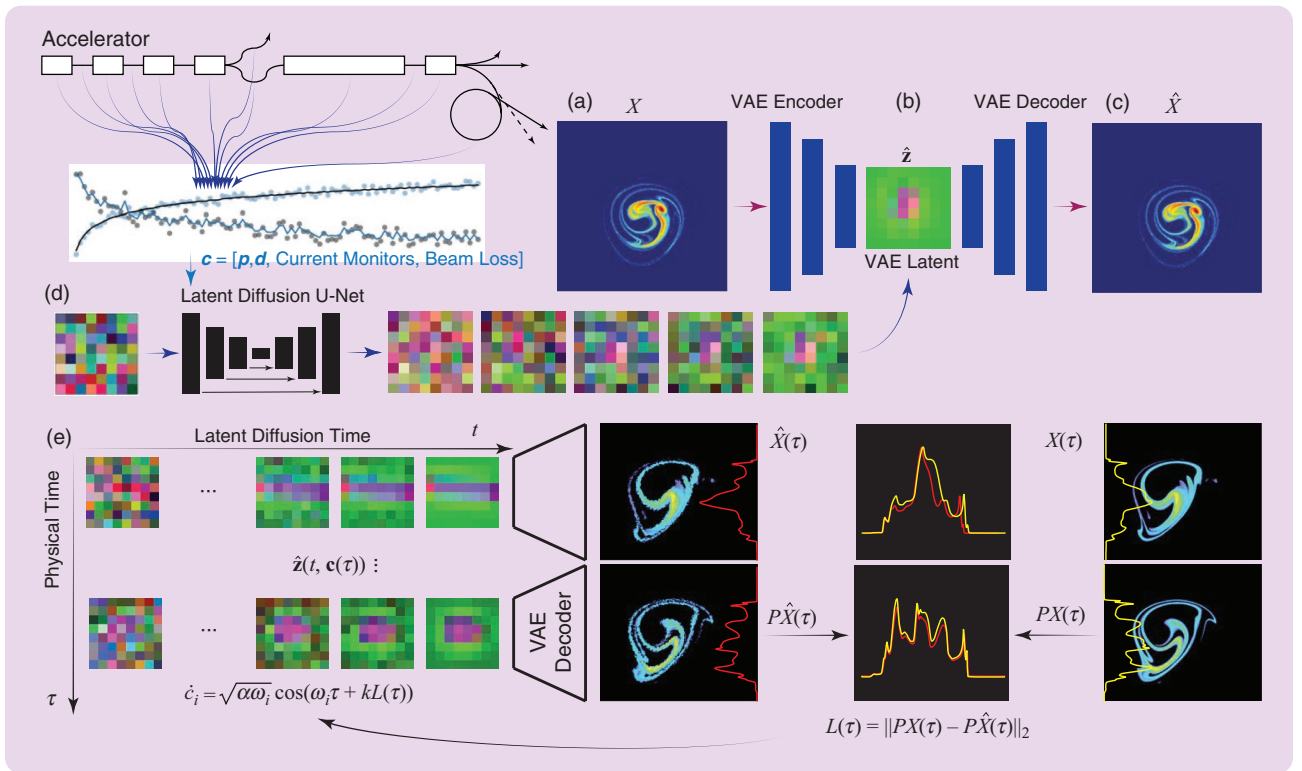


FIGURE 16 The encoder (a) of a VAE is trained to compress high-resolution input images $X \in \mathbb{R}^{256 \times 256}$ down into much smaller $\mathbf{z} \in \mathbb{R}^{8 \times 8 \times 3}$ latent representations (b), from which the decoder (c) regenerates an estimate \hat{X} of X . The latent diffusion model is then trained to take conditional vectors as inputs and generate estimates $\hat{\mathbf{z}}$ of the latent vectors (d), from which the decoder can then generate \hat{X} . After the VAE is trained, the encoder is no longer used; the only needed inputs are conditional vectors consisting of beam loss and current monitors and accelerator parameter settings (d), which work directly with the latent diffusion model and the decoder of the VAE. To add adaptive feedback to this setup, we look at noninvasively measurable 1D projections, such as the energy spread, of the beam, and we compare those projections to the same projections from the diffusion-generated image (e). ES is then used to adaptively tune the conditional vector to track the time-varying beam distribution.

diffusion model is then conditioned on a vector of measurable accelerator parameters, such as magnet and RF settings, together with a vector of beam loss or beam current measurements along the length of a kilometer-long accelerator. Although the noisy 1D loss/current measurements are incomprehensible to a person, the conditional diffusion model is found to be able to map them to detailed images of the LANSCE beam's phase space. This setup can be described as follows.

We consider the true state of a 2D phase space projection of the beam, such as the time versus energy (t, E) longitudinal phase space, which we denote as $X(\tau)$. Note that this is a two-time-scale approach in which τ represents physical time relative to which a dynamic system is evolving, and t represents synthetic algorithm runtime within a computer. We utilize noninvasive beam and accelerator measurements that include accelerator component setpoints and beam measurements, such as beam current or beam loss monitor readings, which are grouped together into a conditional vector $\mathbf{c}(\tau)$. Note that some of the components of $\mathbf{c}(\tau)$ are only approximations, such as magnetic fields in the accelerators, because we cannot set them directly. Rather, we set things like voltages or currents

of power supplies, which then generate electromagnetic fields. This conditional vector is used by the latent diffusion model to generate a latent estimate $\hat{\mathbf{z}}(t, \tau) = \mathcal{D}(\mathbf{c}(\tau), t)$, and that is converted into an estimate $\hat{X}(\tau)$ of $X(\tau)$ by the decoder part of the VAE. The direct measurement of $X(\tau)$ might be unavailable or may require a lengthy, destructive process that would interrupt accelerator operations. However, we may have access to a 1D projection $PX(\tau)$, such as projecting (z, E) to its E -axis. The measured projection is compared to a projection of the generated beam image, and the difference is a cost function $L = \|PX(\tau) - P\hat{X}(\tau)\|_2$ that is minimized by adjusting the conditional vector $\mathbf{c}(\tau)$. A schematic of the overall approach is shown in Figure 16. As shown in [67], this algorithm is able to track a time-varying beam with only limited noninvasive 1D beam projection measurements, and it is robust to noise-corrupted measurements of $\mathbf{c}(\tau)$ and $PX(\tau)$.

CONCLUSION

This article provides a very brief introduction to deep learning, focusing on several generative deep learning methods and ways in which they are being combined with adaptive feedback control to develop robust virtual diagnostics for

high-dimensional complex systems, such as particle accelerators. The approaches shown are general in nature and can be applied to a wide variety of complex high-dimensional systems. Bounded ES has proven to be a powerful control technique that complements the generative strengths of deep neural network-based models, offering robustness to noise and the ability to handle time-varying systems. The approach of adding adaptive feedback directly within the low-dimensional latent embeddings or conditional vectors of generative models enables real-time virtual diagnostics of time-varying systems, and these diagnostics in turn enable control of large complex objects, such as high-resolution images, whose iterative tuning would be otherwise impractical.

ACKNOWLEDGMENT

This work was supported by the U.S. Department of Energy, Los Alamos National Laboratory Directed Research and Development Program Directed Research project 20220074DR.

AUTHOR INFORMATION

Alexander Scheinker (ascheink@lanl.gov) received the Ph.D. degree in dynamic systems and control theory from the University of California San Diego and also has master's degrees in math and physics. He is a research scientist at Los Alamos National Laboratory, where he is the Adaptive Machine Learning team leader in the Applied Electrodynamics Group. His focus is on incorporating adaptive feedback control algorithms and hard physics constraints within the architectures of deep generative models to enable their use for virtual diagnostics and automatic control of complex time-varying dynamic systems. He has applied his methods to particle accelerators around the world and also for material science applications, such as 3D coherent X-ray diffraction imaging. He is a Senior Member of IEEE.

REFERENCES

[1] A. Vaswani, "Attention is all you need," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2017, pp. 6000–6010. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf

[2] J. Achiam et al., "GPT-4 technical report," 2023, *arXiv:2303.08774*, doi: [10.48550/arXiv.2303.08774](https://arxiv.org/abs/2303.08774).

[3] A. Dubey et al., "The Llama 3 herd of models," 2024, *arXiv:2407.21783*.

[4] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2256–2265. [Online]. Available: <https://proceedings.mlr.press/v37/sohl-dickstein15.html>

[5] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2020, pp. 6840–6851. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf

[6] J. Song, C. Meng, and S. Ermon, "Denosing diffusion implicit models," 2020, *arXiv:2010.02502*.

[7] T. Karras, M. Aittala, S. Laine, and T. Aila, "Elucidating the design space of diffusion-based generative models," in *Proc. 36th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2022, pp. 26,565–26,577. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/a98846e9d9cc01cfb87eb694d946ce6b-Paper-Conference.pdf

[8] Z. Guo et al., "Diffusion models in bioinformatics and computational biology," *Nature Rev. Bioeng.*, vol. 2, no. 2, pp. 136–154, 2023, doi: [10.1038/s44222-023-00114-9](https://doi.org/10.1038/s44222-023-00114-9).

[9] J. Abramson et al., "Accurate structure prediction of biomolecular interactions with AlphaFold 3," *Nature*, vol. 630, no. 8016, pp. 493–500, 2024, doi: [10.1038/s41586-024-07487-w](https://doi.org/10.1038/s41586-024-07487-w).

[10] S. Yuan and S. V. Dordevic, "Diffusion models for conditional generation of hypothetical new families of superconductors," *Sci. Rep.*, vol. 14, no. 1, 2024, Art. no. 10275, doi: [10.1038/s41598-024-61040-3](https://doi.org/10.1038/s41598-024-61040-3).

[11] A. Scheinker, "Conditional guided generative diffusion for particle accelerator beam diagnostics," *Sci. Rep.*, vol. 14, no. 1, 2024, Art. no. 19210, doi: [10.1038/s41598-024-70302-z](https://doi.org/10.1038/s41598-024-70302-z).

[12] A. Scheinker, A. Edelen, D. Bohler, C. Emma, and A. Lutman, "Demonstration of model-independent control of the longitudinal phase space of electron beams in the Linac-coherent light source with femtosecond resolution," *Phys. Rev. Lett.*, vol. 121, no. 4, 2018, Art. no. 044801, doi: [10.1103/PhysRevLett.121.044801](https://doi.org/10.1103/PhysRevLett.121.044801).

[13] A. Scheinker, F. Cropp, S. Paiguaga, and D. Filippetto, "An adaptive approach to machine learning for compact particle accelerators," *Sci. Rep.*, vol. 11, no. 1, 2021, Art. no. 19187, doi: [10.1038/s41598-021-98785-0](https://doi.org/10.1038/s41598-021-98785-0).

[14] A. Scheinker, "Model independent beam tuning," in *Proc. 4th Int. Particle Accel. Conf., (IPAC)*, Shanghai, China, 2013, pp. 1862–1864. [Online]. Available: <https://accelconf.web.cern.ch/ipac2013/papers/tupwa068.pdf>

[15] C. Benjamins et al., "CARL: A benchmark for contextual and adaptive reinforcement learning," 2021, *arXiv:2110.02102*.

[16] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1126–1135. [Online]. Available: <https://proceedings.mlr.press/v70/finn17a.html>

[17] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," in *Proc. 4th Int. Conf. Learn. Representations*, San Juan, Puerto Rico, 2016.

[18] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel, "RL2: Fast reinforcement learning via slow reinforcement learning," 2016, *arXiv:1611.02779*.

[19] H. Kurniawati, "Partially observable Markov decision processes and robotics: A survey," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 5, no. 1, pp. 253–277, 2022, doi: [10.1146/annurev-control-042920-092451](https://doi.org/10.1146/annurev-control-042920-092451).

[20] A. Hallak, D. Di Castro, and S. Mannor, "Contextual Markov decision processes," 2015, *arXiv:1502.02259*.

[21] J. H. Cho, V. Jayawardana, S. Li, and C. Wu, "Model-based transfer learning for contextual reinforcement learning," in *Proc. 38th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2024, pp. 88,279–88,319. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2024/file/a10c3d85879c29331ba4d73ede56c7d3-Paper-Conference.pdf

[22] J. Köhler and M. N. Zeilinger, "Predictive control for nonlinear stochastic systems: Closed-loop guarantees with unbounded noise," *IEEE Trans. Autom. Control*, vol. 70, no. 11, pp. 7382–7397, Nov. 2025, doi: [10.1109/TAC.2025.3571575](https://doi.org/10.1109/TAC.2025.3571575).

[23] J. Degraeve et al., "Magnetic control of tokamak plasmas through deep reinforcement learning," *Nature*, vol. 602, no. 7897, pp. 414–419, 2022, doi: [10.1038/s41586-021-04301-9](https://doi.org/10.1038/s41586-021-04301-9).

[24] B. Recht, "A tour of reinforcement learning: The view from continuous control," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 2, no. 1, pp. 253–279, 2019, doi: [10.1146/annurev-control-053018-023825](https://doi.org/10.1146/annurev-control-053018-023825).

[25] J. Zhang, S. Z. Yong, and D. Panagou, "Safety-critical control with off-line-online neural network inference," 2024, *arXiv:2408.00918*.

[26] Z. Wang, Y. Gao, S. Wang, M. M. Zavlanos, A. Abate, and K. H. Johansson, "Policy evaluation in distributional LQR," in *Proc. 5th Annual Learn. Dyn. Control Conf.*, PMLR, 2023, pp. 1245–1256. [Online]. Available: <https://proceedings.mlr.press/v211/wang23c.html>

[27] H. Hasanbeig, D. Kroening, and A. Abate, "Certified reinforcement learning with logic guidance," *Artif. Intell.*, vol. 322, Sep. 2023, Art. no. 103949, doi: [10.1016/j.artint.2023.103949](https://doi.org/10.1016/j.artint.2023.103949).

[28] A. Scheinker, "Physics-constrained superresolution diffusion for six-dimensional phase space diagnostics," *Physical Rev. Res.*, vol. 7, no. 2, 2025, Art. no. 023091, doi: [10.1103/PhysRevResearch.7.023091](https://doi.org/10.1103/PhysRevResearch.7.023091).

[29] L. Huang, J. Lygeros, and F. Dörfler, "Robust and kernelized data-enabled predictive control for nonlinear systems," *IEEE Trans. Control Syst. Technol.*, vol. 32, no. 2, pp. 611–624, Mar. 2024, doi: [10.1109/TCST.2023.3329334](https://doi.org/10.1109/TCST.2023.3329334).


[30] J. Yan, A. Chakrabarty, A. Rupenyan, and J. Lygeros, "MPC of uncertain nonlinear systems with meta-learning for fast adaptation of neural predictive models," in *Proc. IEEE 20th Int. Conf. Automat. Sci. Eng. (CASE)*, Piscataway, NJ, USA: IEEE Press, 2024, pp. 1910–1915, doi: [10.1109/CASE59546.2024.10711717](https://doi.org/10.1109/CASE59546.2024.10711717).

- [31] M. Fazlyab, M. Morari, and G. J. Pappas, "Safety verification and robustness analysis of neural networks via quadratic constraints and semi-definite programming," *IEEE Trans. Autom. Control*, vol. 67, no. 1, pp. 1–15, Jan. 2022, doi: [10.1109/TAC.2020.3046193](https://doi.org/10.1109/TAC.2020.3046193).
- [32] E. Tolstaya, F. Gama, J. Paulos, G. Pappas, V. Kumar, and A. Ribeiro, "Learning decentralized controllers for robot swarms with graph neural networks," in *Proc. Conf. Robot Learn.*, PMLR, pp. 671–682, 2020. [Online]. Available: <https://proceedings.mlr.press/v100/tolstaya20a.html>
- [33] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A general safety framework for learning-based control in uncertain robotic systems," *IEEE Trans. Autom. Control*, vol. 64, no. 7, pp. 2737–2752, Jul. 2019, doi: [10.1109/TAC.2018.2876389](https://doi.org/10.1109/TAC.2018.2876389).
- [34] S. Bansal and C. J. Tomlin, "Deepreach: A deep learning approach to high-dimensional reachability," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 1817–1824, doi: [10.1109/ICRA48506.2021.9561949](https://doi.org/10.1109/ICRA48506.2021.9561949).
- [35] D. A. Nix and A. S. Weigend, "Estimating the mean and variance of the target probability distribution," in *Proc. IEEE Int. Conf. Neural Netw. (ICNN)*, vol. 1, Piscataway, NJ, USA: IEEE Press, pp. 55–60, doi: [10.1109/ICNN.1994.374138](https://doi.org/10.1109/ICNN.1994.374138).
- [36] H. Bourlard and Y. Kamp, "Auto-association by multilayer perceptrons and singular value decomposition," *Biol. Cybern.*, vol. 59, nos. 4–5, pp. 291–294, 1988, doi: [10.1007/BF00332918](https://doi.org/10.1007/BF00332918).
- [37] P. Baldi and K. Hornik, "Neural networks and principal component analysis: Learning from examples without local minima," *Neural Netw.*, vol. 2, no. 1, pp. 53–58, 1989, doi: [10.1016/0893-6080\(89\)90014-2](https://doi.org/10.1016/0893-6080(89)90014-2).
- [38] E. Plaut, "From principal subspaces to principal components with linear autoencoders," 2018, *arXiv:1804.10253*.
- [39] C. Behrens et al., "Few-femtosecond time-resolved measurements of X-ray free-electron lasers," *Nat. Commun.*, vol. 5, no. 1, 2014, Art. no. 3762, doi: [10.1038/ncomms4762](https://doi.org/10.1038/ncomms4762).
- [40] A. Scheinker, "cDVAE: VAE-guided diffusion for particle accelerator beam 6D phase space projection diagnostics," *Sci. Rep.*, vol. 14, no. 1, 2024, Art. no. 29303, doi: [10.1038/s41598-024-80751-1](https://doi.org/10.1038/s41598-024-80751-1).
- [41] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*, vol. 10. Upper Saddle River, NJ, USA: Pearson, 2010.
- [42] A. Scheinker and M. Krstić, "Minimum-seeking for CLFs: Universal semiglobally stabilizing feedback under unknown control directions," *IEEE Trans. Autom. Control*, vol. 58, no. 5, pp. 1107–1122, May 2013, doi: [10.1109/TAC.2012.2225514](https://doi.org/10.1109/TAC.2012.2225514).
- [43] M. Guay and M. Benosman, "Finite-time extremum seeking control for a class of unknown static maps," *Int. J. Adapt. Control Signal Process.*, vol. 35, no. 7, pp. 1188–1201, 2021, doi: [10.1002/acs.3123](https://doi.org/10.1002/acs.3123).
- [44] S. Koga, M. Benosman, and J. Borggaard, "Extremum seeking-based observer design for reduced order models of coupled thermal and fluid systems," *Int. J. Adapt. Control Signal Process.*, vol. 35, no. 7, pp. 1316–1335, 2021, doi: [10.1002/acs.3252](https://doi.org/10.1002/acs.3252).
- [45] A. Scheinker, M. Bland, M. Krstić, and J. Audia, "Extremum seeking-based optimization of high voltage converter modulator rise-time," *IEEE Trans. Control Syst. Technol.*, vol. 22, no. 1, pp. 34–43, Jan. 2014, doi: [10.1109/TCST.2013.2240387](https://doi.org/10.1109/TCST.2013.2240387).
- [46] Q. Han, Z. Liu, H. Su, and X. Liu, "Distributed I&I adaptive output feedback control of uncertain second-order systems with output constraint and input saturation," *Inf. Sci.*, vol. 710, Aug. 2025, Art. no. 122086, doi: [10.1016/j.ins.2025.122086](https://doi.org/10.1016/j.ins.2025.122086).
- [47] M. Abdelgalil and J. I. Poveda, "On lie-bracket averaging for hybrid dynamical systems with applications to model-free control and optimization," *IEEE Trans. Autom. Control*, vol. 7, no. 7, pp. 4655–4670, Jul. 2025, doi: [10.1109/TAC.2025.3529375](https://doi.org/10.1109/TAC.2025.3529375).
- [48] Z. Zhang, Y. Zhang, and J. Sun, "Piecewise constant tuning gain-based singularity-free MRAC with application to aircraft control systems," *IEEE Trans. Cybern.*, vol. 55, no. 8, pp. 4026–4037, Aug. 2025, doi: [10.1109/TCYB.2025.3567177](https://doi.org/10.1109/TCYB.2025.3567177).
- [49] P. Gardonio, G. K. Rodrigues, L. D. Bo, and E. Turco, "Extremum seeking online tuning of a piezoelectric vibration absorber based on the maximisation of the shunt electric power absorption," *Mech. Syst. Sig. Process.*, vol. 176, Aug. 2022, Art. no. 109171, doi: [10.1016/j.ymsp.2022.109171](https://doi.org/10.1016/j.ymsp.2022.109171).
- [50] X. Cheng, Y. Tan, and I. Mareels, "Stability and robustness analysis of switched vibrational control," in *Proc. 59th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2020, pp. 3304–3309, doi: [10.1109/CDC42340.2020.9304398](https://doi.org/10.1109/CDC42340.2020.9304398).
- [51] X. Sun, W. Chen, and H. Dai, "Event-triggered adaptive control for a class of non-linear systems with multiple unknown control directions," *IET Control Theory Appl.*, vol. 12, no. 5, pp. 629–637, 2018, doi: [10.1049/iet-cta.2017.0662](https://doi.org/10.1049/iet-cta.2017.0662).
- [52] B. Zitte, B. Hamroun, F. Couenne, and I. Pitault, "Extremum-seeking based distributed optimization of heat exchangers network," *IFAC-PapersOnLine*, vol. 51, no. 23, pp. 331–336, 2018, doi: [10.1016/j.ifacol.2018.12.057](https://doi.org/10.1016/j.ifacol.2018.12.057).
- [53] J. Sharafi, W. H. Moase, and C. Manzie, "Multiplexed extremum seeking for calibration of spark timing in a CNG-fuelled engine," *Control Eng. Pract.*, vol. 72, pp. 42–52, Mar. 2018, doi: [10.1016/j.conengprac.2017.11.005](https://doi.org/10.1016/j.conengprac.2017.11.005).
- [54] Z. Cao, R. Gondhalekar, E. Dassau, and F. J. Doyle, "Extremum seeking control for personalized zone adaptation in model predictive control for type 1 diabetes," *IEEE Trans. Biomed. Eng.*, vol. 65, no. 8, pp. 1859–1870, Aug. 2018, doi: [10.1109/TBME.2017.2783238](https://doi.org/10.1109/TBME.2017.2783238).
- [55] A. Scheinker, S. Baily, D. Young, J. S. Kolski, and M. Prokop, "In-hardware demonstration of model-independent adaptive tuning of noisy systems with arbitrary phase drift," *Nucl. Instrum. Methods Phys. Res., Sect. A*, vol. 756, pp. 30–38, Aug. 2014, doi: [10.1016/j.nima.2014.04.026](https://doi.org/10.1016/j.nima.2014.04.026).
- [56] A. Scheinker, E.-C. Huang, and C. Taylor, "Extremum seeking-based control system for particle accelerator beam loss minimization," *IEEE Trans. Control Syst. Technol.*, vol. 30, no. 5, pp. 2261–2268, Sep. 2022, doi: [10.1109/TCST.2021.3136133](https://doi.org/10.1109/TCST.2021.3136133).
- [57] A. Scheinker and D. Scheinker, "Extremum seeking for optimal control problems with unknown time-varying systems and unknown objective functions," *Int. J. Adapt. Control Signal Process.*, vol. 35, no. 7, pp. 1143–1161, 2021, doi: [10.1002/acs.3097](https://doi.org/10.1002/acs.3097).
- [58] A. Scheinker, "100 years of extremum seeking: A survey," *Automatica*, vol. 161, Mar. 2024, Art. no. 111481, doi: [10.1016/j.automatica.2023.111481](https://doi.org/10.1016/j.automatica.2023.111481).
- [59] A. Scheinker, "Adaptive machine learning for time-varying systems: Low dimensional latent space tuning," *J. Instrum.*, vol. 16, no. 10, 2021, Art. no. P10008, doi: [10.1088/1748-0221/16/10/P10008/meta](https://doi.org/10.1088/1748-0221/16/10/P10008/meta).
- [60] A. Scheinker, F. Cropp, and D. Filippetto, "Adaptive autoencoder latent space tuning for more robust machine learning beyond the training set for six-dimensional phase space diagnostics of a time-varying ultrafast electron-diffraction compact accelerator," *Phys. Rev. E*, vol. 107, no. 4, 2023, Art. no. 045302, doi: [10.1103/PhysRevE.107.045302](https://doi.org/10.1103/PhysRevE.107.045302).
- [61] A. Scheinker and R. Pokharel, "Adaptive 3D convolutional neural network-based reconstruction method for 3D coherent diffraction imaging," *J. Appl. Phys.*, vol. 128, no. 18, p.184901, Nov. 2020, doi: [10.1063/5.0014725](https://doi.org/10.1063/5.0014725).
- [62] A. Scheinker and R. Pokharel, "Enabling dynamic 3D coherent diffraction imaging via adaptive latent space tuning of generative autoencoders," *npj Comput. Mater.*, vol. 10, no. 1, 2024, Art. no. 283, doi: [10.1038/s41524-024-01482-5](https://doi.org/10.1038/s41524-024-01482-5).
- [63] P. Dhariwal and A. Nichol, "Diffusion models beat GANs on image synthesis," in *Proc. 35th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2021, pp. 8780–8794. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2021/file/49ad23d1ec9fa4bd8d77d02681df5cfa-Paper.pdf
- [64] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans, "Cascaded diffusion models for high fidelity image generation," *J. Mach. Learn. Res.*, vol. 23, no. 47, pp. 1–33, 2022. [Online]. Available: <http://jmlr.org/papers/v23/21-0635.html>
- [65] C. Saharia, J. Ho, W. Chan, T. Salimans, D. J. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 4, pp. 1–14, Apr. 2022, doi: [10.1109/TPAMI.2022.3204461](https://doi.org/10.1109/TPAMI.2022.3204461).
- [66] O. Ronneberger, P. Fischer, and T. Brox, U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2015, pp. 234–241, doi: [10.1007/978-3-319-24574-4_28](https://doi.org/10.1007/978-3-319-24574-4_28).
- [67] A. Scheinker and A. Williams, "Latent diffusion can map beam loss to two-dimensional phase space projections," *Physical Rev. Accelerators Beams*, vol. 28, no. 9, 2025, Art. no. 094602, doi: [10.1103/rqg9-g3dp](https://doi.org/10.1103/rqg9-g3dp).
- [68] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-resolution image synthesis with latent diffusion models," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 10,684–10,695. [Online]. Available: https://openaccess.thecvf.com/content/CVPR2022/html/Rombach_High-Resolution_Image_Synthesis_With_Latent_Diffusion_Models_CVPR_2022_paper
- [69] O. Avrahami, O. Fried, and D. Lischinski, "Blended latent diffusion," *ACM Trans. Graphics*, vol. 42, no. 4, pp. 1–11, 2023, doi: [10.1145/3592450](https://doi.org/10.1145/3592450).
- [70] D. Podell et al., "SDXL: Improving latent diffusion models for high-resolution image synthesis," 2023, *arXiv:2307.01952*.

Extremum Seeking: The Simple Feedback That Will Outlive Us All

In this issue of the magazine, my interview with Professor Miroslav Krstic about the history of extremum seeking is presented. Miroslav Krstic works on nonlinear, adaptive, and infinite dimensional control theory and applications. He is a co-author of 19 books, including the classic *Nonlinear and Adaptive Control Design* (1995), one of the two most cited control monographs, the graduate textbook *Backstepping control of PDEs* (2008), which has won the Chestnut Textbook Prize from IFAC, and the single-authored *Delay Compensation for Nonlinear, Adaptive, and PDE Systems* (2009). Miroslav

has played a central role in the inception and evolution of the field of extremum seeking. The applications of his work in extremum seeking have enabled technology developments worth billions of dollars in photolithography for chip manufacturing and in landing jet aircraft on the newest class of aircraft, “super-carriers” Gerald Ford. His algorithms have also been employed in the management of Lithium-ion battery systems, oil drilling, and charged particle accelerators.

Anuradha Annaswamy 

MIROSLAV KRSTIC

Anu: When and how did you become interested in extremum seeking (ES) control?

Miroslav: The encounter with ES is one of the fondest, most stressful, and most consequential experiences of my career. The year was 1997. Adaptive nonlinear control, the subject on which I had written my dissertation, many papers, and the book [1], was commanding the attention of a significant slice of the control community. The feeling was that if we just keep advancing the complex tools of feedback design, we will be conquering ever more difficult and relevant problems in control of nonlinear uncertain systems.

Then your work, Anu [2], in control of combustion, attracted my attention. I began consulting for United Technologies on control of combustion instabilities in jet engines. A chance discussion with a combustion specialist drew my attention to an algorithm called “peak seeking.”



Miroslav Krstic

Peak seeking wasn’t new, although I was learning of it for the first time. The most creative contributions to peak seeking, also known as ES, date to the 1950s–1960s. In fact, it was invented in 1922 [3], the same year as PID control [4]. And the algorithm wasn’t complicated; it was the antithesis to the kind of nonlinear unknown systems it was meant to control, containing just one integrator, combined with a sinusoidal signal used both as a perturbation and for demodulation.

Coming fresh off years of research on control of nonlinear systems with unknown parameters, built around differential geometric transformations and global Lyapunov techniques, I was shocked and paralyzed for a long moment. How would such a feedback law, which would take two lines to code, be capable of doing something useful for nonlinear systems with unknown models?

True, it was not meant for stabilizing unstable nonlinear systems, like adaptive backstepping. Its name, “peak seeking,” made it clear that the algorithm’s goal was optimization. But the magic mix of an integrator and an additive sinusoid at the input, applied also multiplicatively at the output, seemed impenetrable.

Frantically, I read the available literature. Draper and Li [5] introduced peak seeking in the United States in 1951, after a half-decade of inventions in the Soviet Union. The literature that followed in the 1950s and 1960s, virtually all in the Soviet Union, abounded with creative heuristics and intuition, but there was no theory in the modern sense: no theorems that state clear

How would such a feedback law, which would take two lines to code, be capable of doing something useful for nonlinear systems with unknown models?

properties, no proofs based on precise assumptions.

About thirty years had passed since the most recent interesting publications in the late 1960s. Astrom and Wittenmark [6], just a couple of years before my dive into the study of ES, put it among the most promising future areas for adaptive control. But how can something that was abandoned more than a generation earlier, and has not lived up to the rising standards of theoretical control, possibly be our field's future?

Anu: ES had existed for decades. You developed its first proof in 1997. What motivated you to seek a theoretical proof for ES?

Miroslav: Saying that I was unsettled by ES would be an understatement. Having my foundation shaken is a more apt description. These theoretically uninterpretable two-line algorithms, although not tackling the same goals, were dealing with the same nonlinear, highly uncertain systems for which I had spent half a decade producing pages-long control laws and parameter update laws.

Just ignoring ES, or accepting that it is “unprovable,” although it works well, was not an option. All my preceding work was challenged. I had to understand not only what the limits of ES are but also how it, provably, does what it does well, even if what it does is limited.

I was fortunate to have had good training in averaging theory and singular perturbations from UC Santa Barbara. But it was not just unclear how to combine these two asymptotic methods to prove something but also unclear *what* to prove and even how

exactly to use ES algorithms for nonlinear *dynamical* systems.

In retrospect, it is obvious that 1) the nonlinear plant has to be stable for all constant inputs; 2) one has to run ES slower than the plant's transients, so the transients settle between significant changes of the inputs; 3) the “exploring” perturbations in ES, while slower than the plant, need to be fast relative to the parameter adjustments of the algorithm; and 4) both the perturbation amplitude and the integrator gain need to be small compared with the perturbation frequency.

Used in this manner, ES lends itself to two nested approximations: a singular perturbation approximation of the plant and an averaging approximation of the ES loop. Local exponential practical stability follows, and the residual bias is to the “flatter” (less lossy) side of the peak of the equilibrium map.

That was the beginning: a paper written in the spring of 1997, over a few days, and published at the CDC that December [7].

Anu: After providing the theory for this method, was your plan then to direct your career into ES?

Miroslav: Not at all. With my serenity restored, I intended to turn my attention back to complicated things and let ES be used by whoever wishes, if anybody cares about such simple algorithms.

But that is not how things unfolded. ES had too much potential. Every idea yielded bountifully compared with the effort demanded: early applications in jet engine compressors, combustors, bioreactors, antilock braking, formation flight,

and, most importantly, autonomous ground and underwater vehicles, theoretical advances for time-varying optimizers, discrete-time and stochastic ES, Newton versions of ES, and, amazingly, seeking of Nash equilibria for unknown payoff functions. Things absolutely did not go according to plan, the plan having been to let ES be and direct all my energy to PDE control.

Anu: How did ES develop after the theoretical foundation and the user guidelines based on it became available? What is the cumulative technological and theoretical accomplishment on this subject in the 28 years since 1997?

Miroslav: For a subject ignored for three decades prior to 1997, the most unlikely thing happened over the subsequent three decades. Slowly but surely, more and more researchers got bitten by the same bug as I was in 1997. They saw useful or interesting things in ES, its use in their applications, with the utmost simplicity of implementation and almost indestructibility when used in conjunction with the theory provided or at least with common sense. The theorists and method developers also noted numerous opportunities to deploy their sophisticated tools and creative faculties into offering advances in the analysis or design of ES.

ES first grew by dozens, then hundreds, then thousands of new papers per year, exponentially. One currently finds about 20,000 papers since 1997.



Books on various directions in extremum seeking (deterministic, stochastic, for stabilization, and for PDEs) by Krstic and co-authors, from the period 2003–2022.

On the industrial side, there is hardly a branch of major technology in which ES has not been at least tried. About 3,000 patents partly testify to this. Any attempt to list the technologies or specific solutions would miss an order of magnitude more than it would mention. So I won't attempt to list "a few" technologies advanced by ES.

But I will mention one ES application that is already known to many: stabilization of extreme ultraviolet light (EUV) for semiconductor photolithography by Cymer, acquired by ASML, following this commercial and technological milestone. EUV's 200-fold increase of density of chips in 2013 was enabled by the discrete-time ES in this early article [8]. Today, in 2025, EUV is a US\$10 billion dollar per year industry, in light sources for photolithography alone, enabling hundreds of billions of dollars in chip manufacturing per year and making NVIDIA's and Apple's most advanced chips.

I would not even dream of trying to review the literature of the size to which ES has grown, tens of thousands of papers and several books. But somebody has already done this. The comprehensive survey by Alex Scheinker in *Automatica* [9] covers most of the achievements in the theory and applications of this method known to me. And allow me, please, to also mention a remarkable book by the guest editor of this special issue, Tiago Oliveira, on ES in infinite dimensions [10].

Anu: Your scientific roots, like mine, are in adaptive control. Is ES, in fact, a form of adaptive control? Or is it, instead, a form of optimal control, or optimization, or nonlinear control?

Miroslav: It is adaptive and nonlinear. And it performs optimization, in a model-free fashion, for dynamical systems.

As I prepare to assume the position of editor-in-chief of *Transactions on Automatic Control (TAC)* in January 2026, I have been reading papers from the journal's earliest years, 1956 onward. *TAC*, the first journal for control sys-

tems, was initially less theoretical than it is now. The gradual beginning of the theoretical style emerges in the early 1960s. In the 1950s, one sees, among other heuristic methods, considerable interest in ES, often referred to as "adaptive control." ES is indeed adaptive, pursuing its optimization goal in a model-free manner. And it is nonlinear even in its simplest manifestation, applying to maps with extrema. This *nonlinear* nature of *adaptation* made ES exciting in the 1950s and early 1960s, such a nature being absent from the "classical" linear methods predominating that period and tied to the Nyquist, Bode, root locus, and emerging LQR methods.

I cannot be sure, but I believe that it was clear to the ES pioneers in the 1950s what the analogy between ES and PID control was. Both feedback laws are model free. PID performs *regulation* to a setpoint, in the face of an unknown constant disturbance, using the integrator. ES, similarly, using its integrator, performs model-based *optimization*, that is, the regulation to an unknown optimizer. What is the difference? PID measures the feedback signal, the tracking error. ES, on the other hand, has to feed back, as an error, the gradient of the map, which is not measured but is meant to be driven to zero. It is with the perturbation and demodulation that this unknown gradient is estimated.

So, ES and PID are siblings. In the parlance of college aptitude tests, which emphasize analogies, "ES is to optimization what PID is to regulation." (The P and D actions, missing from basic ES, can always be added; in fact, ES with phase-lead compensators, emulating the PD action, has been introduced in the early 2000s.)

The interest in ES waned in the late 1960s as Lyapunov methods for model-based adaptive control began to emerge, lending themselves to theoretical proofs, while ES continued to defy that goal.

Anu: With the emergence of artificial intelligence (AI) and data-driven

control, it is often asked whether ES, which performs model-free control, is a form of data-driven control. Is ES "data-driven control"? In the era of AI, does it have any relation to machine learning (ML)?

Miroslav: Yes, ES is a form of data-driven control, and one can explain its relation to AI and ML with reasonable precision.

But first, let me note that the *sinusoid-driven* peak seeking extends back untold hundreds of millions of years. In botany, plant seedlings have evolved to use sinusoidal probing for climbing the *light gradients* (phototropism) during the first hours of growth. One also finds ES "discovered" by invertebrate animals, like sea urchins, whose sperm use sinusoidal action of the flagellum to perform circular motions and climb the gradient of the chemoattractant of an egg.


ES, as a human-conceived technology, is a method for model-free online optimization. Its distinguishing features are that 1) it is supported by theoretical guarantees and 2) it is applicable to dynamical systems, which is best exemplified by the use of ES in "source seeking" for nonholonomic autonomous vehicles, for which other methods, suitable for static optimization, and not envisioned for use in feedback with dynamical systems, would be at best artificial.

But, speaking more broadly, this method of model-free optimization, which employs only measurements, namely, data, is indeed a data-driven control method. It is also a learning-based control method, but it aims to learn only the optimizing input, not the entire maps.

So, is $ES \subset AI$? ES has puzzled many in the era of ML. It has defied categorization. Given that ES is a small-data (optimizer-seeking) rather than a big-data (model-learning) technology, is AI "a variant of AI" or a "mere" control technique, defined more by theorems than by broad, easy, and heuristic use?

(continued on p. 190)

IEEE Control Systems welcomes suggestions for books to be reviewed in this column. Please contact either Scott R. Ploen, Hong Yue, or Thomas Schön, associate editors for book reviews.



DATA-DRIVEN METHODS FOR DYNAMIC SYSTEMS

by JASON J. BRAMBURGER

Reviewed by Yorie Nakahira 

Data-Driven Method for Dynamic Systems by Prof. Jason J. Bramburger is a timely book that integrates the rigor of classical dynamical systems with the flexibility of modern computational tools. For readers who want to understand how mathematics, data, and computation come together to analyze time-

dependent systems, this book offers an insightful roadmap.

The opening chapter, “Dynamical Systems—Old and New,” introduces the field of dynamical systems analysis and the transformative influence of computers. It presents proper orthogonal decomposition and demonstrates how it provides a way of decomposing space–time data via an optimal basis expansion that is obtained using only the singular values. The chapter also outlines the book’s goal to explore such data-driven methods, showing how they complement traditional analysis by drawing inspiration from dynamical systems theory to interpret dynamic data.

The second chapter, “Linear Evolution Models,” thoughtfully describes methods for creating linear models of evolution from data collected from nonlinear systems, centered around dynamic mode decomposition (DMD). DMD is a data-driven technique that identifies a linear transformation to approximate the temporal evolution of system “snapshots,” enabling computationally efficient forecasting via matrix multiplication. The chapter details several extensions to address specific

challenges: physics-informed DMD, that constrains the matrix to incorporate known physical laws; sliding window DMD, that better handles multiscale data; and time-delay coordinates that help capture more information from low-dimensional measurements. These methods are presented as practical approximations of the underlying Koopman operator theory, which recasts nonlinear dynamics as an infinite-dimensional linear system acting on functions of the state space, known as *observables*. It further details how methods, such as extended DMD and kernel DMD, use dictionaries of observable functions to create finite-dimensional approximations of the Koopman operator.

The third chapter, “Identifying Nonlinear Dynamics,” transitions to the methods for identifying nonlinear dynamical systems from data. It begins with the sparse identification of nonlinear dynamics (SINDy) algorithm, which frames model discovery as a sparse regression problem and seeks to identify an accurate and parsimonious dynamical system that fits time-series data. The chapter then illustrates the application of the SINDy algorithm in learning nonlinear discrete-time mappings that iterate through coarse-grained solutions of continuous-time systems, to stabilize periodic orbits, and to discover conserved quantities.

The fourth chapter, “Data-Driven Polynomial Optimization,” demonstrates practical techniques for extracting high-level information from a system without first needing to identify an explicit model. This is done through polynomial optimization in data-driven methods. The core of this method lies in identifying auxiliary functions that can be used to produce the information or statistics about a system. It shows how to use sums of squares and semidefinite programming techniques to find Lyapunov functions, and extends it to the case of control Lyapunov function. It also shows the use of extended DMD in approximating Lie derivatives acting on dictionaries of observables, allowing one to extract information about a system directly from data without the intermediary step of model discovery.

The fifth chapter, “Learning Dynamics With Neural Networks,” introduces the use of neural networks for the analysis of dynamical systems. It begins with the concept of universal approximation theorems, which state that neural networks can approximate any continuous function to arbitrary accuracy. The text explains the architecture of neural networks, including layers, widths, and depths, and details the training process through minimizing a loss

Digital Object Identifier 10.1109/MCS.2026.3659234
Date of current version: 26 March 2026

function via gradient descent and backpropagation. It also introduces physics-informed neural networks (PINNs), which incorporate the governing differential equations of a system directly into the loss function to approximate solutions to ordinary differential equations and partial differential equations. The chapter concludes by exploring the use of PINNs for data-driven discovery of partial differential equations, highlighting the benefits of automatically calculating derivatives while cautioning about challenges, such as the nonuniqueness of potential models.

The final chapter, "Autoencoder Neural Networks," explores the use of autoencoder neural networks to discover simpler forms of dynamical systems. It first explains topological conjugacies and then demonstrates how the autoencoder neural network architecture can approximate conjugacies between systems. The chapter also demonstrates an innovative combination of autoencoders with the SINDy algorithm, in which the autoencoder performs dimensionality reduction, transforming the dynamics of a complex system into a low-dimensional space where a simple, interpretable model can be discovered.

Together, these chapters provide a systematic and modern overview of an evolving field. By covering techniques from DMD and sparse model discovery to

For readers who want to understand how mathematics, data, and computation come together to analyze time-dependent systems, this book offers an insightful roadmap.

polynomial optimization and neural networks, the book successfully illustrates how a new generation of computational approaches can be used to analyze, interpret, and simplify complex dynamical systems using data.

REVIEWER INFORMATION

Yorie Nakahira (ynakahir@andrew.cmu.edu) is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 USA.




Share Your Preprint Research with the World!

TechRxiv is a free preprint server for unpublished research in electrical engineering, computer science, and related technology. TechRxiv provides researchers the opportunity to share early results of their work ahead of formal peer review and publication.

BENEFITS:

- Rapidly disseminate your research findings
- Gather feedback from fellow researchers
- Find potential collaborators in the scientific community
- Establish the precedence of a discovery
- Document research results in advance of publication

Upload your unpublished research today!

 Follow us @TechRxiv_org
Learn more techrxiv.org

TechRxiv™
Powered by IEEE

Riotous Assembly in San Diego

Ninth IEEE Conference on Control Technology and Applications

Riotous is a glorious word to describe an energetic gathering of like-minded but diverse people with a common purpose—update while having a good time. Get the convivial and intellectual sides of a conference aligned and in synch, and the rest takes care of itself; the delegates make sure of that. Such was CCTA2025 in San Diego this past August at the Westin Gaslamp Quarter Hotel because of the following:

- » set in summer—long days, warm weather, nearby beaches
- » in an ideal contiguous and accommodating space on a single floor of adjacent session rooms, plenary room, registration desk, exhibits, meeting areas with wide corridors and good seating, and daylight, all within easy reach from the hotel rooms and coffee shops
- » with full-time access to the open-air rooftop Garden Terrace for tête-à-tête meetings and the venue for all receptions and dinners, each of which was buffet style with an open bar featuring the best of San Diego craft brewing products
- » moseying distance to the San Diego nightlife and restaurant areas of the Gaslamp Quarter, Little Italy, and Seaport Village
- » an easy commute from San Diego International Airport
- » supported by the San Diego Surf School to complete one's professional accreditation at Pacific Beach.

The physical side of the event was set. And, once comfortable, the attendees were treated to an attractive and varied technical program spread over two-and-a-half days, staggered to accommodate jetlag and attention spans by juggling end times, breaks, and the prospect of dining al fresco on the roof.

SUNDAY 24 AUGUST

- » Sunday saw three attractive workshops with attendance included in the registration fee.
 - “Learning, Estimation and Control for PDE Systems,” organized by Shu-Xia Tang, Mamadou Diagne, and Leobardo Camacho Solorio
 - “Advances to the Theory and Application of Aerospace Con-

trol,” organized by Andrea L’Affilitto and Gokhan Inhan.

- “Harmonic Control Arrays: Design, Simulation and Implementation,” organized by Murat Dogruel.
- » The workshops were exceptionally well attended, began at 9 a.m., and concluded by 5 p.m.
- » Sunday was also the first appearance of the Westin Gaslamp’s signature donut wall at the morning break ([Figure 1](#)).
- » The evening’s Welcome Reception on the rooftop was taco carts, of course ([Figure 2](#)), accompanied by local wines and beers. (The tireless general chair’s job includes selecting the drinks, naturally. This required quite some *research*.)



FIGURE 1 Finance Chair Jorge Cortés and CSS President Carolyn Beck being riotous in front of the Donut Wall at morning break.



FIGURE 2 Tacos on the rooftop Garden Terrace.

MONDAY 25 AUGUST

The conference proper comprised five parallel sessions plus plenaries each day. In retrospect, this was exactly right-sized for the participants, very Goldilocks with a good balance between choice and coverage. The attendees were the ones to voice this assessment.

- » Monday was gentle.
- » A pick-me-up plenary at 8:30 a.m. was presented by Liane Matthes (Figure 3) from ASML: “A Glimpse Into the Automation and Control Powering ASML’s EUV Light Sources.”
- » Morning sessions from 10 to 12 a.m. were followed by a student-industry mingle over lunch.
- » Anything but anodyne: a 1:30 p.m. plenary was presented by Guy



FIGURE 3 Liane Matthes and Stephanie Stocker.

Dumont from UBC: “Control Challenges Toward the Adoption of Closed-Loop Anaesthesia and Therapeutic Interventions as Standard of Care.”

- » Afternoon sessions concluded at 4:40 p.m. to permit the delegates some sunshine.
- » NextCom, the next generation of CSS leaders, who apparently like to play games—Jenga, Cornhole,

Ladder Golf—organized a catered session from 5:30 to 8 p.m. at the Embarcadero Marina Park, a short walk from the hotel.

- » The Recognition Dinner for the navvies, movers, and shakers (perhaps the past generation of CSS leaders) was held at Coasterra on San Diego Harbor on Monday evening (Figure 4).

TUESDAY 26 AUGUST

No rest for the wicked. The hard work began in earnest, accompanied by the Industrial Panel (Figure 5), followed by hard partying and follow-up back on the roof, armed with a pasta plate and a craft brew in the evening twilight [Figure 6(a)–(c)].

- » Alf Isaksson from ABB kicks off and regaled us with his morning plenary overview: “Industrial Control Systems: A Retrospective and the Road Ahead.”
- » Morning sessions ran from 10 a.m. to noon.
- » A lunchtime Industry Panel was organized and moderated by Industry Chair Mehmet



FIGURE 4 The San Diego skyline from Coasterra during the Recognition Dinner.



FIGURE 5 Industry Panel: Sam Crisafulli (ASML), GB Singh (Solar Turbines), and Alf Isaksson (ABB).

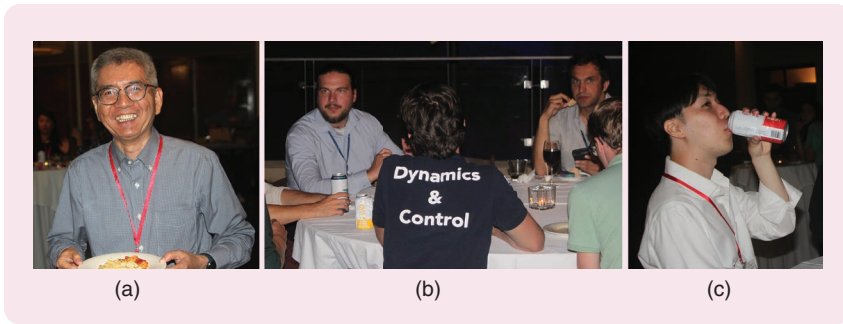


FIGURE 6 (a)–(c) Dinner after dark on the roof. That is CCTA2026 Vancouver General Chair Ryozo Nagamune in (a).

Mercangöz with speakers Sam Crisafulli of ASML, GB Singh of Solar Turbines, and Alf Isaksson of ABB. Attendance was excellent. So it was great that the event took place in the plenary space. Mehmet worked with the participants in advance to develop a cogent program focused on employment and training questions.

- » Coaxing them back from lunch, Allison Okamura from Stanford delivered the afternoon plenary: “Haptics Anywhere: Enabling Mobile Telepresence.”
- » Two afternoon sessions of six papers and then four papers rounded out the long day. The reward was a table of lemon cakes and pastries between the two stints. A welcome sugar hit.
- » Marathon Tuesday adjourned at 6:20 p.m., and we headed to the rooftop, drawn by the bar and crowned by the CCTA2025 dinner in the open air.

WEDNESDAY 27 AUGUST

Crescendo day! And a sprint to the finish.

- » Morning sessions ran from 8:30 to 10:30 a.m.
- » The final donut wall appeared.



FIGURE 7 The CCTA2025 bottle opener: utility in publicity, CDC2023 Singapore, home of Tiger Beer.

- » Thomas Parisini of Imperial College London bewitched the delegates with his Transition to Practice Award plenary, always a highlight of CCTA: “Digital Twins for a Greener and More Resilient Future in the Metal Industry.”
- » The CCTA2025 awards ceremony announced the Student Prize Papers and presented Thomas’s Transition to Practice Plaque.
- » We returned to the Garden Terrace for the farewell reception as the attendees either
 - a) Prepared to head to the airport and then home.
 - b) (Wet-)suited up for an afternoon surfing lesson, notably if you were an NTNU grad student.

PRE- AND POST-EVENT ACTIVITIES

The careful stage management of CCTA2025, of course, began a few years earlier. Bob Bitmead was invited to steer this ship forward after its original 2021 version went online during COVID-19. The hotel contract had been negotiated in 2019. The revamped Operating Committee (OpCom) was brought together a serious while before 2025, with Bob and indefatigable Program Chair Colin Jones of EPFL being approved by the CSS Board of Governors in 2022. From here, the team was accreted quickly and enthusiastically with the connivance of Stephanie Stockar, our diversity chair, to help us keep things blended.

The well-oiled machine that is OpCom is as follows:

- » Bob Bitmead (gopher)
- » Saverio Bolognani (website)
- » Airlie Chapmani (publications)
- » Jorge Cortés (finance)
- » Stefano di Cairano (TCEB)
- » Mamadou Diagne (exhibits)
- » Florian Dörfler (publicity)
- » Colin Jones (program)
- » Steffi Knorn (student activities)
- » Justin Koeln (invited sessions)
- » He Kong (registration)
- » Johanna Mathieu (workshops)
- » Mehmet Mercangöz (industry)
- » Pradeep Misra (digital systems)
- » Behrouz Touri (local arrangements)
- » Stephanie Stockar (diversity).

Publicity launched for CCTA2025 at CCTA2023 in Barbados, with follow-ups at CDC2023 and CDC2024. Florian and Saverio drove this forward via the website and written material. Bob designed the bottle opener, a task within his grasp, literally (Figure 7). The early invitations and acceptances from the plenary speakers were important too, to build excitement and promotion.

Early February 2025 was the closing date for paper submissions. Colin and Justin shifted into higher gear, working with Stefano on reviews, audits, ethics checks, and related tasks. Airlie was the interface with IEEE publications, which is the eventual publisher of presented papers. Her interactions included the bureaucratic efforts to keep IEEE on board and the papers on time. The program team then moved on to formulating the sessions and the program proper. This included Colin’s judgment on the five parallel sessions to maintain an attractive and manageable program.

Johanna wrangled the workshops from solicitation and selection through to logistics. She was responsible for the suggestion to include workshop participation in the registration fees. Bravo! This gave a strong impetus to student involvement in the workshops, ensuring their subscription. Steffi oversaw the student activities in total, from managing the CSS Student Travel

Awards with CSS Student Activities Chair Francesca Parise through to the implementation of the Student Best Paper Prizes. These tasks required dedication because following through in ensuring payments to the hotel is also required. Steffi provided the kinetic energy.

CCTA is the CSS applications conference focused on the nexus between industry practice and the intellectual juggernaut that is control theory. CCTA is the venue to nurture these connections and support practitioners in the field. So, the industry chair is central to keeping the focus. For CCTA2025, this was Mehmet Mercangöz, who exemplifies the connections between the two sides. Mehmet was a driver of the industrial panel and the student-industry mingle. He and Steffi helped maintain the introductions and interactions between the students and industrial attendees. ASML, our Platinum Sponsor, was delighted with how this operated.

Registration Chair He was a dynamo from Shenzhen in working with Pradeep on the registration process and in responding to the many queries. Because He was unable to travel, Mamadou stepped in as a Johnny-on-the-spot with the registrations ... and a lot of other things. Local Arrangements Chair Behrouz was the fixer of many things and general roustabout way beneath his talents. The registration desk was manned/womanned/studented/peopled by an array of willing helpers, including Bob's better half, Jan, and a cadre of UCSD grad students, from early Sunday morning to Wednesday midday. The envelope stuffing took place on Jan's and Bob's back deck (Figure 8). They know the alphabet really well now.

Post-conference, the work continued unabated or maybe even amplified (Figures 9 and 10). The publication process needs to proceed eventually to IEEE Xplore. Again, Airlie drove this to completion. Pradeep Misra was a font of assistance, decoding and wisdom



FIGURE 8 The Bitmead back deck on Friday, 22 August.



FIGURE 9 Workers Behrouz Touri, Steffi Knorn, Ari Partanen (an industry ring-in), and Carolyn Beck kick back late.

in all things digital systems oriented from submissions, reviews, decisions, uploads, registrations, billing, the online app, etc. As general chair of CDC 2006, also in San Diego, he really knows the drill.

So, who is the last man standing? Who is still slaving in the office when the rest of OpCom is in the pub as "last drinks" is called? Jorge Cortés, as finance chair, is that person. The accounts need to be finalized, tabulated, and reported to IEEE for their accounting and audit. We are hopeful that this finalized budget has been uploaded and will pass muster with the bean counters at IEEE headquarters. But wait! It gets worse. Jorge is also the general chair for CDC 2026 in Honolulu, and CCTA2025 is merely his apprenticeship. If you see



FIGURE 10 Allison Okamura and Bob Bitmead "Reclaim the Red Hat."

him in Honolulu, please buy him a beer. Of course, the same goes for the whole OpCom.

(continued on p. 195)

2025 IEEE CSS Awards

Every year, IEEE and the IEEE Control Systems Society (CSS) recognize the outstanding contributions by individuals in our technical community through a number of awards. The Society is very appreciative of the work that the respective subcommittee chairs and members devote to the selection process. The 2025 IEEE CSS Awards chair was Thomas Parisini, and the subcommittee chairs are as follows:

- » Denis Dochain, Antonio Ruberti Young Researcher Prize
- » Murli Salapaka, IEEE Control Systems Society Transition to Practice Award
- » Stephen Duncan, Control Systems Technology Award
- » Lars Gruene, George S. Axelby Outstanding Paper Award
- » Andrew J. Fleming, IEEE TCST Outstanding Paper Award
- » Mehran Mesbahi, IEEE Control Systems Magazine Outstanding Paper Award
- » Antonis Papachristodoulou, IEEE Control Systems Letters Outstanding Paper Award
- » Alessandro Astolfi, IEEE TCNS Outstanding Paper Award
- » Prashant Mehta, IEEE CDC Best Student Paper Award
- » Sophie Tarbouriech, Roberto Tempo Best CDC Paper Award
- » Naira Hovakimyan, IEEE CSS Award for Technical Excellence in Aerospace Control
- » Toru Namerikawa, IEEE CCTA Best Student Paper Award
- » Ravi Banavar, IEEE Control Systems Society Graduate Collaboration Fellowship.

More details about the CSS awards, the nomination process, and past recipients can be found on the IEEE CSS website: <http://ieeecss.org/awards/awards-program>.

CSS DISTINGUISHED MEMBER AWARDS

CSS annually confers Distinguished Member Awards to selected members of our community who have made significant technical contributions and provided outstanding long-term service to CSS. The 2025 awards recognize Amir Aghdam, Dennice Gayme, and Afef Fekih for their significant technical contributions and outstanding long-term service to CSS.

OUTSTANDING CHAPTER AWARD

This award recognizes a chapter for demonstrating a high level of activity, innovation, or growth. Vice-President of Member Activities Necmiye Ozay was responsible for the selection process for this award. The 2025 award went to the Control, Robotics, Cybernetics Joint Chapter of IEEE Vancouver Section (Canada), chaired by Dejan Kihias, in recognition of their efforts in organizing many controls-focused technical and social events for their membership in 2024.

CDC OUTSTANDING STUDENT PAPER AND BEST STUDENT PAPER AWARDS

The CDC Outstanding Student Paper and Best Student Paper Awards recognize excellence in a paper presented at the IEEE Conference on Decision and Control (CDC) whose primary author is a Student Member of IEEE. One of the Outstanding Student Paper awardees will be selected

as the winner of the Best Student Paper Award and will receive that award in lieu of the Outstanding Student Paper Award. The awards are based on the paper's originality, clarity, and potential impact on practical applications or theoretical foundations of control. Selected from 65 nominations, the CDC Outstanding Student Paper Award winners are as follows:

- » Joudi Hajar and Reza Ghane (advisor: Babak Hassibi) for "A Bregman Divergence Approach to H^∞ Control," by Joudi Hajar, Reza Ghane, and Babak Hassibi
- » Mohamad Shehab and Alperan Terkan (advisor: Necmiye Ozay) for "Efficient Reward Identification in Max Entropy Reinforcement Learning With Sparsity and Rank Priors," by Mohamad Shehab, Alperan Terkan, and Necmiye Ozay
- » Amir Shakouri (advisor: Kanat Camlibel) for "A New Perspective on Willems' Fundamental Lemma: Universality of Persistently Exciting Inputs," Amir Shakouri, Henk J. van Waarde, and M. Kanat Camlibel
- » Leonardo Toso (advisor: James Anderson) for "Learning Stabilizing Policies via an Unstable Subspace Representation," by Leonardo Toso, Lintao Ye, and James Anderson.

The Best Student Paper Award winners are as follows:

- » Joudi Hajar and Reza Ghane (advisor: Babak Hassibi) for "A Bregman Divergence Approach to H^∞ Control," by Joudi Hajar, Reza Ghane, and Babak Hassibi.

Figure 1 shows the CDC 2025 awardees.

CCTA OUTSTANDING STUDENT PAPER AND BEST STUDENT PAPER AWARDS

The CCTA Outstanding Student Paper and Best Student Paper Awards recognize excellence in a paper presented at the IEEE Conference on Control Technology and Applications (CCTA) whose primary author is a Student Member of IEEE. One of the Outstanding Student Paper awardees will be selected as the winner of the Best Student Paper Award and will receive that award in lieu of the Outstanding Student Paper Award. The awards are based on the paper's originality, clarity, and potential impact on practical applications of control. Selected from 24 nominations, the CCTA Outstanding Student Paper Award winners are as follows:

- » Philipp Santer (advisor: Knut Graichen) for "A Model Predictive Control Approach to Trajectory Tracking With Human-Robot Collision Avoidance," by Philipp Santer, Andreas Völz, and Knut Graichen
- » Yingjie Hu (advisor: Karl Berntorp) for "Optimal Measurement Projection in GNSS-RTK Factor Graph Optimization," by Yingjie Hu, Stefano Di Cairano, and Karl Berntorp
- » Felix Brändle (advisor: Melanie N. Zeilinger) for "On the Effects of Angular Acceleration in Orientation Estimation Using Inertial Measurement Units," by Felix Brändle, David Meister, Marc Seidel, Robin Strässer, and Frank Allgöwer.

The winner of the Best Student Paper Award is as follows:

- » Yingjie Hu (advisor: Karl Berntorp) for "Optimal Measurement Projection in GNSS-RTK Factor Graph Optimization," by Yingjie Hu, Stefano Di Cairano, and Karl Berntorp.

IEEE CONTROL SYSTEMS MAGAZINE OUTSTANDING PAPER AWARD

The IEEE Control Systems Magazine Outstanding Paper Award recognizes an outstanding article or column pub-



FIGURE 1 The CDC 2025 award winners.

lished during the two calendar years prior to the year of the award and is based on impact and benefit to CSS members. The 2025 award was given to the article "Quantum Computing Through the Lens of Control: A Tutorial Introduction," coauthored by Julian Berberich and Daniel Fink, *IEEE Control Systems*, vol. 44, no. 6, pp. 24–49, 2024.

IEEE TRANSACTIONS ON CONTROL SYSTEMS TECHNOLOGY OUTSTANDING PAPER AWARD

The IEEE Transactions on Control Systems Technology Outstanding Paper Award recognizes an outstanding paper published during the two calendar years prior to the year of the award and is based on originality, relevance of the application, clarity of exposition, and demonstrated impact on control systems technology. The 2025 award was given to the paper "A Modular Framework for Task-Agnostic, Energy Shaping Control of Lower Limb Exoskeletons," coauthored by Jianping Lin, Gray C. Thomas, Nikhil V. Divekar, Vamsi Peddinti, and Robert D. Gregg, *IEEE Transactions on Control Systems Technology*, vol. 32, no. 6, pp. 2359–2375, 2024.

IEEE CONTROL SYSTEMS LETTERS OUTSTANDING PAPER AWARD

The IEEE Control Systems Letters Outstanding Paper Award recognizes an outstanding paper published in *IEEE*



FIGURE 2 Jacquelin M. A. Scherpen

Control Systems Letters during the two calendar years preceding the year of the award, based on originality, potential impact on the theoretical foundations of control, importance and practical significance in applications, and clarity. The 2025 award was given to the paper "Temporal Forward-Backward Consistency, not Residual Error, Measures the Prediction Accuracy of Extended Dynamic Mode Decomposition," by Masih Haseli and Jorge Cortés, *IEEE Control Systems Letters*, vol. 7, pp. 649–654, 2023.

GEORGE S. AXELBY OUTSTANDING PAPER AWARD

The George S. Axelby Outstanding Paper Award recognizes an outstanding paper published in *IEEE Transactions on Automatic Control* during

the two calendar years prior to the year of the award and is based on originality, clarity, potential impact on the theoretical foundations of control, and practical significance in applications. The 2025 award was given to the paper “Nonlinear Opinion Dynamics With Tunable Sensitivity,” coauthored by Anastasia Bizyaeva, Alessio Franci, and Naomi Ehrich Leonard, *IEEE Transactions on Automatic Control*, vol. 68, no. 3, pp. 1415–1430, 2023.

IEEE TRANSACTIONS ON CONTROL OF NETWORK SYSTEMS OUTSTANDING PAPER AWARD

The IEEE Transactions on Control of Network Systems Outstanding Paper



FIGURE 3 Malcolm C. Smith

Award recognizes an outstanding paper published during the two calendar years prior to the year of the award and is based on originality, potential impact on the foundations of network systems, importance and practical significance in applications, and clarity. The 2025 award was given to the paper “The END: Estimation Network Design for Games Under Partial-Decision Information,” coauthored by Mattia Bianchi and Sergio Grammatico, *IEEE Transactions on Control of Network Systems*, vol. 11, no. 4, pp. 2200–2212, 2024.

ROBERTO TEMPO BEST IEEE CDC PAPER AWARD

This award is given in honor of Roberto Tempo, 44th president of CSS. The Tempo Award Committee selects the best paper from the previous year’s CDC based on originality; potential impact on any aspect of control theory, technology, or implementation; and clarity of writing. The 2025 award was given to the 2024 CDC paper “Average Number of Mistakes in Sequential Risk-Averse Scenario Decision-Making,” by Simone Garatti, Algo Carè, and Marco C. Campi, *63rd IEEE Conference on Decision and Control*, pp. 8245–8250, 2024.

IEEE CSS AWARD FOR TECHNICAL EXCELLENCE IN AEROSPACE CONTROL

The Award for Technical Excellence in Aerospace Control recognizes an outstanding paper or patented idea based on originality of technical innovation, significance/relevance to the aerospace community, aerospace application, and potential impact on the practice of aerospace engineering. The award can be conferred on an individual or a team. The 2025 award was given to Ilya V. Kolmanovskiy for contributions to constrained and model predictive control of aerospace systems.

CONTROL SYSTEMS TECHNOLOGY AWARD

The Control Systems Technology Award recognizes outstanding contributions to control systems technology, either in design and implementation or project management. This award can be conferred on an individual or a team. The 2025 award was given to José Luis Guzmán, Manuel Berenguel, F. Gabriel Acien, and Ana Sánchez Zurano for the application of automatic control as an enabling technology for sustainable bioprocesses based on microalgae in a blue economy.

IEEE CSS TRANSITION TO PRACTICE AWARD

The Transition to Practice Award recognizes outstanding collaborative scientific interactions between industry or research laboratories and academic communities that transition basic controls and system theory to practical systems for the benefit of society at large. The winner of the 2025 CSS Transition to Practice Award is Reza Moheimani for pioneering advancements in control of the scanning tunneling microscope and for the successful transfer to industry. The Transition to Practice Award comes with an invitation to deliver a plenary lecture at the 2026 IEEE CCTA.



FIGURE 4 Malcolm C. Smith receiving the 2025 Control Systems Award.

Malcolm C. Smith's Acceptance Speech

Thank you. It is a most extraordinary honor to be chosen as the 2025 IEEE Control Systems Award winner. It is something I never imagined. I would like to thank the IEEE Control Systems Society president, Carolyn Beck, for her very generous remarks and for being here to present me with this award. There are many others that I wish to thank, and my regret is that five minutes will be too short to name even a fraction of them.

It is tempting when looking back to rationalize one's career, but in truth, many of the most important things happen by chance. Opportunities that my parents never had, and their wise support, enabled me to be the first in my family to go to university. At Cambridge, three years of mathematics and lots of organ playing flashed by so quickly that I looked for an opportunity to stay for *just one more year*. The M.Phil. in control engineering and operational research seemed perfect, even though I had no idea what control engineering was. It was the charismatic figure of Alistair MacFarlane that drew me into the field, and I stayed on for the Ph.D.

The Cambridge group had state-of-the-art computing facilities, a single mainframe that could calculate *and* plot Nyquist diagrams in less than a minute, and more than one person could use it at once. The push in Cambridge was to generalize the classical framework for multivariable design; that early exposure to classical control was vital for me. But there was also the new kid on the block: *modern control*, the state-space methods used so successfully in the NASA program. There was appreciation of their mathematical elegance, but also awareness of possible limitations when there is substantial model uncertainty.

I gained much when I was younger from philosophical discussions about our subject: with Gerhard Kreisselmeier on adaptive control, the "two-cultures" debates with Peter Caines and George Zames, and on systems theory fundamentals with Jan Willems. I recall Zames saying on the origins of H-infinity control that he really wanted to do adaptive control but felt we needed to *understand* classical control first.

I am indebted to the irrepressible Allen Tannenbaum for drawing me to H-infinity control of time-delay systems at McGill, but even more for introducing me to Tryphon Georgiou. Later on, it was Tryphon's wonderful paper connecting the gap metric with H-infinity that led to our close friendship and collaboration. Those seem like golden years, as we honed our style and devoted ourselves to fundamental problems. Looking back at those papers, you can, just occasionally, detect the aroma of Greek coffee from the cafe in Oropos.

After a fruitful time at The Ohio State University, I returned to Cambridge as a lecturer after Keith Glover's appointment as the new professor of control engineering. Keith's inspired leadership created space and encouraged one to try to do good work. And I also gained much from exceptional colleagues in the group, such as Jan Maciejowski, Glenn Vinnicombe, and, later, Rodolphe Sepulchre.

During my first summer back in Cambridge, when everyone was on holiday, Tryphon was visiting, and we were working flat out to finish a paper before his return flight on the Saturday. A call came through *for Keith* from Paddy Lowe at the Williams Formula One team. The call was transferred to me. Paddy was looking for some help with an active suspension problem, and they had a test at Silverstone next Tuesday. I half-jokingly replied that my first free day was Sunday; it wasn't taken as a joke.

One is fortunate to find people to work with, like Paddy Lowe, and later on, Dick Glover at McLaren. Control has a lot to say when you are designing the plant as well as the controller. Paddy and Dick were superb engineers, but were also very open to "systems thinking."

The eventual ban on active suspension in Formula One did have an unexpected silver lining. For a control theorist, it seemed rather unremarkable to go from the class of all stabilizing controllers to the class of all passive impedances. And then the snag of realizing *all* positive real functions mechanically was just a stone in my shoe, holding up a theory paper. I had no idea of how much would follow from persisting with that question.

More than twenty years later, it still surprises me that people outside our field become impatient with the more abstract, systems explanation of the inerter. I even had some comments along the lines of "if you can't explain it more simply than that, you don't understand it." When news about the inerter broke in the Formula One world, grandprix.com published a newflash: "The first details were published in 2002 in the *obscure* Institute of Electrical and Electronics Engineers' publication called Transactions on Automatic Control."

When I look at academia now, and the control field in particular, I sometimes wonder if I would survive if I were starting out now. This is not a new feeling. Questioning the importance of our field is not new either. Forty years ago, at the 24th CDC in Fort Lauderdale, George Zames gave the acceptance speech for the Control Systems Award. He recalled that as one of the students at MIT in the late 1950s, Gordon Brown had just pronounced controls dead and ordered them to go into computers.

I feel very fortunate to have found controls as a field of work and research. Our subject contacts physics, biology, and the social sciences in a serious way. What is also in our favor is the very wide range of engineering applications and their contrasting character. The mathematical underpinnings are quite mature, but there is still great scope, in my opinion, for new problem formulations. This makes me optimistic about the future.

My time is up, and there are so many colleagues and former students whom I have not been able to thank by name. I am deeply indebted to all of them. I will close by thanking the CDC organizing committee for a most memorable conference here in Rio. It is an enormous privilege to be standing here.

Thank you.

—M. C. Smith

11 December 2025

ANTONIO RUBERTI YOUNG RESEARCHER PRIZE

The Antonio Ruberti Young Researcher Prize recognizes distinguished cutting-edge contributions by a young researcher to the theory or application of systems and control. The 2025 Ruberti Prize was given to Thinh T. Doan for fundamental contributions to the optimization and control theoretic foundations of machine learning.

IEEE CSS GRADUATE COLLABORATION FELLOWSHIP

The fellowship award seeks to promote and provide support to initiate collaborative research with a hosting faculty member of well-established institutions in the area of systems and control for outstanding graduate students in developing geographical regions by

visiting leading institutions in countries outside their own. The fellowship consists of a plaque and funding up to a maximum of US\$10,000 to defray costs associated with visiting host institutions. There were three recipients of the 2025 Fellowship Awards: Felipe Cinto will be hosted by Ricardo Sanfelice at the University of California, Santa Cruz, USA; Youssef Ait-Si will be hosted by Sadeh Soudjani at the Max Planck Institute, Germany; and Kaiyun Xie will be hosted by Gaoxi Xiao at NTU, Singapore.

HENDRIK W. BODE LECTURE PRIZE

The IEEE CSS Hendrik W. Bode Lecture Prize recognizes distinguished contributions to control systems science or engineering. The chair of

the Bode Lecture Prize was CSS Past President Magnus Egerstedt. The 2025 Bode Lecture Prize was given to Jacquelin M. A. Scherpen (Figure 2) for groundbreaking contributions to nonlinear balancing, model reduction, and passivity-based control.

IEEE CONTROL SYSTEMS AWARD

The IEEE Control Systems Award recognizes outstanding contributions to control systems engineering, science, or technology. The chair of the IEEE Control Systems Award for 2025 was Alberto Isidori. The 2025 IEEE Control Systems Award was given to Malcolm C. Smith (Figures 3 and 4) for contributions to robust control, network synthesis, and applications to mechanical systems.



» HISTORIES OF CONTROL (continued from p. 179)

Training of neural networks in ML is typically conducted by variants of stochastic optimization techniques. The stochastic algorithms of Robbins and Munro [11] and Kiefer and Wolfowitz [12] are of the same vintage as the sinusoid-driven peak seekers [5], meant for feedback use. ES doesn't use decaying gains and doesn't pursue unbiased convergence, because real systems, and their optimal operating points, are expected to drift in real time, so the optimizers must not "go to sleep." Apart from this difference, which results in ES convergence being practical but exponential and stochastic approximation (SA) having unbiased but nonexponential convergence, ES and SA are fundamentally in the same broad category of model-free optimization tools.

It is recognized that ES can be employed in training ML models. This is evidenced in the paper by Alexander Scheinker in this special issue. And ES

is, in fact, already being employed in training of large language models, by researchers in the Machine Learning Department at Carnegie Mellon University.

ES will see increasing integration with ML and AI in the future.

Anu: Thanks Miroslav, for speaking with me! I am certain the magazine readers will find this interview highly educational and inspirational.

Miroslav: My pleasure.

REFERENCES

- [1] M. Krstic, I. Kanellakopoulos, and P. V. Kokotovic, *Nonlinear and Adaptive Control Design*. New York, NY, USA: Wiley, 1995.
- [2] A. M. Annaswamy and A. F. Ghoniem, "Active control in combustion systems," *IEEE Control Syst. Mag.*, vol. 15, no. 6, pp. 49–63, Dec. 1995, doi: 10.1109/37.476386.
- [3] M. Leblanc, "Sur l'électrification des chemins de fer au moyen de courants alternatifs de fréquence élevée," *Rev. Gen. L'Electricite*, vol. 12, no. 8, pp. 275–277, 1922.
- [4] N. Minorsky, "Directional stability of automatically steered bodies," *J. Amer. Soc. Nav. Eng.*, vol. 34, no. 2, pp. 280–309, 1922.

[5] C. S. Draper and Y. T. Li, "Principles of optimizing control systems and an application to the internal combustion engine," in *Optimal and Self-Optimizing Control*, 1st ed., vol. 1, R. Oldenburger, Ed., Cambridge, MA, USA: MIT Press, 1951.

[6] K. J. Astrom and B. Wittenmark, *Adaptive Control*, 2nd ed. Reading, MA, USA: Addison-Wesley, 1995.

[7] M. Krstic and H.-H. Wang, "Design and stability analysis of extremum seeking feedback for general nonlinear systems," in *Proc. 36th IEEE Conf. Decis. Control*, vol. 2, pp. 1743–1748, 1997, doi: 10.1109/CDC.1997.657809.

[8] J.-Y. Choi, M. Krstic, K. Ariyur, and J.-S. Lee, "Extremum seeking control for discrete-time systems," *IEEE Trans. Autom. Control*, vol. 47, no. 2, pp. 318–323, Feb. 2002, doi: 10.1109/9.983370.

[9] A. Scheinker, "100 years of extremum seeking: A survey," *Automatica*, vol. 161, Mar. 2024, Art. no. 111481, doi: 10.1016/j.automatica.2023.111481.

[10] T. R. Oliveira and M. Krstic, *Extremum Seeking through Delays and PDEs*. Philadelphia, PA, USA: SIAM, 2022.

[11] H. Robbins and S. Monro, "A stochastic approximation method," *Ann. Math. Statist.*, vol. 22, no. 3, pp. 400–407, 1951, doi: 10.1214/aoms/1177729586.

[12] J. Kiefer and J. Wolfowitz, "Stochastic estimation of the maximum of a regression function," *Ann. Math. Statist.*, vol. 23, no. 3, pp. 462–466, 1952, doi: 10.1214/aoms/1177729392.



European Control Conference 2025

The 23rd European Control Conference (ECC25) was held from 24 to 27 June 2025 in Thessaloniki, Greece, at the Thessaloniki Concert Hall (Figure 1). The ECCs are organized under the auspices of the European Control Association (EUCA) and aim to bring together academic and industrial professionals in the field of systems and control and to promote scientific cooperation and exchanges within Europe and between Europe and other parts of the world. The ECC has been held biannually from 1991 to 2013 and annually thereafter. The 2025 local organization was the Aristotle University of Thessaloniki (AUTH). Markos Papageorgiou (Technical University of Crete, Greece) served as the general chair, and George Rovithakis (AUTH, Greece) and Zoe Doulgeri (AUTH, Greece) served as the general cochairs. The conference program chair was Christoforos Hadjicostis (University of Cyprus, Cyprus), with cochair Charalampos Bechlioulis (University of Patras, Greece). The

The conference featured nine exciting plenary and semiplenary lectures by renowned international experts in the field of control systems.

conference editorial board chair was Antonella Ferrara (University of Pavia, Italy). Technical cosponsorship of the event was provided by the IEEE Control Systems Society (CSS) and by the International Federation of Automatic Control (IFAC). Further details are available at <https://ecc25.euca-ecc.org/>.

TECHNICAL PROGRAM

ECC25 received 747 contributions from 59 countries. Considering the corresponding authors of those contributions, the top 10 contributing countries were as follows: Italy (67), Germany (64), France (53), the United States of America (43), The Netherlands (40), Sweden (31), Greece (27), China (22), the United Kingdom (21), and Denmark (17). In total, 1,789 authors and coauthors contributed to the program.

The top 10 contributing countries were as follows: Italy (246), Germany (188), France (187), The Netherlands (135), the United States of America (127), China (92), Sweden (84), Greece (75), the United Kingdom (68), and Denmark (65). The program included 460 contributed papers, 57 invited session papers, 17 tutorial papers, and six industry abstracts.

The theme of ECC25 was *Resilient Control Systems with Performance*. The conference featured nine exciting plenary and semiplenary lectures by renowned international experts in the field of control systems as follows:

- » “Robust Adaptive Control and the Greeks Who Made it Possible,” Miroslav Krstic, Distinguished Professor, University of California, San Diego, USA [Figure 2(a)]

Digital Object Identifier 10.1109/MCS.2026.3659281
Date of current version: 26 March 2026



FIGURE 1 (a) and (b) The Thessaloniki Concert Hall was the venue of ECC25. (c) The registration desk.



FIGURE 2 ECC25 plenary speakers. (a) Miroslav Krstic, (b) Yongduan Song, (c) Antonis Papachristodoulou, (d) Luca Zaccarian, (e) Frank Allgöwer, (f) Stefano Stramigioli, and (g) Jing Zhou.

- » “Control of Power Conversion Systems: Challenges, Theory, and Applications,” Dr. Tobias Geyer, corporate executive engineer, ABB System Drives, Switzerland
- » “Dissipativity, Optimal Control, and Turnpikes—How to Get From Willems to Wasserstein and Deep Learning,” Prof. Timm Faulwasser, School of Electrical Engineering, Computer Science and Mathematics, Hamburg University of Technology
- » “NN Learning-Driven Automatic Control and Automatic Control for Machine Learning,” Prof. Yongduan Song, Chongqing University, China [Figure 2(b)]

- » “Analysis and Design of Safe Control Laws,” Prof. Antonis Papachristodoulou, University of Oxford, U.K. [Figure 2(c)]
- » “Uncovering Reset-Endowed LTI Flow Dynamics in Hybrid Control of Mechatronic Systems,” Prof. Luca Zaccarian, LAAS-CNRS and the University of Trento, Italy [Figure 2(d)]
- » “Event-Based Control: When to Act, When to Wait, and Why It Matters,” Frank Allgöwer, director of the Institute for Systems Theory and Automatic Control and professor in the Mechanical Engineering Department, University of Stuttgart, Germany [Figure 2(e)]
- » “Port-Based Thinking in Modeling and Control of all Physical Systems,” Prof. Stefano Stramigioli, University of Twente, The Netherlands [Figure 2(f)]
- » “Adaptive Control With Quantization,” Prof. Jing Zhou, Department of Engineering Sciences, University of Agder, Norway [Figure 2(g)].

3.24 review reports per paper were collected. By the end of the review process, 433 papers graded B+ or B (or B-) were accepted and included in the program. An additional 84 papers graded C (out of 156) were also accepted based on their relevance to the conference sessions. In total, 517 contributed and invited session papers were accepted, representing an acceptance rate of 69.86%.

The program was organized into nine parallel tracks, each comprising from eight to 10 sessions with six presentations each. There were 76 contributed sessions and 11 invited sessions. The conference also featured four tutorial sessions on the following:

- a) “Dynamics and Control Enable New Incentive Schemes”
- b) “Homotopy Optimization and Hybrid Learning for Inference and Control”
- c) “Control-Theoretic Approaches to the Analysis and Design of Optimization Algorithms”
- d) “From Research to Reality in Control.”

To facilitate industry collaboration, ECC25 dedicated a session for six industry abstracts. Additionally, a plenary lecture by Dr. Tobias Geyer (corporate executive engineer,



FIGURE 3 Dr. Tobias Geyer delivering his plenary lecture.



FIGURE 4 The Women in Control Lunch Session.

ABB System Drives, Switzerland) on “Control of Power Conversion Systems: Challenges, Theory, and Applications” was presented (Figure 3). The program was further enriched by an industry panel discussion, “The Use of Data-Driven Techniques in Industry.”

The conference’s five lunch sessions offered attendees valuable networking opportunities in a relaxed environment. This year’s lunch events focused on themes relevant to both emerging and experienced professionals in control systems. A special Women in Control lunch took place on the final day, offering women a platform for networking, sharing perspectives on their roles in control engineering, and discussing opportunities and personal experiences (Figure 4).

On the day prior to the official opening, ECC25 hosted seven pre-conference workshops on cutting-edge topics such as algorithmic fairness, sparsity in control, reinforcement learning for robotics, systems theory of optimization, safe control, control for cyberphysical-human systems, and certified control synthesis for robotics. These workshops aimed to



FIGURE 5 EUCA President M. Elena Valcher presenting the 2025 European Control Award to Timm Faulwasser.



FIGURE 6 From right to left: (a) the finalists of the ECC25 Best Student Paper award and (b) the winner of the ECC25 Best Student Paper award.

foster discussion and collaboration, providing participants with advanced knowledge and practical skills.

EUCA AWARDS

During the banquet on 26 June, the awards ceremony took place. With the European Control Award, EUCA recognizes outstanding contributions by a young researcher in the area of systems and control. EUCA president M. Elena Valcher presented the 2025 European Control Award to Prof. Timm Faulwasser (Figure 5). The ECC25 Best Student Paper Award recognizes excellence in a paper presented at the ECC whose primary author is a student at the time of submission (Figure 6). The award was presented to Riccardo Brumali for “Distributed Learning and Optimization of a Multi-Agent Macroscopic Probabilistic Model.” Best Student Paper Award certificates were also issued to the finalists, Youssef AIT SI for “Symbolic Control: Unveiling Free Robustness Margins” and Julius Petrus Jakob Krebbekx for “Scaled Relative Graph Analysis of Lur’e Systems and the Generalized Circle Criterion.” The ECC25 Best Paper Award was presented to Igor Vladimirov and Ian R. Retersen for “Decoherence Time Control by Interconnection of Finite-Level Quantum Memory Systems”



FIGURE 7 Pierre-Alexandre Bliman receives the ECC25 Best Paper award.




FIGURE 8 Recognition of the organizing committee by EUCA President M. Elena Valcher. (a) Markos Papageorgiou, (b) George Rovithakis, (c) Zoe Doulgeri, (d) Christoforos Hadjicostis, and (e) Charalampos Bechlioulis.

and to Pierre-Alexandre Bliman for “Basic Offspring Number and Robust Feedback Design for the Biological Control of Vectors by Sterile

Insect Release Technique” (Figure 7). Also, EUCA President M. Elena Valcher recognized the organizing committee (Figure 8).

Conference attendees were rewarded with a welcome reception at the Thessaloniki City Hall (Figure 9). In addition, a conference banquet was held at the Ippikos Members Club (Figure 10). Many opportunities were provided for attendees to socialize during breaks (Figure 11). The student volunteers pitched in to keep the events running well (Figure 12).



FIGURE 9 Conference attendees at the welcome reception at the Thessaloniki City Hall.



FIGURE 10 Conference banquet at Ippikos Members Club.



FIGURE 11 (a) and (b) Socializing during a break.



FIGURE 12 Student volunteers.

SPECIAL ISSUE OF EUROPEAN JOURNAL OF CONTROL ON ECC25

Authors of 56 selected articles presented at ECC25 were invited to prepare an extended version of their work to be considered for publication in a special issue of *European Journal of Control* devoted to the conference. Prospective papers were chosen by Editor-in-Chief Thomas Parisini, based on the evaluations received by the EUCA-CEB.

ECC26

ECC25 was a successful event, maintaining high levels of engagement throughout. The next ECC26 will be held in Reykjavik, Iceland, from 7 to 10 July 2026. More information on ECC26 can be found at <https://ecc26.euca-ecc.org/>.

ACKNOWLEDGEMENT

The ECC25 organizers thank EUCA, IEEE CSS, and IFAC for their support and the opportunity to organize ECC25 in Thessaloniki. We extend special thanks to EUCA President Maria Elena Valcher, EUCA Secretary Federica Garin, and EUCA Past President Karl Henrik Johansson for their invaluable help in the conference organization and implementation. The organization of ECC25 required the efforts of many volunteers. All the members of the organizing committee contributed significantly to addressing all tasks and issues with

timeliness and expertise. A warm thank you to all of them for making the organization of ECC25 possible. ECC25 organizers wish to express our sincere gratitude to the sponsors—the AUTH, MATHWORKS, YOTTAVOLT, KENOTOM, ABB, and the *Journal of the Franklin Institute*—for their support. Sincere thanks to Pradeep Misra and his team at PeprCept. We also gratefully acknowledge Vicky Papadimitriou and the team of SYM VOLI (a conference organization company) for the superb management of the financial and logistical details of the conference. Last but not least, we would like to thank the many students and volunteers of the Aristotle University of Thessaloniki for their enthusiastic support during the event.

George Rovithakis
General Cochair, ECC25



» CONFERENCE REPORTS (continued from p. 185)

SIGNING OFF

You can appreciate from the previous discussion that conferences are spread over time and space way beyond their immediate location. The OpCom is central to a good event—and CCTA2025 was a great event—but there are many others who grease the wheels and deserve to be called out. The CSS Executive Committee and Board of Governors, or Committees and Boards, since there have been a number over the journey, have been singularly supportive and helpful.

The Westin Gaslamp Quarter was of enormous assistance in the preparation and logistics of CCTA2025. Given that we were resurrecting a contract

from 2019 and for 2021, one might have expected some friction in settling on terms and expectations. To a person, the crew at the hotel was accommodating of our organization and requirements. That made the general chair's life a whole lot simpler. His ignorance of several operational aspects was preserved intact.

2025 coincided with a new regime in the United States, or the reanimation of a bleak and discordant ancien régime, which caused problems for a number of people who were able (and willing) to travel to San Diego for CCTA2025. Program Chair Colin and Registration Chair He were among them. Nevertheless, and perhaps courtesy of training during

COVID, we were all able to operate effectively at a distance. Where speakers were unable to travel, they were invited to submit prerecorded videos of their presentations. There were seven such instances, and so not a great proportion. The session audiences were attentive and engaged for these presentations, and many communicated with the recorded presenters afterward. The success of the conference was unimpeded.

You should have been there!

Robert Bitmead
CCTA 2025 General Chair



In Memoriam: Neng (Eva) Wu, 1956–2025

Neng (Eva) Wu passed away on 19 September 2025 in Binghamton, NY, USA, at the age of 68. She was a professor in the Department of Electrical and Computer Engineering at Binghamton University, where she devoted 38 years of dedicated service to teaching, research, and the academic community. Eva was a respected member of the control systems community, known for her scholarly contributions, commitment to education, and generous collegial spirit.

Eva was born on 24 December 1956, and her early life and education were shaped by her family and her own experience during the Cultural Revolution in China. After graduating from high school, she was sent to the countryside for “reeducation,” an experience marked by severe hardship. Despite these challenges, she demonstrated resilience and an enduring love of learning that would define her life. Even during those years, she found ways to inspire and uplift others, often sharing stories and literature with her peers after long days of labor, as her lifelong friend Liqing Gu recalled.

Her academic journey formally began at the Northwestern Telecommunication Engineering Institute in Xi’an, China. She was a member of the class of 1977, the first cohort of college students admitted after the Cultural Revolution. She earned her B.S. in electrical engineering in 1982, specializing in signal processing. Later that year, Eva moved to the United States to pursue graduate studies at the University of Minnesota. She earned her M.S. in electrical engineering in 1983, with



Dr. Neng (Eva) Wu.

JOHNATHAN COHEN, BINGHAMTON UNIVERSITY

research focused on noise in semiconductor devices, and completed her Ph.D. in 1987 under the supervision of Prof. Bruce Lee, specializing in control science and dynamic systems.

After completing her doctoral studies, Eva joined the Department

of Electrical and Computer Engineering at Binghamton University in 1987. She remained there throughout her entire academic career, contributing through research, teaching, and service to the department and university. Her long-standing commitment reflected a deep sense of responsibility to the institution and to the broader academic profession.

Eva’s research focused on control and dynamic systems, with particular emphasis on fault detection and fault-tolerant control. In later years, her work expanded to include applications in power systems. Her scholarship was supported by funding from the National Science Foundation, the National Institutes of Health, and NASA and was recognized by the State University of New York Chancellor’s Award for Excellence in Scholarship and Creative Activities (2008–2009). Throughout her career, Eva approached research



Young Eva Wu (second from the right) in the countryside, 1976.

LIQING GU

problems with rigor and persistence, and she remained intellectually engaged and curious well into her later professional life.

Teaching was central to Eva's identity as an academic. She was deeply committed to undergraduate and graduate education and was known for her patience, clarity, and dedication to student development. Generations of students benefited from her careful mentoring, her high standards, and her genuine concern for their success. Colleagues frequently remarked on her willingness to share ideas, offer thoughtful feedback, and support others, making her a valued collaborator and mentor within her department and beyond. As her longtime colleague and former department chair, Prof. Douglas Summerville, reflected at her funeral, "She was deeply loved by her students, who admired not only her intellect but also her compassion and her genuine commitment to their learning. Her teaching extended far beyond the classroom, inspiring numerous students to approach research with curiosity, confidence, and—of course—the highest integrity."

Eva was also an active participant in the control systems community. She was a regular presence at professional conferences, including the American Control Conference (ACC) and the IEEE Conference on Decision and Control (CDC). Through these activities, she contributed not only technically but also through sustained professional engagement, fostering collegiality and connection within the field.

Beyond her academic pursuits, Eva found great joy in music and the performing arts. She was an active member of the Binghamton Downtown Singers and viewed music as an essential part of a balanced and meaningful life. Her enthusiasm for music reflected the same passion, discipline, and joy that characterized her professional endeavors.

I first met Eva at the 1991 ACC. At the time, we were both early in our ac-



Return to the University of Minnesota.

ademic careers, working in the same field during an era when women faculty in control engineering were still relatively few. From that initial meeting, we became longtime colleagues and friends, often sharing rooms at conferences such as the ACC and CDC. Over more than three decades, I benefited greatly from her collegiality, intellectual curiosity, and generosity of spirit.

Several qualities of Eva remain especially vivid to me. First was her warmth, expressed through her genuine smile and soothing voice, which had a remarkable ability to put others at ease. Second was her dedication: she gave fully to her work, her students, and her family, approaching every responsibility with care and seriousness. Third was her passion for learning and creativity. She was intellectually engaged, full of ideas, and deeply enthusiastic about research, teaching, and music. Finally, Eva was defined by her generosity and kindness. She gave her time, wisdom, and encouragement and was always more focused on supporting others rather than on personal recognition.

Eva Wu will be remembered as an accomplished scholar, a devoted educator, and a compassionate colleague. Her contributions to control systems research, her impact on generations of students, and her quiet but enduring presence in the professional com-



Eva Wu visiting San Francisco in 2024.

munity constitute a lasting legacy. She will be deeply missed by all who had the privilege of knowing her.

The author gratefully acknowledges Dennis S. Bernstein, Mogens Blanke, Jie Chen, Tongwen Chen, Guoxiang Gu, Pramod Khargonekar, Hitay Özbay, Mario Rotea, Jakob Stoustrup, and Kemin Zhou for sharing their memories of Eva Wu, as well as Eva's family and her childhood friend Liqing Gu for providing photographs.

Jing Sun^{ID}



IEEE CSS Board of Governors Meeting December 9, 2025

MOTIONS FOR 2026 APPOINTMENTS PREPARED BY CAROLYN BECK

CONSENT AGENDA

Motion to approve *new and continuing Standing Committee Chairs for 2026* (* Indicates *ex officio* positions).

- » **Endorsed by:** Executive Committee
- » **Financial Impact:** None.

Standing Committees

- » Awards [Parisini, Thomas](#) (25)
- » Awards Nomination Chair [Sepulchre, Rodolphe](#) (26)
- » Chapter Activities [Guoqiang, Hu](#) (24)
- » Conference Editorial Board (CDC, ACC) [Prieur, Christophe](#) (26)
- » Conference Editorial Board (CCTA) [Dabbene, Fabrizio](#) (26)
- » Conference Operations [Chong, Edwin](#) (19), [Jabbari, Faryar](#) (20)
- » Conference Publications [Galeani, Sergio](#) (22)
- » Distinguished Lecturer [Lall, Sanjay](#) (25)
- » Electronic Information [Pasqualetti, Fabio](#) (22)
- » Ethics in Publishing [Serrani, Andrea](#) (24)
- » Executive Committee [Beck, Carolyn](#) (25)
- » Fellow Evaluation [Seiler, Peter](#) (25)
- » Fellow Nominations [Papachristodoulou, Antonis](#) (24)
- » History [Padoan, Alberto](#) (26)
- » Industry Activities [Auret, Lidia](#) (25)
- » Long Range Planning [How, Jonathan](#) (26)*

- » Standards [Padariya, Riddhi](#) (25)
- » Student Activities [Uribe, César](#) (25)
- » Women in Control [Franco, Elisa](#) (26), [Pena, Patrícia](#) (26)
- » Diversity, Outreach & Development Task Force [Paré, Philip](#) (26)

Motion to approve new and continuing Long Range Planning Committee Members for 2026.

- » **Endorsed by:** Executive Committee
- » **Financial Impact:** None.

APPOINTED MEMBERS:

- » [Coogan, Sam](#) (26), [Knorn, Steffi](#) (26), [Prieur, Christophe](#) (26)

As an information item, the members *ex officio* are determined as of Bylaws Article IV, Section 6: Chair:

- » [Beck, Carolyn](#) (24)
- » [Ozay, Necmiye](#) (25)
- » [Ferrara, Antonella](#) (23)
- » [Jungers, Raphaël](#) (26)
- » [Serrani, Andrea](#) (24)
- » [Balakrishnan, Venkataraman](#) (Ragu) (26)
- » [Di Bernardo, Mario](#) (25)
- » [Mastellone, Silvia](#) (26)
- » [How, Jonathan](#) (23)

Motion to approve new and continuing Awards Subcommittee Chairs for 2026.

- » **Endorsed by:** Executive Committee
- » **Financial Impact:** None.

AWARDS SUBCOMMITTEES

- » Antonio Ruberti Young Researcher Prize [Broucke, Mireille](#) (26)

- » Control Systems Technology Award [Onori, Simona](#) (26)
- » Transition to Practice Award [Hovakimyan, Naira](#) (26)
- » George S. Axelby Outstanding Paper Award [Grüne, Lars](#) (24)
- » TCST Outstanding Paper Award [Ferrara, Antonella](#) (26)
- » CSM Outstanding Paper Award [Dochain, Denis](#) (26)
- » CDC Best Student Paper Award [Giordano, Giulia](#) (26)
- » CCTA Best Student-Paper Award [Faulwasser, Timm](#) (26)
- » TCNS Outstanding Paper Award [Sandberg, Henrik](#) (26)
- » CSS Aerospace Award [Batten, Belinda](#) (26)
- » Roberto Tempo Best CDC Paper Award [Tarbouriech, Sophie](#) (25)
- » CCTA Outstanding Paper Award [Ringwood, John](#) (26)
- » L-CSS Outstanding Paper Award [Vamvoudakis, Kyriakos](#) (26)
- » CSS Graduate Collaboration Fellowship Award [Banavar, Ravi](#) (25)
- » Bode Award Members (*ex-officio*): [Beck, Carolyn](#) (25), [How, Jonathan](#) (26), [Parisini, Thomas](#) (25)

Motion to approve new and continuing Technical Committee Chairs for 2026.

- » **Endorsed by:** Executive Committee
- » **Financial Impact:** None.

TECHNICAL COMMITTEES

- » Aerospace Controls [Inalhan, Gokhan](#) (22)

- » Automotive Controls [Katriniok, Alexander \(24\)](#)
- » Control Education [Hedengren, John \(21\)](#)
- » Discrete-Event Systems [Yin, Xiang \(23\)](#)
- » Distributed Parameter Systems [Diagne, Mamadou \(26\)](#)
- » Energy Systems [Almassalkhi, Mads \(23\)](#)
- » Health and Medical Systems [Pequito, Sergio \(24\)](#)
- » Hybrid Systems [Liu, Jun \(26\)](#)
- » Intelligent Control [Sarsilmaz, S. Burak \(26\)](#)
- » Manufacturing Automation & Robotic Control [Cao, Yongcan \(23\)](#)
- » Networks and Communication Systems [Dell'Anese, Emiliano \(25\)](#)
- » Nonlinear Systems and Control [Sanyal, Amit \(26\)](#)
- » Process Control [Mesbah, Ali \(21\)](#)
- » Quantum Computing, Systems, and Control [Dong, Daoyi \(24\)](#)
- » Robust and Complex Systems [Yong, Sze Zheng \(26\)](#)
- » Security and Privacy [Zhu, Quanyan \(21\)](#)
- » Smart Cities [Salazar, Mauro \(25\)](#)
- » Stochastic Systems and Control [Pasik-Duncan, Bozenna \(23\)](#)
- » Systems and Synthetic Biology [Singh, Abhyudai \(24\)](#)
- » Systems Identification and Adaptive Control [Piga, Dario \(26\)](#)
- » Variable Structure and Sliding Mode Control [Usai, Elio \(25\)](#)
- » Control and Optimization in Food and Agriculture [Garone, Emanuele \(26\)](#)

Motion to approve new and continuing Liaison Representatives for 2025.

- » **Endorsed by:** Executive Committee
- » **Financial Impact:** None.

LIAISON REPRESENTATIVES

- » AACC Director [Beck, Carolyn \(25\)](#)
- » AACC Alternate Director [How, Jonathan \(26\)](#)

- » EUCA [Valcher, Elena \(24\)](#)
- » IEEE Nanotechnology Council (2) [Ren, Juan \(18\)](#), [Moheimani, Reza \(24\)](#)
- » IEEE Systems Council (2) [Sipahi, Rifat \(22\)](#), [Yucelen, Tansel \(24\)](#)
- » IEEE Press [Serrani, Andrea \(24\)](#)
- » IEEE Society on Social Implications of Technology [Pasik-Duncan, Bozenna \(13\)](#)
- » IEEE Transactions on Image Processing [Serrani, Andrea \(24\)](#)
- » IEEE PAMI Editorial Board [Serrani, Andrea \(24\)](#)
- » IEEE Committee on Women in Engineering [TBD](#)
- » SIAM/SIAG [Demetriou, Michael \(14\)](#)
- » SICE (Inst & Control - Japan) [Fujimoto, Kenji \(22\)](#)
- » INFORMS [Nedich, Angelia \(15\)](#)
- » IEEE TAB Magazine Committee [Serrani, Andrea \(24\)](#)
- » IEEE TAB Newsletters Committee [Serrani, Andrea \(24\)](#)
- » IEEE TAB Transactions Committee [Serrani, Andrea \(24\)](#)
- » IEEE Smart Cities [Farid, Amro \(18\)](#)
- » IEEE Energy Systems (2) [Qu, Zhihua \(18\)](#), [Hiskens, Ian \(19\)](#)
- » Steering Committee of Life Science Technical Council (LSTC) [Jayawardhana, Bayu \(19\)](#)
- » IEEE Young Professionals [Noorani, Erfau \(25\)](#)
- » IEEE IoT Initiative [Cao, Yongcan \(24\)](#)
- » IEEE Biometrics Council [Zhu, Quanyan \(24\)](#)
- » IEEE Transportation Electrification Council [Siri, Silvia \(24\)](#)

Motion to approve new and continuing IEEE-USA Committee Representatives for 2026.

- » **Endorsed by:** Executive Committee
- » **Financial Impact:** None.

MEMBERS ON IEEE—USA COMMITTEES

- » Communication Policy [Sinopoli, Bruno \(21\)](#)
- » Energy Policy Committee [Moutis, Panos \(21\)](#)

- » Intellectual Property Committee [Jagannathan, Sarangapani \(19\)](#)
- » Transportation and Aerospace Policy [Inalhan, Gokhan \(22\)](#)

Motion to approve the continuing arXiv Moderators for 2026.

- » **Endorsed by:** Executive Committee
- » **Financial Impact:** None.
- » Chair: [Lovera, Marco \(20\)](#)
- » Moderators: [Ferrari, Riccardo \(20\)](#), [Shi, Guodong \(16\)](#), [Wang, Yuan \(11\)](#)

Biographies for new appointments are in Appendix I

Motion: Move to approve the new and continuing members of the Editorial Board for the IEEE Transactions on Automatic Control (TAC), effective January 1st 2026.

- » **Editor-in-Chief:** [Krstic, Miroslav \(26\)](#)

SENIOR EDITORS

- » [Cortes, Jorge \(24\)](#)
- » [Dullerud, Geir \(24\)](#)
- » [Guay, Martin \(25\)](#)
- » [Hadjicostis, Christoforos \(23\)](#)
- » [Ito, Hiroshi \(26\)](#)
- » [Nedich, Angelia \(23\)](#)
- » [Nesic, Dragan \(26\)](#)
- » [Yuksel, Serdar \(23\)](#)

ASSOCIATE EDITORS

- » [Amin, Saurabh \(23\)](#)
- » [Antunes, Duarte \(21\)](#)
- » [Arslan, Gurdal \(25\)](#)
- » [Avranchenkov, Konstantin \(24\)](#)
- » [Bako, Laurent \(24\)](#)
- » [Baldi, Simone \(26\)](#)
- » [Basile, Francesco \(25\)](#)
- » [Bekiaris-Liberis, Nikos \(26\)](#)
- » [Bobtsov, Alexey \(23\)](#)
- » [Bombois, Xavier \(23\)](#)
- » [Bresch-Pietri, Delphine \(22\)](#)
- » [Briat, Corentin \(26\)](#)
- » [Broucke, Mireille \(22\)](#)
- » [Charalambous, Themistoklis \(24\)](#)
- » [Chen, Jiming \(25\)](#)
- » [Chen, Tianshi \(24\)](#)
- » [Chen, Xiang \(23\)](#)
- » [Chen, Xinkai \(22\)](#)
- » [Como, Giacomo \(26\)](#)

- » Dani, Ashwin (25)
- » Faulwasser, Timm (21)
- » Franceschelli, Mauro (25)
- » Gao, Chuanhou (21)
- » Halder, Abhishek (26)
- » Huang, Minyi (26)
- » Huang, Yulong (25)
- » Jones, Colin (23)
- » Kamgarpour, Maryam (25)
- » Kan, Zhen (21)
- » Kishida, Masako (24)
- » Lazar, Mircea (21)
- » Lee, Ti-Chung (24)
- » Lestas, Ioannis (23)
- » Li, Tao (26)
- » Li, Wuquan (24)
- » Liu, Tengfei (23)
- » Mahajan, Aditya (22)
- » Malikopoulos, Andreas (21)
- » Mironchenko, Andrii (26)
- » Mohajerin Esfahani, Peyman (23)
- » Moreno, Jaime (22)
- » Necoara, Ion (22)
- » Olshevsky, Alexander (21)
- » Pavel, Lacra (23)
- » Pola, Giordano (24)
- » Rapaport, Alain (22)
- » Oliveira, Tiago Roux (26)
- » Rovithakis, George (23)
- » Sassano, Mario (26)
- » Scruggs, Jeff (25)
- » Seiler, Peter (25)
- » Shen, Dong (25)
- » Shi, Yang (22)
- » Shu, Zhan (21)
- » Tanwani, Aneel (21)
- » Trélat, Emmanuel (23)
- » Vamvoudakis, Kyriakos (24)
- » Wang, Bohui (26)
- » Wang, Ji (26)
- » Wang, Yongqiang (22)
- » Xu, Xiangru (26)
- » Yang, Jun (26)
- » Yucelen, Tansel (21)
- » Zamani, Majid (24)
- » Zeilinger, Melanie (25)
- » Zemouche, Ali (21)
- » Zhao, Xueyan (26)
- » Zhou, Bin (25)
- » Zhou, Jing (24)
- » Zhu, Minghui (24)

- » **Financial Impact:** None
- » **Endorsed by:** Executive Committee

Biographies for new TAC appointments are listed in Appendix II.

Motion: Move to approve the new and continuing members of the Editorial Board for the IEEE Transactions on Control Systems Technology (TCST), effective January 1st 2026.

- » **Editor-in-Chief:** Kolmanovsky, Ilya (24)

SENIOR EDITORS

- » Abbaszadeh, Masoud (26)
- » Braslavsky, Julio (23)
- » Ferrari-Trecate, Giancarlo (24)
- » Glielmo, Luigi (26)
- » Gravdahl, Jan Tommy (26)
- » Hatanaka, Takeshi (24)
- » Hellstrom, Erik (25)
- » Kellett, Christopher (25)
- » Loria, Antonio (26)
- » Maggiore, Manfredi (25)
- » Stoustrup, Jakob (24)
- » Toth, Roland (24)

ASSOCIATE EDITORS

- » Berntorp, Karl (24)
- » Boiko, Igor (21)
- » Borhan, Ali (25)
- » Bradley, Justin (25)
- » Capello, Elisa (23)
- » Consolini, Luca (26)
- » Corno, Matteo (26)
- » Danielson, Claus (25)
- » Delle Monache, Maria Laura (25)
- » Durand, Helen (26)
- » Evangelou, Simos (21)
- » Fekih, Afef (26)
- » Fourati, Hassen (24)
- » Hamed, Kaveh Akbari (25)
- » Haskara, Ibrahim (25)
- » Hoelzle, David (25)
- » HosseinNia, Hassan (26)
- » Ishizaki, Takayuki (25)
- » Jones, Colin (22)
- » Karimi, Alireza (26)
- » Katrinok, Alexander (25)
- » Lee, Taeyoung (23)
- » Lefeber, Erjen (25)
- » Li, Shengbo Eben (26)
- » Li, Zhaojian (25)
- » Macchelli, Alessandro (21)
- » Mammarella, Martina (25)
- » Maruta, Ichiro (26)
- » Meurer, Thomas (25)

- » Mesbah, Ali (23)
- » Nielsen, Christopher (23)
- » Onori, Simona (23)
- » Orosz, Gabor (21)
- » Ossareh, Hamid (26)
- » Pavlov, Alexey (26)
- » Phillips, Sean (25)
- » Pisano, Alessandro (25)
- » Prodan, Ionela (24)
- » Puig, Vicenç (23)
- » Rabbath, Camille Alain (26)
- » Romagnoli, Raffaele (26)
- » Rosolia, Ugo (25)
- » Ruderman, Michael (25)
- » Santini, Stefania (25)
- » Sato, Masayuki (23)
- » Scott, Joseph (26)
- » Siegel, Jason (24)
- » Toffanin, Chiara (24)
- » Velni, Javad (25)
- » Worthmann, Karl (24)
- » Yang, Chenguang (21)
- » Yang, Tao (21)
- » You, Fengqi (21)

- » **Financial Impact:** None
- » **Endorsed by:** Executive Committee

Biographies for new TCST appointments are listed in Appendix III.

Motion: Move to approve the continuing members of the Editorial Board for the IEEE Transactions on Control of Network Systems (TCNS), effective January 1st 2026.

- » **Editor-in-Chief:** Pavel, Lacra (25)

SENIOR EDITORS

- » Como, Giacomo (23)
- » Dey, Subhrakanti (20)
- » Dimarogonas, Dimos (23)
- » Gasparri, Andrea (26)
- » Lestas, Ioannis (25)
- » Li, Na (26)
- » Marden, Jason (25)
- » Russo, Giovanni (24)
- » Schenato, Luca (19)
- » Shames, Iman (25)
- » Wang, Yongqiang (26)
- » Zheng, Wei Xing (24)

ASSOCIATE EDITORS

- » Ahn, Hyo-Sung (24)

- » Alizadeh, Mahnoosh (24)
- » Cai, Kai (25)
- » Chen, Lijun (24)
- » Consolini, Luca (25)
- » Cristofaro, Andrea (26)
- » Dall'Anese, Emiliano (26)
- » Farokhi, Farhad (24)
- » Frasca, Paolo (26)
- » He, Jianping (22)
- » Hovakimyan, Naira (23)
- » Kia, Solmaz (22)
- » Kibangou, Alain (22)
- » Laleg, Meriem (24)
- » Liu, Ji (26)
- » Margellos, Kostas (26)
- » Martins, Nuno (24)
- » Mo, Yilin (23)
- » Motee, Nader (24)
- » Mojica-Nava, Eduardo (23)
- » Ocampo-Martinez, Carlos (23)
- » Pan, Ya-Jun (25)
- » Panagou, Dimitra (23)
- » Parisio, Alessandra (22)
- » Quevedo, Daniel (24)
- » Qin, Jiahu (22)
- » Reveliotis, Spyros (23)
- » Rizzo, Alessandro (23)
- » Scardovi, Luca (25)
- » Sundaram, Shreyas (24)
- » Taha, Ahmad (22)
- » Taylor, Josh (21)
- » Touri, Behrouz (23)
- » Wang, Xiaofeng (26)
- » Wen, Guanghui (25)
- » Yang, Tao (24)
- » You, Keyou (21)
- » Zelazo, Daniel (26)
- » Zhu, Hao (24)

- » **Financial Impact:** None
- » **Endorsed by:** Executive Committee

*Motion: Move to approve the **new** and **continuing** members of the Editorial Board for the IEEE Control Systems Letters (L-CSS), effective January 1st 2026.*

- » **Editor-in-Chief:** Menini, Laura (25)

SENIOR EDITORS

- » Aguiar, Pedro (23)
- » Baldi, Simone (25)
- » Giordano, Giulia (25)

- » Girard, Antoine (25)
- » **Liu, Shu-Jun** (26)
- » Loria, Antonio (25)
- » Manzie, Chris (23)
- » Morarescu, Irinel-Constantin (25)
- » Olaru, Sorin (23)
- » Oomen, Tom (23)
- » **Pepe, Pierdomenico** (26)
- » Savla, Ketan (23)
- » Tesi, Pietro (23)
- » Vazquez, Rafael (26)
- » Z

ASSOCIATE EDITORS

- » Ahmed Ali, Tarek (23)
- » Airimitoae, Tudor (25)
- » Albea, Carolina (26)
- » Auriol, Jean (26)
- » Back, Juhoon (21)
- » Besselink, Bart (23)
- » Boem, Francesca (25)
- » Borri, Alessandro (26)
- » Braun, Philipp (23)
- » **Bridgeman, Leila** (26)
- » Busoniu, Lucian (26)
- » Califano, Claudia (24)
- » Castaldi, Paolo (26)
- » Cherukuri, Ashish (24)
- » Chong, Michelle (24)
- » Coogan, Samuel (23)
- » Davila, Jorge (23)
- » Del Vecchio, Carmen (25)
- » **Di Ferdinando, Mario** (26)
- » Dinh, Thach N. (24)
- » Fang, Huazhen (22)
- » Fabiani, Filippo (25)
- » **Feng, Zhiguang** (26)
- » Ferrante, Francesco (22)
- » Fiacchini, Mirko (25)
- » Fiore, Davide (25)
- » Fosson, Sophie (24)
- » Garin, Federica (22)
- » Gonzalez, Alejandro (26)
- » Haesaert, Sofie (25)
- » Hale, Matthew (25)
- » Hori, Yutaka (25)
- » Iervolino, Raffaele (25)
- » Incremona, Gian Paolo (25)
- » Jungers, Marc (23)
- » Kamalapurkar, Rushikesh (24)
- » Kashima, Kenji (23)
- » Kergus, Pauline (25)
- » Knorn, Steffi (22)
- » Lavaei, Abolfazl (26)
- » Liu, Lei (26)

- » Liu, Zhixin (22)
- » Lucia, Walter (25)
- » Maghenem, Mohamed (26)
- » Maity, Dipankar (25)
- » Manchester, Ian (23)
- » Massioni, Paolo (23)
- » ~~Matni, Nikolai~~ (25)
- » Mattioni, Mattia (24)
- » Mazo, Manuel Jr. (23)
- » Mehr, Negar (26)
- » Mylvaganam, Thulasi (23)
- » Natarajan, Vivek (22)
- » **Ni, Yuan-Hua** (26)
- » **Nicotra, Marco** (26)
- » Nuño, Emmanuel (25)
- » Nurdin, Hendra (24)
- » Ocampo-Martinez, Carlos (25)
- » Oliva, Gabriele (22)
- » Pal, Debasattam (26)
- » Palumbo, Pasquale (26)
- » Possieri, Corrado (24)
- » Poveda, Jorge (23)
- » Pu, Shi (26)
- » Raffo, Guilherme (26)
- » Rapisarda, Paolo (24)
- » Ravazzi, Chiara (21)
- » Rios, Hector (26)
- » Rizzo, Alessandro (23)
- » Oliveira, Tiago Roux (24)
- » Roy, Spandan (25)
- » Sassano, Mario (23)
- » Schmid, Robert (21)
- » Schoukens, Maarten (23)
- » Siri, Silvia (23)
- » Ticozzi, Francesco (22)
- » Tong, Yin, (24)
- » Trimpe, Sebastian (23)
- » Van Waarde, Henk (24)
- » **Wang, Bing-Chang** (26)
- » Wang, Lin (25)
- » Wang, Yuanzhe (25)
- » Wang, Zheming (25)
- » Wardi, Yorai (21)
- » **Wisniewski, Rafal** (26)
- » Xiang, Weiming (23)
- » **Yang, Tao** (26)
- » Yin, Xiang (23)
- » Yong, Sze Zheng (26)
- » Yuan, Shuai (25)
- » Zhu, Yanzheng (23)
- » Zino, Lorenzo (26)

- » **Financial Impact:** None
- » **Endorsed by:** Executive Committee

Biographies for new L-CSS appointments are listed in Appendix IV.

Motion: Move to approve the continuing members of the Editorial Board for the IEEE Open Journal of Control Systems (OJCS), effective January 1st 2026.

- » Editor-in-Chief: Martinez, Sonia (21)

SENIOR EDITORS

Smith, Steven (25)
Seiler, Peter J. (25)

ASSOCIATE EDITORS

- » Briat, Corentin (22)
- » Canale, Massimo (22)
- » Carli, Ruggero (25)
- » Gambuzza, Lucia Valentina (24)
- » Iannelli, Luigi (22)
- » Innocenti, Giacomo (24)
- » Kishida, Masako (25)
- » Le Ny, Jerome (26)
- » Li, Xiuxian (26)
- » Liu, Jun (25)
- » Mathieu, Johanna (25)
- » Mehr, Nigar (25)
- » Miao, Fei (25)

- » Nowzari, Cameron (25)
- » Ocampo-Martinez, Carlos (21)
- » Ramirez Llanos, Eduardo Jose (26)
- » Reniers, Michael (22)
- » Oliveira, Tiago Roux (22)
- » Speranzon, Alberto (21)
- » Srivastava, Vaibhav (25)
- » Tedesco, Francesco (26)
- » Tallapragada, Pavankumar (21)
- » Wang, Gang (22)
- » Zheng, Wei Xing (26)
- » Zhu, Minghui (21)

- » Financial Impact: None
- » Endorsed by: Executive Committee

Motion: Move to approve the continuing members of the Editorial Board for the IEEE Control Systems Magazine, effective January 1st 2026.

- » Editor-in-Chief: Annaswamy, Anuradha (25)
- » Senior Technical Writer: Norton, Sandra (25)
- » Co-Associate Editors for Book Reviews: Ploen, Scott (13); Georgiou, Tryphon (25); Sojoudi, Somayeh (25)

- » Associate Editor (History): Vacant
- » Associate Editor (Education): Pates, Richard (23)
- » Corresponding Editor (Conference Activities): Rizvi, Syed Zeeshan (21)

ASSOCIATE EDITORS

- » Abramovitch, Daniel (24)
- » Alizadeh, Mahnoosh (25)
- » Banavar, Ravi (20)
- » Bazanella, Alexandre (20)
- » Bonnabel, Silvere (21)
- » Chen, Jie (22)
- » Engelberg, Shlomo (21)
- » Fregene, Kingsley (25)
- » Jovanovic, Mihailo (24)
- » Nowzari, Cameron (25)
- » Pates, Richard (23)
- » Sepulchre, Rodolphe (25)
- » Wang, Gang (25)
- » Xie, Lihua (23)
- » Zeilinger, Melanie (22)

- Financial Impact: None
- Endorsed by: Executive Committee

Minutes of the Control Systems Society Board of Governors Meeting December 9, 2025 Rio de Janeiro, Brazil

CALL TO ORDER AND APPROVAL OF AGENDA

President Carolyn Beck welcomed all attendees to CDC 2025 and initiated all-participant introductions. She called the meeting of the IEEE Control Systems Society (CSS) Board of Governors (BoG) to order at 1:17 p.m. BRT on December 9, 2025. She then re-

mind the BoG about the objectives of the meeting. Meeting procedures were reviewed, including the mechanism for voting on motions. The following members of the BoG were in attendance:

- » Beck, C.
- » Coogan, S.
- » Di Bernardo, M.
- » Dong, D.
- » Egerstedt, M. (late)
- » Fekih, A.
- » Ferrara, A.

- » Fregene, K.
- » How, J.
- » Isaksson, A.
- » Johansson, K.
- » Jungers, R.
- » Mesbahi, M
- » Nair, G.
- » Ozay, N.
- » Palycarpou, M.
- » Pena, P.
- » Rudie, K.
- » Sandberg, H.

- » Serrani, A.
- » Vernekar, P. (online)
- » Xie, L.

Additionally, the following visitors attended the meeting:

- » Aghdam, A.
- » Allgöwer, F.
- » Annaswany, A.
- » Baillieul, J.
- » Balakvishnan, V.
- » Basilo, J. C.
- » Fromm, J.
- » Kolmanovsky, I.
- » Krstić, M.
- » Paschalidis, Y.
- » Pavel, L.
- » Prandini, M.
- » Prieur, C.
- » Zaccarian, L.

The agenda is available at <http://bog-excom.ieeecss.org/>.

The agenda was approved unanimously.

Before the official meeting began, the CDC 2025 General Chairs, João Carlos Basilio and Yannis Paschalidis, provided brief opening remarks regarding the conference, including expressions of appreciation to attendees.

President Beck clarified the expectations for BoG members:

She reminded BoG members of their commitment to serve by attending two in-person meetings per year:

- » ACC or (or ECC or CCTA) and CDC.

She emphasized the importance of thoroughly reading the materials provided prior to the meeting. She also provided an overview of meeting procedures, including the use of *Robert's Rules of Order* and the process for voting on motions using name tents. Background information was given on how these rules were adopted by the Society.

Logistical Details

The following logistical items were reviewed:

- » Quorum: Established
- » Voting Method: Name tents were used to facilitate voting
- » Consent Agenda: Introduced as part of efficient meeting management

» Voting Requirement: A majority of present members are required to pass simple motions

» Travel Support: CSS provides \$1,200 for domestic travel and \$1,600 for international travel. Members must retain and submit all receipts via Concur for reimbursement.

Beck introduced a revised meeting format, already adopted in the previous two years, beginning with a financial summary to better inform the group about the current financial standing and set the stage for strategic decisions.

Approval of Minutes

Quorum was established, and the meeting agenda approved. Next, a motion to approve the minutes of the BoG meeting from July 7, 2025 (Denver, Colorado) was presented.

Seconded by: Fekih

Abstention: none

The motion passed unanimously.

Approval of Consent Agenda

Motion to approve the consent agenda

Called by Johansson

Seconded by: Xie

The motion passed unanimously.

ACTION ITEMS

VP-Finance Activities: Jonathan How

CSS Financial—Financial Overview:

How began with a dual-purpose presentation: to report on the Society financial status and to provide that data in a clear and engaging format.

How reported the following financial details:

- » Overall, CSS financial status is very healthy.
- FM14 23: Revenue: **\$8,653K**, Expenses: **\$7,131K**, Surplus: **\$1,488K** (budget **\$603K**)
- FM14 24: Revenue: **\$8,337K**, Expenses: **\$6,940K**, Surplus: **\$1,398K** (budget **\$388K**)

How reviewed what these figures represent for the Society financial health and discussed future fiscal planning.

Conference Financials:

CDC 2024

- » Surplus: €143K
- » Additional Expenses: Over €150K (plus €70K for student hotel costs)
- » CSS Share of Costs: Approx. €90K
- » Obligated 2025 Initiatives: Already exceed \$355K

Reserves and Society Funds:

» Reserves (End of 2023): \$27,133K (up 14% from 2022)

» Society Fund (End of 2024): \$732K

How explained that while the reserves serve as a cushion, funds in the reserve are difficult to access directly.

How noted that IEEE rules on finances have changed significantly since he assumed the role. He emphasized that the majority of revenue is generated through journals and conferences.

CSm was identified as operating at a loss due to publication costs; however, it was emphasized that this is an intentional investment made in service to the Society.

Motions:

President Beck presented the following motion:

- » **Motion:** *To fund the CSS Family Care Grant Program for years 2026–2029.*

Financial Impact: Up to \$20K per year for 4 years

Endorsed By: Executive Committee

It was noted that “DODA” appears in blue on the slide, as the program name is likely to change. President Beck provided background on the grant and outlined application guidelines. Priority will be given to applicants who have not previously received the grant, as well as to those presenting at conferences.

Fregene asked how the funds have been used and directed the question to Johansson, however Beck clarified that this is a new grant under the Dependent Care Act and noted that the program has not reached its funding limit. She explained that approval of funds by the BoG is required before submitting the proposal to IEEE and

210230-Control Systems Summary Report For Period Ending Oct-25	Total Revenue			Total Expense			Net		
	Annual Budget	YTD Budget	YTD Actual	Annual Budget	YTD Budget	YTD Actual	Annual Budget	YTD Budget	YTD Actual
All Products									
Magazines	169K	143K	149K	429K	330K	366K	259K	188K	217K
Journals Transactions Rev	3,296K	2,763K	2,758K	2,264K	1,880K	2,019K	1,032K	883K	739K
Letters	566K	472K	521K	408K	336K	384K	158K	135K	136K
Periodical Packaged Prod	1K	1K	1K	6K	5K	5K	5K	4K	4K
Conference Events	1,241K	297K	229K	1,007K	219K	218K	234K	78K	11K
Conference Proceeding	2,662K	2,219K	2,319K	1,540K	1,283K	1,455K	1,122K	935K	864K
Society Membership	135K	135K	139K	4K	4K	4K	131K	131K	135K
Total Product	8,070K	6,029K	6,116K	5,658K	4,059K	4,451K	2,412K	1,970K	1,665K
Cost Centers									
ExCom	0K	0K	0K	0K	0K	128K	0K	0K	128K
Membership Committee	0K	0K	0K	2K	0K	2K	2K	0K	2K
Awards Committee	0K	0K	0K	8K	1K	1K	8K	1K	1K
Executive Office	0K	0K	0K	0K	0K	0K	0K	0K	0K
Societies Operations	0K	0K	0K	1,789K	1,331K	1,146K	1,789K	1,331K	1,146K
Meetings /Conference	18K	15K	70K	243K	190K	236K	225K	175K	166K
Total Cost Center	18K	15K	70K	2,042K	1,522K	1,512K	2,024K	1,507K	1,443K
Total From Operations	8,088K	6,044K	6,186K	7,699K	5,581K	5,964K	388K	463K	222K
Grand Total	8,088K	6,044K	6,186K	7,699K	5,581K	5,964K	388K	463K	222K

	Revenue		Expense		Net	
	budget	actual	budget	actual	budget	actual
Control Systems Magazine	143K	149K	(330K)	(366K)	(188K)	(217K)
Trans on Automatic Control	1,797K	1,829K	(1,162K)	(1,196K)	636K	633K
Trans on Control Systems	631K	600K	(436K)	(474K)	195K	125K
Transactions on Control of	300K	303K	(233K)	(230K)	68K	73K
Open Journal of Control Sy	34K	27K	(49K)	(119K)	(15K)	(92K)
Control Systems Letters	472K	521K	(336K)	(384K)	135K	136K
Conference Distributed Pa	2,219K	2,319K	(1,283K)	(1,455K)	935K	864K
Conferences Total	2,516K	2,548K	(1,502K)	(1,673K)	1,014K	875K

expressed confidence that IEEE approval would be granted.

Fregene asked whether the allocated amount should be increased, given comments about difficulty accessing reserve funds. Beck noted that in the one time this program was enacted the requests were far below the limit. How suggested that the BoG revisit the funding level if the situation arises that more than the amount allocated is requested.

Beck noted that the program would begin at CCTA the following year, and that additional funding could be requested through a future motion if needed. Rudie expressed strong support for the initiative and suggested increased advertising through Women in Control (WiC), noting that earlier equity-focused initiatives had low participation.

Called by: Serrani

Seconded by: Coogan

The motion passed unanimously.

President Beck presented the following motion:

» **Motion:** *To approve funds for continued CSS website management by Conference Catalysts*

Financial Impact: \$31,800 (an increase of \$1,800 from previous)

Endorsed By: Executive Committee

It was noted that Conference Catalysts manages conference websites and that the existing contract expires at the end of 2025. The scope of services was summarized.

How asked whether the increased cost compared to the prior Statement of Work should be expected annually.

Beck explained that a 3–4% increase is typically built in every few years but acknowledged that this increase was approximately 6% and suggested it was to account for higher cost-of-living adjustments.

Called by: Serrani

Seconded by: How

Abstention: Palycarpou

The motion passed.

VP for Finance Activities: Jonathan How

VP for Finance Activities, Jonathan How, presented the following motion:

» **Motion:** *to reduce the cost for IEEE CSS Membership from \$25 to \$20 (student rate reduced from \$5 to \$1)*

Financial Impact: Depends on by how much we reduce the cost

Endorsed By: Executive Committee

The motivation for this proposal stemmed from a comparison of IEEE, CIS, and RAS membership costs during renewal. A table showing the proposed rates was presented:

		Student Rate		
		1	5	10
Member Rate	5	32.3	37.8	44.8
	10	63.2	68.8	75.7
	15	94.1	99.7	106.6
	20	125.0	130.6	137.5
	25	156.0	161.5	168.4

How noted that while the Society cannot provide significant financial benefits, it can reduce the cost of membership. Once membership fees are lowered, it may be difficult to increase them again. It was suggested that the rates could be reevaluated the following year.

Pena noted that documentation she received referenced a reduction from US\$25 to US\$15. Ferrara clarified that an updated version of the proposal was distributed following the ExCom meeting to reflect subsequent updates. How clarified that the intent was to aim for US\$15 over a two-year period, with the possibility of revisiting the issue if the first-year proceeds smoothly.

Fregene asked whether the reduction was intended to compete with RAS membership rates. How responded that the decision was driven by the Society's surplus and not by competitive intent. He reiterated that the Society cannot provide large incentives, but reducing membership cost is a tangible benefit.

Pena asked whether the IEEE membership cost of US\$239 varied by country. How responded that the figure was taken directly from the IEEE website but confirmed that membership costs vary by location.

Annaswany reminded the Board of previous limited-time promotions and asked whether there was data on the proportion of non-CSS or non-IEEE members attending confer-

ences. How explained that authors may become members during paper submission, but that there is no broad outreach effort targeting students at random times. He emphasized focusing on member appreciation initiatives, potentially during CSS Days. Beck expressed support for revisiting this idea in May.

Fregene proposed an amended membership fee of US\$18. How noted that this change would affect the accuracy of the financial projections shown.

Fregene formally proposed the amendment to change the membership to US\$18.

Seconded by: Fekih

All members voted in favor of the amendment.

The amended motion was restated:

» **Motion:** *To reduce the cost for IEEE CSS Membership from \$25 to \$18 (student rate reduced from \$5 to \$1)*

Financial Impact: approximately \$50K based on current membership distribution

Endorsed By: Executive Committee

Seconded by: Ferrara

The motion passed unanimously.

VP for Publication Activities: Andrea Serrani

VP for Publication Activities, Andrea Serrani, presented the following motions

Serrani reviewed the two types of contracts used by the Society: contracts through a university or company, and independent contractor agreements.

» **Motion:** *To approve a new contract for the CSM editorial assistant for 2026*

Financial Impact: \$70,297 for one year

Endorsed By: Executive Committee

The Editor-in-Chief, Annaswany, requested renewal of the contract with the current EA. The editorial services provided for CSM differ from other publications and involve extensive editorial reviews of the entire publication.

Called by: Ferrara

Seconded by: Coogan

The motion passed unanimously.

» **Motion:** *To approve a new contract for the TCNS editorial assistant for 2026*

Financial Impact: \$28,840 for one year

Endorsed By: Executive Committee

The Editor-in-Chief requested retention of the same Editorial Assistant.

Called by: Fekih

Seconded by: How

The motion passed unanimously.

» **Motion:** *To approve a new contract for the TCST editorial assistant for 2026*

Financial Impact: \$56,650 for one year

Endorsed By: Executive Committee

The publication transitioned from a university-based agreement to an independent contractor. Following termination of the prior contract from the university, the current EA was hired in August 2025.

Editor-in-Chief, Ilya Kolmanovsky, expressed strong satisfaction with the EA's performance and requested renewal with a 3% adjustment for 2026.

Rudie proposed that the name of the contact person from the University be removed from any publicly posted documents. The motion was updated accordingly. Fregene brought up the point of whether names should be removed from publicly posted minutes. Beck agreed and stated that minutes would be reviewed to remove names prior to posting. Ferrara emphasized the importance of BoG review of minutes before publication in CSM.

Called by: Fekih

Seconded by: Ferrara

The motion passed unanimously.

» **Motion:** *To approve a new contract for a combined editorial assistant for TAC and OJ-CSYS (at 100% FTE), and to terminate the current editorial service contract for OJ-CSYS*

Financial Impact: \$117,429 for one year

Endorsed By: Executive Committee

Serrani explained that combining editorial services was logical. The EiC of OJ-CSYS, Sonia Martínez, and the upcoming EiC of TAC, Miroslav Krstić, have agreed to share the Editorial

Assistant, with a percentage of shared effort to be discussed, and a formula was developed to allocate costs across publications:

Allocation Formula:

$$E_{ojcs} = 0.1 + \frac{N_{ojcs}}{N_{ojcs} + N_{tac}} * 0.8$$

$$E_{tac} = 0.1 + \frac{N_{tac}}{N_{ojcs} + N_{tac}} * 0.8$$

N_{ojcs}: number of submissions in 2025 for OJ-CSYS

N_{tac}: number of submissions in 2025 for TAC

With the current figures, the allocation of effort is **12% OJ-CSYS and 88% TAC**

A three-year contract was initially proposed to IEEE; however, the total exceeded the threshold for internally provided IEEE services. How presented the proposal to the IEEE FinCom and addressed questions. IEEE approved a one-year contract through 2026, with alternatives to be discussed for 2027.

How summarized FinCom discussions and noted that while initial feedback was negative, a one-year solution was reached. How and Serrani will work with Krstić on next steps.

Serrani added that a complete re-evaluation is required and that this presents an opportunity to establish a more uniform process.

Palycarpou asked whether IEEE is pushing societies toward internal editorial assistants and how this might affect CSS.

Serrani explained that IEEE's rates are highly competitive and that CSS uniquely uses PaperPlaza, making a return to ScholarOne infeasible

Called by: Fekih

Seconded by: How

The motion passed unanimously.

» **Motion:** *To approve the page budgets of the CSS journals for 2026*

Financial Impact: \$25,828 for one year
Endorsed By: Executive Committee

Questions were raised regarding exceeding page limits and backlog management. Serrani explained that exceeding page limits incurs

additional costs and that any increases would require an ExCom motion.

Krstić noted a 15% increase in TAC submissions over the past year. Serrani stated that the current page budget was discussed at ACC and would be reevaluated after one year.

Called by: Fekih

Seconded by: Di Bernardo

The motion passed unanimously.

» **Motion:** *To approve a roll-over of the 39 unused APC waivers for OJ-CSYS expiring in 2025 over 2026 and 2027*

Financial Impact: \$3,757 over the next two years (estimated)

Endorsed By: Executive Committee

An overview of the waivers was provided. A chart illustrating waiver usage was presented:

Serrani believes the cost is worth the promotion of the journal.

Egerstedt expressed strong support for the waivers and asked whether there was insight into why a roll-over was necessary. Serrani suggested that OJ-CSYS is still a relatively young journal and that a combination of limited awareness and publication timing may be contributing factors.

Fregene commented that the Society appears to be managing many editorial initiatives simultaneously. Serrani clarified that the structure of these initiatives is governed by IEEE and is not fully under Society control.

Fregene asked whether the same issue would arise again once the waivers expire in two years.

Serrani responded that the journal status would be reevaluated at that time, with consideration given to alternative support mechanisms, including potential assistance for authors from less financially stable countries.

Beck noted that IEEE has previously suggested similar approaches and that the Society has implemented them in the past. She stated that allowing continued support is a natural step.

Serrani noted the figure that was presented was just a rough estimate. Judging from the trend, it indicates increased usage in 2026

Called by: Fekih

Seconded by: Di Bernardo

The motion passed unanimously.

» **Motion:** *To renew the co-sponsorship of two publications managed by the Computer Society*

» IEEE Transactions on Computational Biology and Bioinformatics (TCBB)

» IEEE Transactions on Big Data (TBD)

Financial Impact: none

Endorsed By: Executive Committee

It was noted that these sponsorships carry no financial obligations for CSS.

Dong asked about the duration of the sponsorship. Beck confirmed that the renewal period is five years

Called by: Fekih

Seconded by: Di Bernardo

The motion passed unanimously.

VP for Conference Activities: Lihua Xie

VP for Conference Activities, Lihua Xie, presented the following motions

» **Motion:** *To approve Antonella Ferrara as Program Chair for CCTA 2028*

Financial Impact: none

Endorsed By: Executive Committee

Ferrara was asked to step out of the room during discussion

Year	APC of single paper	New Cost	Cost Difference = New Cost – Previously approved cost
2026	\$2,160	24W @ 50% = \$25,920	\$1,920
2027	\$2,245	15 W @ 50% = \$16,837.5	\$1,837.5
Total		\$42,757.5	\$3,757.5

An overview of her qualifications was provided. No questions or comments were raised by the Board.

Called by: Fregene

Seconded by: Beck

The motion passed unanimously.

Ferrara entered back into the meeting

» **Motion:** *To approve Murat Arcak as Program Chair for CDC 2029.*

Financial Impact: none

Endorsed By: Executive Committee

Xie gave an overview of Arcak's qualifications. No comments were raised.

Called by: Palycarpou

Seconded by: Johansson

The motion passed unanimously.

» **Motion:** *To approve Yanlong Zhao as Program Chair for CDC 2030*

Financial Impact: none

Endorsed By: Executive Committee

Xie gave an overview of Zhao's qualifications. No comments were raised.

Called by: Isaksson

Seconded by: Dong

The motion passed unanimously.

VP for Technical Activities: Mario di Bernardo

VP for Technical Activities, Mario di Bernardo, presented the following motions

» **Motion:** *To approve the change of name of the TC on Networks and Communication Systems to TC on Network Systems*

Financial Impact: none

Endorsed By: Executive Committee

No comments were raised.

Called by: Johansson

Seconded by: Ferrara

The motion passed unanimously.

» **Motion:** *To approve Emanuele Garone as TC Chair and Takeshi Hatanaka as TC Vice-Chair of the new TC in Control and Optimization in Food and Agriculture*

Financial Impact: none

Endorsed By: Executive Committee

No comments were raised.

Called by: Xie

Seconded by: Johansson

The motion passed unanimously.

» **Motion:** *To approve CSS Days 2026 budget for streaming and video recording*

Financial Impact: \$12,000

Endorsed By: Executive Committee

How clarified that US\$12,000 was allocated for services, with an additional US\$3,000 reserved for unexpected expenditures. Di Bernardo confirmed this understanding. Egerstedt suggested that the motion be revised to approve only the US\$12,000, noting that additional expenditures up to US\$3,000 could be approved by the President.

The wording in parentheses was removed to reflect this clarification.

Called by: Fregene

Seconded by: Johansson

The motion passed unanimously.

» **Motion:** *To approve Ian Petersen and Jing Sun as members together with current VP TA of the CSS-OTC Award 2026*

Financial Impact: none

Endorsed By: Executive Committee

It was noted that this award will be presented for the first time next year and will evaluate the work of Technical Committees as a whole.

Called by: Ferrara

Seconded by: Beck

The motion passed unanimously.

» **Motion:** *To approve Prof. Yoshio Ebihara as the new CSS Liaison Representative to SICE*

210230-Control Systems Summary Report For Period Ending Oct-25	Total Revenue			Total Expense			Net		
	Annual Budget	YTD Budget	YTD Actual	Annual Budget	YTD Budget	YTD Actual	Annual Budget	YTD Budget	YTD Actual
All Products									
Magazines	169K	143K	149K	429K	330K	366K	259K	188K	217K
Journals Transactions Rev	3,296K	2,763K	2,758K	2,264K	1,880K	2,019K	1,032K	883K	739K
Letters	566K	472K	521K	408K	336K	384K	158K	135K	136K
Periodical Packaged Prodt	1K	1K	1K	6K	5K	5K	5K	4K	4K
Conference Events	1,241K	297K	229K	1,007K	219K	218K	234K	78K	11K
Conference Proceeding	2,662K	2,219K	2,319K	1,540K	1,283K	1,455K	1,122K	935K	864K
Society Membership	135K	135K	139K	4K	4K	4K	131K	131K	135K
Total Product	8,070K	6,029K	6,116K	5,658K	4,059K	4,451K	2,412K	1,970K	1,665K
Cost Centers									
ExCom	0K	0K	0K	0K	0K	128K	0K	0K	128K
Membership Committee	0K	0K	0K	2K	0K	2K	2K	0K	2K
Awards Committee	0K	0K	0K	8K	1K	1K	8K	1K	1K
Executive Office	0K	0K	0K	0K	0K	0K	0K	0K	0K
Societies Operations	0K	0K	0K	1,789K	1,331K	1,146K	1,789K	1,331K	1,146K
Meetings /Conference	18K	15K	70K	243K	190K	236K	225K	175K	166K
Total Cost Center	18K	15K	70K	2,042K	1,522K	1,512K	2,024K	1,507K	1,443K
Total From Operations	8,088K	6,044K	6,186K	7,699K	5,581K	5,964K	388K	463K	222K
Grand Total	8,088K	6,044K	6,186K	7,699K	5,581K	5,964K	388K	463K	222K

Financial Impact: none
 Endorsed By: Executive Committee

No comments were raised.
 Called by: Xie
 Seconded by: Johansson
The motion passed unanimously

Meeting breaks at 3:04 p.m.
 Meeting returned at 3:30 p.m.

Project Name	2026 Budget (\$K)
CSS Child-care and Disability Assistance	20
CSS Outreach	160
CSS Student Travel Awards	165
NextCom	30
Women in Control Lunch	20
Total	395

INFORMATIONAL ITEMS

IEEE CSS VP for Financial Activities

Jonathan P. How

Financial summary

How provided a deeper dive into the numbers discussed earlier:

It was noted that blue figures indicate positive performance, although values are lower than historical norms. Revenue is based on download activity from IEEE *Xplore*.

Questions were raised regarding whether revenue levels remain consistent with prior years. It was stated that the lower figures are not cause for immediate concern.

Editorial Services Assistant (ESA)

How spoke about the journal that had outstanding invoices totaling approximately US\$120,000, representing two years of unpaid services. A decision was made to pay the two outstanding invoices in the current year, resulting in a total of three ESA invoices paid in 2025.

2025 inputs from IEEE

It was reported that prior IEEE budgeting rules, including the 1%/3%/50% rules, have been eliminated.

IEEE has also eliminated initiative-based spending, with all expenses now routed through the operational budget. IEEE identified significant funds held in Society reserve accounts and expressed interest in seeing those funds actively utilized.

What Were Known as Society Initiatives

How noted that the 2026 budget has not been communicated very clearly. He continued saying it is more of an indicator of how we think we will spend that money. Emphasis was placed on forecasting rather than treating these figures as fixed allocations.

How advised that, when reviewing the budget, we should be increasingly mindful of how funds are spent as we approach the margin. Operating near zero is considered acceptable; however, if expenses exceed revenue and move into a deficit, reevaluation will be necessary.

Egerstedt commented that the Society has intentionally moved away from generating profit from conferences, noting that the previous 20% margin target is no longer in place. He questioned whether the observed reduction reflects this fact. How stated that he is not concerned, emphasizing that breaking even is a positive outcome.

How noted that CCTA reflected a loss of approximately US\$50,000, which he considered acceptable.

Beck asked whether the Society should reevaluate its approach if conferences consistently incur into losses, noting that some expenses have shifted into the operational budget. How responded that a US\$50,000 variance is close to breakeven and does not represent a significant loss. He emphasized that the Society is acting appropriately as long as conferences break even or remain close. He further stated that while revenue streams cannot be directly controlled, expenses can be managed.

Coogan asked whether IEEE provides click data. How clarified that

IEEE provides download counts, not click data.

Security Awareness

How presented an example of phishing. Board members were urged to remain vigilant and cautious when responding to suspicious emails.

IEEE CSS VP for Conference Activities

Lihua Xie

Xie gave a summary of recent and upcoming CSS-sponsored conferences

For CCTA 2029, a proposal was received that requires revisions. The proposal will be reevaluated once revisions are submitted.

CSS Technically Co-Sponsored Conferences 2025/26

Approved to date (21 for 2025 and 9 for 2026 so far)

- » ICRCA 25: 2025 the 9th International Conference on Robotics, Control and Automation, Shanghai, China, 10–12 Jan, without PA
- » ANZCC25: The 2025 Australian and New Zealand Control Conference, Gold Coast, Australia, 30–31 Jan., with PA
- » EECI 25: EECI International Graduate School on Control, International, 27 Jan–4 July, without PA
- » CCDC 2025: The 37th Chinese Control and Decision Conference, Xiamen, China, May 16–19, with PA
- » WiOpt 25: International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks, Linköping, Sweden, May 26–29, with PA

Conference	Location	Dates	General Chair	Program Chair(s)
ACC 2025	Denver, CO	Jul 8–10	Carolyn Beck	Jonathan How
CCTA 2025	San Diego, CA	Aug 18–20/25–27	Robert Bitmead	Colin Jones
CDC 2025	Rio de Janeiro, Brazil	Dec 10–12	Yannis Paschalidis, João Carlos Basilio	N. Li, C. Hadjicostis
ACC 2026	New Orleans, LA	May 27–29	Randy Beard	Eric Frew
CCTA 2026	Vancouver, Canada	Aug 12–14	Ryozo Nagazume	Santosh Devisia
CDC 2026	Honolulu, Hawaii	Dec 15–18	Jorge Cortes	Lacra Pavel
ACC 2027	Philadelphia, PA	July 6–9	Rajesh Rajamani	Junmin Wang
CCTA 2027	Japan	Aug 23–25	M. Nagahara, D. Quevedo	Kai Cai
CDC 2027	Lisbon, Portugal	Dec 14–17	João Hespanha	Paulo Tabuada
ACC 2028	San Francisco	May 31–June 2	Zongli Lin	Amir Aghdam
CCTA 2028	Berlin, Germany	TBD	Frank Allgöwer Steffi Knorn	[Antonella Ferrara]
CDC 2028	Sydney, Australia	December 12–15	Girish Nair	Henrik Sandberg
CDC 2029	TBD	TBD	Geir Dullerud	[Murat Arcak]
CDC 2030	Hong Kong, China	TBD	Lei Guo	[Yanlong Zhao]

- » ICUAS 25: 2025 International Conference on Unmanned Aircraft Systems, UNC Charlotte, Charlotte, NC, USA, 14–17 May, with PA
- » MED 25: 33rd Mediterranean Conference on Control and Automation, Tangier, Morocco, Jun 10–15, with PA
- » CPDE25: Control of Systems Governed by Partial Differential Equations, Beijing, China, 18–20 June, without PA
- » ECC 2025: 23rd ECC, Thessaloniki, Greece, 24–27 June, 2026, with PA
- » ICCA 2025: The 19th IEEE International Conference on Control & Automation, Tallinn, Estonia, June 30–July 3, with PA
- » ECC 2025: ECC, Thessaloniki, Greece, June 24–27, with PA
- » qCCL 25: 2025 IEEE International Conference on Quantum Control, Computing and Learning, Hong Kong, Jun 25–28, without PA
- » CCC 2025: The 44th Chinese Control Conference, Chongqing, China, 28–31 July 2024, with PA

- » CoDIT25: 11th International Conference on Control, Decision and Information Technologies, Split, Croatia, 15–18 July, with PA
- » MMAR 2025: 29th International Conference on Methods and Models in Automation and Robotics 2024, Miedzyzdroje, Poland, 26–29 Aug. 2025, with PA
- » ICAT 25: 30th International Conference on Information, Communication and Automation Technologies, Sarajevo, Bosnia and Herzegovina, 2–4 Sept., with PA
- » SysTol 25: 6th Conference on Control and Fault-Tolerant Systems, Ayia Napa, Cyprus, Jun 6–8, with PA
- » ICSTCC 2025: 29th International Conference on System Theory, Control and Computing, Cluj-Napoca, Romania, Oct. 9–11, with PA
- » IMC-SSGP 25: 2025 IEEE International Conference on Smart Systems and Green Process, Hammamet, Tunisia, Oct. 30–Nov. 2, without PA

- » ICSC 25: IEEE 13th International Conference on Systems and Control, Marrakech, Morocco, 22–24 Oct., with PA
- » ICC 25: 2025 Eleventh Indian Control Conference, Bengaluru, India, Dec. 18–20, with PA
- » CCDC 2026: The 38th Chinese Control and Decision Conference, Nanjing, China, May 15–18, 2026, with PA
- » ICCA 2026: The 20th IEEE International Conference on Control & Automation, Almaty, Kazakhstan, 16–19 June, 2026, with PA
- » ICUAS 26: 2026 International Conference on Unmanned Aircraft Systems, June 15–18, 2026, Corfu, Greece, with PA
- » WODES 2026: 18th Workshop on Discrete Event Systems, Eindhoven, The Netherlands, June 8–10, 2026, with PA
- » PCC 26: 22nd Polish Control Conference, Poznań, Poland, July 1–3, 2026, without PA
- » qCCL 26: 2026 IEEE International Conference on Quantum Control, Computing and Learning, July 1–3, 2026, without PA
- » VSS 2026: 18th International Workshop on Variable Structure Systems, Exeter, UK, July 27–30, 2026, with PA
- » AdCONIP 26: International Symposium on Advanced Control of Industrial Processes, Auckland, New Zealand, Nov. 29–Dec. 1, 2026, with PA
- » ICSC 26: IEEE 2026 14th International Conference on Systems and Control, Dec. 16–18, 2026, Porto, Portugal, with PA

Beck noted that conferences are evaluated not only on quality but also on alignment with the Society scope. It was stated that approximately 20 conferences are reviewed each year.

Jungers asked whether the US\$500 conference fee could be waived. Beck responded that the fee cannot be waived, as it is imposed by IEEE. IEEE charges APPROXIMATELY US\$1,500, AND IF CSS PAYS the Publication Agreement fee,

an additional charge of approximately US\$22–US\$23 per paper is applied by IEEE *Xplore*. This applies only to non-IEEE conferences.

IEEE CSS VP for Member Activities

Necmiye Ozay

New Distinguished Lecturers

New Distinguished Lecturers were announced as part of a cost-shared program that enables local chapters to engage with distinguished members of the control community.

The selection committee is chaired by Sanjay Lall and collects nominations from chapters worldwide.

Distinguished Lecturers (DLs)—3-year term

- » Naira Hovakimyan (Univ. of Illinois)
- » Tara Javidi (UC San Diego)
- » Na Li (Harvard)

Distinguished Industry Lecturers (DILs)—2-year term

- » Kingsley Fregene (Lockheed Martin)
- » Sampriya Bhattacharyya (Navierboat)

See the CSS website for the full list of DLs/DILs.

Ozay urges those involved in local chapters to send names/advertise for future DLs/DILs.

Outstanding Chapter Award

Control, Robotics, Cybernetics Joint Chapter of IEEE Vancouver Section (CH07088)

Di Bernardo asked how many chapters currently exist. Ozay responded that there are approximately 80 chapters.

It was noted that most newly approved chapters are student chapters. A discussion was suggested with IEEE regarding chapter review and renewal processes.

Annaswany stated that correlating chapter activity across organizations is challenging, particularly for joint chapters where CSS participation may be limited. Di Bernardo suggested exploring additional engagement strategies.

CSS Member Statistics

Fregene requested membership data broken down by region. Ozay stated that she would attempt to obtain this information.

IEEE CSS VP for Technical Activities

Mario di Bernardo

CSS Days 2026 Program Structure and Possible Themes

The proposed structure and possible themes for CSS Days 2026 were presented.

It was stated that the goal is to avoid overly packed schedules and instead distribute activities throughout the week.

An example schedule and structure was shown:

Monday, October 19, 2026

Tuesday–Friday, October 20–23, 2026

Di Bernardo urged board members interested in participating to reach out with suggestions or offers to assist.

Di Bernardo showed data from the previous CSS Days event was reviewed, including viewership of recorded content. He noted that a call for proposals is planned, with feedback and results to be reevaluated in March

CSS Days 2026 Task Force

As with previous CSS Days events, a Task Force chaired by the Vice President for Technical Activities will be established to organize the program and activities for October 2026.

Task force will consist of:

VP TA: Mario di Bernardo (Chair)

VP DODA: Karl Johansson

VP Membership: Necmiye Ozay

Future DODA Task Force: Philip Pare

BoG member from industry: Pratik Vernekar

Industry Activities Chair on the Member Activities Board: Lidia Auret
Future WiC Co-Chairs: Elisa Franco, Patricia Pena

Mauro Salazar (Netherlands, TC on Smart Cities), Emiliano Dell'Anese (Boston, TC Networks)

Giovanni Russo (Senior Editor TCNS) Board involvement was again encouraged.

Egerstedt commended the effort, noting that the work reflects careful planning and strong attention to detail.

CSS Website Structural Update and Redesign

Conference Catalysts was asked to provide quotes, and highlights of their proposal were presented FFto the Board of Governors.

Board members were invited to identify website features they consider important and to communicate those priorities accordingly.

CSS Marketing Management


Since the previous meeting, comparisons were conducted between CSS and both the RAS and SMC. It was noted that RAS and SMC maintain highly active LinkedIn accounts, whereas CSS social media platforms currently exhibit little to no activity. This lack of engagement was identified as a significant issue. It was emphasized that numerous CSS activities are underway and that improved communication channels are needed to promote them effectively. Marketing efforts were noted to be closely linked with the restructuring of the EiC role.

The proposed budget includes support for a social media manager and associated services.

Fregene asked about the earliest anticipated timeline for implementation. Di Bernardo responded that implementation would likely occur the following year after ACC.

Fregene emphasized the goal of harmonizing social media postings across platforms.

Pena shared her experience working on social media for CDC, explaining that initial content was created from scratch until it was discovered that an existing CDC account already existed. Utilizing the existing account allowed them to begin with approximately 500 followers.



Session	Organizer	Description
Welcome & President's Message	CSS President & VP TA	Opening remarks, CSS updates, and strategic direction. The VP TA will also introduce the structure and highlights of CSS Week 2026.
President's Forum	CSS President	Dialogue with past and current CSS leaders on future directions.
CSS Publications Showcase	VP Publications and CSS Journals EiCs Task force member: Giovanni Russo	Each Editor presents the editorial vision, policy, and scope of their journal.
CSS History Session	History Committee	Curated reflections on the evolution of control and the Society.
Distinguished Lecturers Highlights	VP Membership & DL Program Chair	Lightning talks or a round table with CSS Distinguished Lecturers.

How commented that Conference Catalysts may not be the most natural choice for social media management, and asked about the outcomes of requests for alternative proposals. Di Bernardo explained that Conference Catalysts has extensive familiarity with IEEE structures but acknowledged that exploring alternative providers remains an option.

It was noted that other website vendors proposed using entirely different

content management systems, which was deemed unacceptable. These alternatives would also have required rebuilding the site from scratch.

Pena added that for CDC, a student was engaged to manage social media, with funding from DODA to support conference attendance. Frege expressed concern about relying on students for a role of this scale, and the preference was reiterated for a

specialized social media management company.

2025 TC-OSPAs Awardees and TC-OSPA Procedure

Di Bernardo reported that 13 of 22 Technical Committees received TC-OSPAs in 2025, compared to 8 committees in 2024. It was expressed that participation is expected to increase further in the following year.

Day	Focus	Lead	Example Activities
Tuesday	Diversity, Outreach & Students Day	VP DODA, WiC, NextCom, DODA, Student Activities Task force members: Philip Pare, Patricia Pena, Elisa Franco	DEI events, outreach stories, student competitions, mentoring sessions.
Wednesday	Industry & Innovation Day	VP Membership & Industry Committee Task force members: Pratik Vernekar, Lidia Auret	Industry Committee session, panels, and success stories on industrial applications and collaboration. Industrial Distinguished Lecturers
Thursday	Chapters Day	VP Membership & Chapter Activities Chair	Local events and webinars around the world showcasing regional initiatives.
Friday	Technical Committees Day	VP TA & TC Chairs Task Force members: Mauro Salazar, Emiliano Dell'Anese	TC showcases, workshops, webinars, and joint sessions between TCs.

TC Chair Terms (Current Bylaws) and TC Committees Reorganization

Di Bernardo explained that under the current process, outgoing TC Chairs notify the VP of their preferred successor, without broader committee involvement or awareness.

He would like to hold an induction day for new TC chairs.

Proposal for a reviewed TC Leadership Structure

Current bylaws:

Article IV—Technical Committees and Working Groups

Section 3. Technical Committee Chairs.

Technical Committee Chairs shall be appointed by the President-Elect when an incumbent is in office, otherwise by the President, subject to Board of Governors approval, **for one-year terms**. These Chairs shall be charged with the guidance of the regular activities of the Technical Committees. They shall be responsible for making yearly reports to the Board of Governors on the activities of their committees and the Working Groups monitored by their committees.

- » Current procedure for appointment of TC chairs vary from TC to TC.
- » Some also have Vice-Chairs and other officers.
- » Most TC Chairs believe their terms is 3 year long.
- » Some TCs are renewing their own terms over and over again

Proposed change to the bylaws to be initiated in 2026

Section 3. Technical Committee Chairs.

Technical Committee Chairs shall be

appointed by the President-Elect when an incumbent is in office, otherwise by the President, subject to Board of Governors approval, **for three-year terms. No Technical Committee Chair shall serve for more than two three-year terms.** These Chairs shall be charged with the guidance of the regular activities of the Technical Committees. They shall be responsible for making yearly reports to the Board of Governors on the activities of their committees and the Working Groups monitored by their committees.

Also, at the next TAB meeting at CDC, we will discuss with all TC Chairs some proposed changes to the guidelines for TC Chair appointments and a possible restructuring of the TC management committees.

Proposal for a reviewed TC Leadership Structure:

OFFICER STRUCTURE

Each TC has three officers: Chair, Chair-Elect, and Secretary. Each serves a 3-year term. All terms end together at the annual TC meeting three years after the Chair's appointment, or between two and three years if the Chair was previously renewed off the annual meeting cycle.

CHAIR RESPONSIBILITIES

The Chair, supported by the other officers, is responsible for:

- a) ensuring timely execution of TC activities to achieve the TC's purpose;
- b) ensuring TC operations comply with CSS Bylaws; and

- c) serving as the primary communication channel between the TC and the CSS Board of Governors.

The Chair presides over Committee meetings and voting, manages all Society interactions, and reports TC activities to the Board of Governors through the Vice President for Technical Activities at least once annually and at least 60 days prior to CDC.

The Chair is a member of the CSS Technical Activities Board (TAB) and attends at least the Annual TAB Meeting, typically held during CDC.

CHAIR APPOINTMENT

The Chair is normally appointed through transition from the Chair-Elect position at the end of their three-year term, or upon resignation/removal of the previous Chair. An out-of-sequence election may be held if the Chair-Elect is unable to assume the Chair role.

CHAIR REMOVAL FOR NON-PERFORMANCE

The Society President may remove a Chair who fails to fulfill their responsibilities, including but not limited to failure to submit the annual TC report, failure to respond to requests from the Vice President for Technical Activities or other Executive Committee officers, or failure to attend required meetings. Before removal, the President shall provide written notice to the Chair specifying the deficiencies and allowing a reasonable period (at least 30 days) for the Chair to remedy the situation. Upon removal of the Chair, the Chair-Elect immediately assumes the role of Chair,

and an election for a new Chair-Elect shall be held within 30 days.

CHAIR DUTIES

The Chair oversees all TC activities, with assistance from officers and members as needed, and is responsible for:

- » soliciting TC nominations for awards and for editorial, Distinguished Lecturer, and CASS representative roles;
- » ensuring preparation of the annual TC report;
- » preparing materials for and participating in periodic TC Reviews;
- » maintaining the TC vision and mission.

For editorial appointments, the Chair monitors editorial board formation schedules and solicits member nominations in a timely manner.

CHAIR-ELECT

The Chair-Elect is elected by ballot by TC members and must have been a regular TC member for the previous two years. The ballot is held at the annual TC meeting or by electronic means if the position becomes vacant during the term (e.g., due to resignation or promotion to Chair).

Nomination process: Candidates' CVs are circulated to all members and to the VP TA, and TC members have at least 15 days to discuss nominations before the ballot is finalized. Self-nominations are permitted. The TC encourages broad participation from all members in both the nomination and election processes. When nominating candidates for Chair-Elect, TC members should actively seek nominees from diverse backgrounds including gender, career stages, geographic regions, and institutional affiliations to ensure the TC leadership reflects the breadth and diversity of the technical community it serves.

The Vice President for Technical Activities may solicit additional nominations from TC members if the initial nominee pool does not adequately reflect the diversity of the TC membership and the broader technical community.

Restrictions: Past TC Chairs cannot be nominated for Chair-Elect if they have: (a) already served as TC Chair more than once; or (b) ended their Chair term less than three years ago.

CHAIR-ELECT RESPONSIBILITIES

The Chair-Elect assists the Chair in all activities, participates in TC review meetings, audits member activities, notifies members of activity lapses, and is responsible for:

- » soliciting nominations for CSS journal best paper awards;
- » quality assessment and maintaining/raising quality of CSS-sponsored conferences in the TC field;
- » other TC activities.
- » Both Chair and Chair-Elect motivate TC members to promote TC-sponsored special sessions and workshops at CDC, CCTA, and other CSS events related to TC activities.

SECRETARY

The Secretary is elected by ballot by TC members and must have been a regular TC member for the previous three years. The ballot is held at the annual TC meeting or by electronic means if the position becomes vacant during the term.

Nomination process: Candidates' CVs are circulated, and TC members have at least 15 days to discuss nominations before the ballot is finalized. Self-nominations are permitted. The TC encourages broad participation from all members in both the nomination and election processes.

SECRETARY RESPONSIBILITIES

The Secretary:

- » records meeting minutes and ensures prompt posting on the TC website;
- » maintains the Roll of Members including email addresses;
- » maintains mailing lists for TC members and officers;
- » maintains records of member attendance at annual meetings and participation in activities;

- » serves as a focal point for maintaining the TC website and coordinating media/web TC presence;
- » acts as TC Treasurer, ensuring timely expenditure of the annual TC allowance on TC-related activities.

OFFICER REMOVAL

Officers may be removed before term-end by a majority anonymous vote of TC regular members, initiated by a motion signed by at least one-fourth of TC regular members.

ADDITIONAL ROLES

Beyond officer roles, other duties may involve periodic tasks or occasional tasks for which temporary positions or working groups may be formed, reporting to the officers.

IMPLEMENTATION OF REVISED PROCEDURES

These revised procedures shall take effect in 2026. To align all TC officer terms with the new structure, all Technical Committees shall hold elections for Chair-Elect and Secretary in 2026. Current Chairs may continue in their roles until their terms expire naturally or until the newly elected Chair-Elect completes their three-year term and transitions to Chair, whichever occurs first. This transition period ensures uniform implementation of the three-year synchronized term structure across all TCs.

Coogan commented that not every Technical Committee requires a Vice Chair and suggested electing Chairs one year earlier to allow for a one-year overlap period.

Di Bernardo stated that the document would be distributed for review and feedback. Once comments are received, the Society will move forward with implementation of the plan.

TC Involvement in DL Program (with VP Membership)

Discussion was held regarding Technical Committee involvement in the Distinguished Lecturer program, in coordination with the Vice President for Membership Activities.

IEEE CSS VP FOR DIVERSITY, OUTREACH, AND DEVELOPMENT ACTIVITIES

Karl H. Johansson

Overview of 2024 Outreach projects

Budget 2024: US\$120,000

Approved 2024 (US\$122,832)

- » Joint Summer School on “Variable structure and sliding mode control” and “Indo-French seminar on advances in robust nonlinear control for uncertain dynamic systems: theory and applications”, India, \$18,209
- » 2024 ACC Self-Driving Car Student Competition, Canada, \$20,000
- » NextCom, Italy, \$17,870+\$3,953
- » 2024 CDC Cup, Italy, \$1,800
- » Series of Control System Workshops 2024–2025, Universiti Sains Malaysia, \$6,000
- » Supporting Students and Early Career Researchers in Control Systems at the 5th IFAC CPHS 2024, Turkey, \$10,000
- » Promoting IEEE Control Systems Society in Nigerian Universities through Automatic Control Workshops, Nigeria, \$20,000
- » Purdue ICON Student Research Conference, USA, \$4,150
- » Inclusive Diversity Lunch Event at IEEE CDC 2024, Italy, \$2,500
- » The future of control engineering—a female perspective, Lund University, Sweden, \$18,350

Others

- » Empower a Billion Lives III 2025, \$100K (from CSS Fund at the IEEE Foundation) + travel expenses, Aranya Chakraborty CSS Responsible, approved by BoG Dec 2024

For 2025, the budget was raised to US\$160,000.

Budget 2025: US\$160,000

Approved applications 2025

(US\$160,393)

- » Summer school on control of distributed parameter systems and

applications, China, \$20,000

- » Workshop on Future PhD in control, Germany, \$14,500
- » When Robotics Meet Controls: Enhancing Education in Africa through the PARC, \$20,000
- » Bridging the Gap: A Gateway to Control Systems Engineering for Lebanese High School Students, \$13,321
- » Coffee and Control, Podcast, \$560
- » 2025 ACC Self-Driving Car Student Competition, USA, \$20,000
- » Strengthening the Control Community in Colombia, \$19,900
- » Workshop Series on Instrumentation, Control, and Automation for Future Indonesian Scientists and Engineers, \$15,400
- » Workshop on Optimal/Robust Control and State Estimation for Aerospace Applications, India, \$14,000
- » CDC 2025 Soccer Cup, Brazil, \$2,500
- » Diversity and Inclusion Day at CDC 2025, Brazil, \$11,912
- » Control Talk Series: Bridging Academia and Industry, USA, \$2,800
- » Returning Mothers Conference, India, \$5,500

Johansson informed the BoG that they are beginning to fill the 2026 requests.

Budget 2026: US\$160,000

Approved applications 2026

(US\$54,500)

- » Workshop on Applications of Geometry, India, \$10,000
- » Launching IEEE CSS Student Branch Chapter at IIT Patna, India, \$1,000
- » Control Systems for Everyone—Low-Cost Hands-On Training for Pre-University Students, India, \$2,500
- » ICON Student Research Conference, USA, \$6,000
- » ACC 2026 Self-Driving Car Student Competition, USA, \$20,000

- » Workshop and Summer School on Variable Structure Systems and Sliding Mode, UK, \$15,000

Planned applications 2026

- » High school student competition at CDC 2026 in Hawai'i, USA

Johansson emphasized that outreach projects are intended to strengthen automatic control activities—not just a sponsorship.

Funding is intended to support activities such as student engagement rather than Distinguished Lecturer visits.

Johansson provided 2 examples of the outreach grants:

- » Workshops on Instrumentation, Control, and Automation for Future Indonesian Scientists and Engineers
- » When Robotics Meet Controls: Enhancing Education in Africa

CSS Graduate Collaboration Fellowship

Johansson reported that only three recipients were selected for the CSS Graduate Collaboration Fellowship this year. Increased participation and expansion of the program were encouraged.

DODA Task Force

Johansson noted that Silvia has provided significant support to the DODA Task Force. She conducts monthly online meetings to solicit and review outreach proposals.

Women in Control Standing Committee

Johansson showed photos from the Women in Control luncheon that was held on Tuesday July 8, 2025, at the ACC.

PRESIDENT'S UPDATES

Request for Bylaws Changes

Article III Section I has already completed the IEEE approval process.

Article I Sections 3–4 remain under review due to language related to DEI.

It was noted that additional CSS bylaws and constitutions also require revision.

Rudie asked whether the Society could clarify its values even if compliance with IEEE requirements is necessary. Beck responded that careful wording is required to maintain CSS values while meeting IEEE requirements.

Awards Nomination Processes

The primary change is the removal of the requirement for supporting letters. While supporting letters may still be submitted, they are no longer mandatory.

For the Roberto Tempo CDC Best Paper and CCTA Best Paper Awards:

- » Nominations will be solicited via the website submission process, by the respective deadline (Jan 15 for CCTA and May 15 for the Tempo award), but NO support letters will be required.
- » Additionally: after the conference, using the audited reviews, a list of top-rated papers (for example, ranging from 5–25 depending on how many nominations received) will be selected by the respective conference Program Chair based on the Associate Editors' recommendations and reviewers' comments, and sent to the Awards Chair. Only the papers, not the reviews, will be sent forward.
- » The Awards Chair will forward the combined list of nominated and top-rated papers to the respective Award Committee Chair.

Mesbahi asked whether self-nominations are permitted. Beck confirmed that authors may nominate their own papers through PaperPlaza.

For the Journal Best Paper Awards:

- » Nominations can be made as before via PaperPlaza, but NO support letters will be required.
- » Additionally: Journal Senior Editors may flag any paper they deem award level quality (over the course of the corresponding eligibility window) for their journal's BP award.
- » By a specified date (Awards Chair and EiCs will select together and communicate yearly to SEs) each SE will be responsible for reviewing their list of flagged papers, editing as needed, and sending those they

would like to be considered forward to their Editor-in-Chief.

- » The EiC will compile the list of papers to be considered and send it to the Awards Chair.

For the Transition to Practice, Control Technology and Ruberti Awards:

- » Nominations can include supporting letters, but these will NOT be required.

With reference to the paper awards, Maria Prandini raised a concern regarding the closure of PaperPlaza after the conference. Beck acknowledged this limitation and stated that the Society is aware of the issue. She noted that it is possible to request that PaperPlaza be reopened for a fee, and that this option should be considered if necessary, but it is preferable that the PCs and GCs make the short list prior to PaperPlaza closing access to the conference paper information.

Beck then showed the newly elected slate of elected BoG members (Stefano di Cairano, Antonis Papanichristodoulou, Fabio Pasqualetti, Larcra Pavel, Daniel Quevedo, and Carla Seatzu), as well as the 3 new appointed BoG members (Aaron Ames, Kenji Fujimoto and Mariela Verrada).

Beck informed the group that one elected BoG member had to step-down with one year left in their term. In this case, they needed to appoint a BoG member for 1 year—this appointed member is Fumin Zhang.

The outgoing Board of Governors and Executive Committee members were formally recognized.

President Beck expressed appreciation for their service and presented certificates in acknowledgment of their contributions.

Other Business

Beck informed the BoG that they would be receiving a gift in recognition of their work, which would be sent to their respective addresses.

She also provided information for the Next BoG meeting, which will be held in conjunction with ACC 2026 in New Orleans, Louisiana, on May 26th, 2026.

Nair raised a question regarding MTNS 2028, noting that it was

originally scheduled for the week prior to CDC but was later withdrawn. He asked whether there was coordination between CSS and MTNS during planning. Beck referred back to conference date information presented earlier and stated that known dates are tracked in advance and can be communicated as needed. She invited members to share additional information during the planning process.

Fregene raised concern regarding penalties incurred for unfilled hotel room blocks at CCTA and asked whether alternative booking strategies are being considered. Beck explained that negotiating for fewer rooms often results in higher rates, and that if attendance exceeds the reserved block, attendees may lose access to the conference rate. She suggested alternative strategies could be considered, such as reducing room block requirements in exchange for higher food and beverage commitments.

Fekih suggested that attendance challenges for the upcoming ACC may be related to the IFAC World Congress years. Beck stated that this was not the case for IFAC and ACC three years prior and that this year's timing for ACC does not overlap closely with the IFAC World Congress in Busan.

Fregene returned to the topic of OJ-CSYS and asked about adopting a tiered payment structure for open access journals. Beck responded that IEEE maintains a three-tier country classification system. Fregene noted that approximately 60% of respondents indicated that publication costs were too high. Beck stated that the tiered structure is based on World Bank data as per IEEE policies.

Beck motioned to adjourn the meeting.

Seconded by: Ferrara

The motion passed unanimously.

The meeting was adjourned at 5:20 p.m.

To be included in the conference calendar for the IEEE Control Systems Society (CSS), send announcements to:

Syed Zeeshan Rizvi
rizvisz@ieee.org

- ▲ Indicates CSS-sponsored conference
- Indicates CSS-cosponsored conference

The complete and current list of CSS-sponsored and cosponsored conferences is available on the CSS website, <http://www.ieeecss.org>. Please visit conference websites for the most up-to-date information.

» 2026

- **THE 24TH INTERNATIONAL SYMPOSIUM ON MODELING AND OPTIMIZATION IN MOBILE, AD HOC, AND WIRELESS NETWORKS**
3–6 June
Columbus, OH, USA
General Chairs: Jia (Kevin) Liu, Yin Sun, and Elif Uysal
Program Chairs: Falko Dressler, I-Hong Hou, and Bo Ji
<https://wiopt2026.github.io/>
- **18TH WORKSHOP ON DISCRETE EVENT SYSTEMS (WODES) 2026**
8–10 June
Eindhoven, The Netherlands
General Chairs: Michel Reniers and Martin Fabian
Program Chairs: Cristian Mahulea, Patricia Pena, and Martijn Goorden
<https://wodes2026.events.tue.nl/>
- **2026 INTERNATIONAL CONFERENCE ON UNMANNED AIRCRAFT SYSTEMS (ICUAS)**
15–18 June
Corfu, Greece
General Chairs: Nikos Vitzilaios and Giuseppe Loianno
Program Chairs: Marco Tognon, Salua Hamaza, and Nitin Sanket
https://uasconferences.com/2025_icuas/
- **22ND POLISH CONTROL CONFERENCE (PCC) 2026**
1–3 July
Poznań, Poland
General Chair: Maciej Michatek
Program Chair: Dariusz Pazderski
<https://pcc2026.put.poznan.pl/>
- **18TH INTERNATIONAL WORKSHOP ON VARIABLE STRUCTURE SYSTEMS (VSS) 2026**
27–30 July
Exeter, Devon, UK
General Chairs: Christopher Edwards and Antonella Ferrara
Program Chair: Prathyush Menon
<https://engineering.exeter.ac.uk/events/variablestructuresystems/>
- **IEEE CONFERENCE ON CONTROL TECHNOLOGY AND APPLICATIONS (CCTA) 2026**
Vancouver, Canada
12–14 August
General Chair: Ryozyo Nagamune
Program Chair: Santosh Devasia
<https://ccta2026.ieeecss.org/>
- **INTERNATIONAL SYMPOSIUM ON ADVANCED CONTROL OF INDUSTRIAL PROCESSES (ADCONIP) 2026**
29 November–1 December
Auckland, New Zealand
General Chairs: Brent Young and David Wilson
Program Chairs: Ikuro Mizumoto and Wei Yu
<https://www.adconip2026.org/>
- ▲ **IEEE CONFERENCE ON DECISION AND CONTROL (CDC) 2026**
15–18 December
Honolulu, HI, USA
General Chair: Jorge Cortés
Program Chair: Lacra Pavel
<https://acc2026.a2c2.org/>
- **38TH CHINESE CONTROL AND DECISION CONFERENCE (CCDC) 2026**
15–18 May
Nanjing, China
General Chairs: Guang-Hong Yang, Bin Jiang, Xiang Ling, and Jia Liu
Program Chairs: Changyun Wen and Zhong-Ping Jiang
<http://www.ccdc.neu.edu.cn/>
- ▲ **AMERICAN CONTROL CONFERENCE (ACC) 2026**
27–29 May
New Orleans, LA, USA
General Chair: Randy Beard
Program Chair: Eric Frew
<https://acc2026.a2c2.org/>

Digital Object Identifier 10.1109/MCS.2026.3659163
Date of current version: 26 March 2026



C. Beck
President
University of Illinois,
Urbana-Champaign
beck3@illinois.edu

M. Egerstedt
Past-President
University of California, Irvine
magnus@uci.edu

L. Xie
Vice President
Conference Activities
Nanyang Technological University
elhxie@ntu.edu.sg

J. How
Vice President
Financial Activities
Massachusetts Institute of
Technology
jhow@mit.edu

N. Ozay
Vice President
Member Activities
University of Michigan
necmiye@umich.edu

A. Serrani
Vice President
Publication Activities
Ohio State University
serrani@ece.osu.edu

M. di Bernardo
Vice President
Technical Activities
University of Naples Federico II
mario.dibernardo@unina.it

A. Ferrara
Director of Operations
The University of Pavia
antonella.ferrara@unipv.it

K. Johansson
Vice-President/Diversity, Outreach,
and Development
KTH Royal Institute of Technology
kallej@kth.se

PUBLICATIONS

IEEE Control Systems
A. Annaswamy
Editor-in-Chief

IEEE Control Systems Letters
L. Menini
Editor-in-Chief

IEEE Open Access J. of Control Systems
S. Martinez
Editor-in-Chief

IEEE Trans. on Automatic Control
A. Astolfi
Editor-in-Chief

IEEE Trans. on Control Syst. Technology
I. Kolmanovsky
Editor-in-Chief

*IEEE Trans. on Control of Network
Systems*
L. Pavel
Editor-in-Chief

Electronic Information
M. Di Benedetto
Editor

CONFERENCES

A. Aghdam
Conference Editorial Board Chair

General Chairs
C. Beck, ACC 2025
Robert Bitmead, CCTA 2025
Y. Paschalidis and J.C. Basilio,
CDC 2025
R. Beard, ACC 2026
R. Nagamune, CCTA 2026
J. Cortes, CDC 2026

**Members of the CSS
Board of Governors**
Term Ending 31 Dec. 2025 (appointed)

S. Coogan
P. Pena
P. Vernekar

Term Ending 31 Dec. 2025
M. Di Bernardo
S. Glavaski
S. Mastellone
G. Nair
M. Oishi
N. Ozay

Term Ending 31 Dec. 2026 (appointed)
H. Chen
K. Fregene
K. Rudie

Term Ending 31 Dec. 2026
G. Giordano
R. Jungers
S. Knorn
S. Onori
A. Serrani
A. Isaksson

Term Ending 31 Dec. 2027
D. Dong
A. Fekih
M. Mesbahi
M. Polycarpou
H. Sandberg
E. Tegling

TECHNICAL COMMITTEE CHAIRS

G. Inalhan
Aerospace Controls
A. Katriniok
Automotive Controls
J. Hedengren
Control Education
X. Yin
Discrete-Event Systems
L. Paunonen
Distributed Parameter Systems
M. Almassalkhi
Energy Systems
S. Pequito
Health and Medical Systems
M. Zamani
Hybrid Systems
T. Yucelen
Intelligent Control
Y. Cao
Manufacturing Automation &
Robotic Control
E. Dell'Anese
Networks and Communication
Systems
A. Serrani
Nonlinear Systems and Control
J. Scruggs
Power Generation
A. Mesbah
Process Control
D. Dong
Quantum Computing, Systems,
and Control
N. Ozay
Robust and Complex Systems
Q. Zhu
Security and Privacy
M. Salazar
Smart Cities
B. Pasik-Duncan
Stochastic Systems and Control
A. Singh
Systems and Synthetic Biology
S. Formentin
Systems Identification and
Adaptive Control

E. Usai
Variable Structure and Sliding
Mode Control

STANDING COMMITTEE CHAIRS

T. Parisini
Awards
C. Tomlin
Awards Nomination
H. Guoqiang
Chapter Activities
A. Aghdam
Conference Editorial Board
(CDC, ACC)
S. Di Cairano
Conference Editorial Board (CCTA)
E. Chong, F. Jabbari
Conference Operations
S. Galeani
Conference Publications
S. Lall
Distinguished Lecturer
F. Pasqualetti
Electronic Information
A. Serrani* ex officio
Ethics in Publishing
C. Beck* ex officio
Executive Committee
P. Seiler
Fellow Evaluation
A. Papachristodoulou
Fellow Nominations
A. Isidori
History
L. Auret
Industry Activities
C. Beck* ex officio
Long Range Planning
R. Padariya
Standards
F. Parise
Student Activities
D. Gayme, A. Fekih
Women in Control
S. Mastellone
Diversity, Outreach &
Development Task Force

WWW.IEEECSS.ORG

>> IEEE CONTROL SYSTEMS SOCIETY

If you are already a member of the IEEE, you can join the IEEE Control Systems Society. As a member of the IEEE Control Systems Society, you receive *IEEE Control Systems* six times per year. Electronic subscriptions to *IEEE Transactions on Automatic Control* (published 12 times per year) and *IEEE Transactions on Control Systems Technology* (published six times per year) are also free with membership. Information on student member and nonmember dues/rates and print subscriptions is available on the CSS website.

Mail to: IEEE Cash Processing Department, 445 Hoes Lane, Piscataway, NJ 08854, U.S.A.

YES! I wish to join the IEEE Control Systems Society for one year at US\$25.00.

IEEE member number _____ Grade _____

Name and mailing address for IEEE mail (please print or type):

Check or money order enclosed payable to IEEE (U.S. dollars, drawn on U.S. bank only) or

VISA MASTERCARD DINER'S CLUB EUROCARD

Charge card number: _____ Exp. Date month/year _____

Full Signature _____

MATLAB FOR AI

Boost system design and simulation with explainable and scalable AI. With MATLAB and Simulink, you can easily train and deploy AI models.

mathworks.com/ai

 **MathWorks®**
Accelerating the pace of engineering and science