

Robust Control: A Workshop on the occasion of Keith Glover's 80th Birthday Talk Titles and Abstracts

Dr Rick Hyde
University of Bristol, Matlab

With the arrival of generative AI what is the future for control engineering as a discipline?

Drawing on my career experience gained in academia, industry, and commercial software tools, I will explore what the advent of generative AI might mean for control engineering as a discipline taught by universities. For me, the skills learnt from when I was one of Keith's PhD students have played a pivotal role in practically everything I have done since. The advent of potentially useful LLMs with OpenAI's GPT-2 in 2019 has already been a disruptor to many industries, and in this presentation I will explore what the impact might have been on my research on control of the VAAC Harrier had it been available at the time. I will also draw on other examples from my career to test, for example, the ability of LLMs to make inferences from one area of engineering science to another. Based on this I will put forward the case that control engineering remains as relevant as it ever was, and that LLMs can be viewed as a productivity tool for the control and/or system modeler rather than as a threat.



Professor Sanjay Lall
Department of Electrical Engineering, Stanford University

Distributed computing using feedback control in the data center

The bittide system is a recent approach to distributed computing, designed to achieve synchronous execution at a large scale without the need for a global clock or traditional wall-clock synchronization. It aims to overcome the complexity and expense of maintaining precise wall-clock time in distributed systems, especially at datacenter scale. The underlying mechanism of bittide uses feedback control to regulate the frequency of the hardware oscillators driving both computation and communication, in such a way as to ensure that all nodes operate in syntony. This allows applications to treat time as purely logical, and to make use of deterministic scheduling and programming methodologies across the entire datacenter. In this talk we present an overview of the bittide system, discussing how the system works, mathematical formulations of the system behaviour, and the consequences of logical synchrony for applications.



Professor Geir Dullerud
Department of Electrical and Computer Engineering, University of Minnesota

Learning and System Identification for Safety and Control Design in Dynamical Processes



AI-based methods have tremendous potential for impacting the performance of autonomous and robotic systems. Such systems include drones, ground- and water-based vehicles, and limbed robots for instance. A barrier to the wide deployment of AI-powered methods in such applications is the risk or unpredictability of algorithm performance. In this presentation we consider the development of safe machine learning (ML) methods for control that provide guarantees about their convergence and performance. First: we present an algorithm and a tool for statistical model checking (SMC) of continuous state space Markov chains initialized to a prescribed set of states.

We observe that it can be formulated as an X-armed bandit problem, and therefore, can be solved using hierarchical optimistic optimization. Our experiments, using our tool HooVer, suggest that the approach scales to realistic-sized problems and is often more sample-efficient compared to other existing tools. Second: We address system identification of a new important class of nonlinear systems: nonlinear systems whose vector fields are real-analytic functions of the state variables and linearly parametrized. Such models are relevant in a wide range of engineered systems. We consider two system identification methods: least-squares estimation (LSE), which is a point estimation method; and set-membership estimation (SME), which estimates an uncertainty set that contains the true parameters. We provide new non-asymptotic convergence rates for LSE and SME under i.i.d. control inputs and control policies with i.i.d. random perturbations, both of which are considered as non-active-exploration inputs. We also illustrate these theory results empirically on some well-known examples.

Professor Kemin Zhou
School of Robotics and Automation, Nanjing University, China

A Framework for High Performance Robust Control

This talk introduces a robust control framework, designated as the multi-objective complementary control (MOCC). This framework holds the potential to circumvent the typical conflicting relationship between desired performance and robustness inherent in traditional control methods. In essence, with the judicious application of MOCC, it becomes feasible to attain specific performance targets without compromising the robustness characteristics of the control systems.



Dr Paul Dickinson
Technical Specialist, CamMotive

Evolution of Automotive Powertrain Controls
(A potted history with cameos from Keith's alumni)



Automotive powertrains have evolved from low numbers of mechanical actuators to complex multivariable systems with a high number of control variables. Recent internal combustion engine (ICE) powertrains have complex systems of significant control dimensions to achieve the legislative emissions requirements. Real-driving-emissions (RDE) require emissions now cover a wide range of driving cycles and style which makes the optimization a significant challenge.

Electric powertrains initially appear much simpler. Tailpipe emissions are no longer a problem, but the controls challenge on dimensionality is significant and key states are not easily observed.

This talk provides a brief overview of how the complexity and approach to automotive powertrain controls have advanced overtime.

Professor Anders Rantzer
Department of Automatic Control, Lund University

On Optimal Dual Control

The interplay between learning and control has a history dating back at least half a century, but a surge of new activities has recently been stimulated by progress in data-driven learning algorithms with strong ties to optimal control. This presentation will focus on concepts and ideas in the intersection between control and computer science. Special attention will be given to recent progress in the theory of dynamic programming, where new ideas give a deeper understanding of dual control and the classical exploration-exploitation tradeoff.



Professor Glenn Vinnicombe
Control Group, Department of Engineering, Cambridge University

The convergence, or otherwise, of Monte-Carlo control

A key way in which reinforcement learning (RL) methods vary is in how heavily they depend on bootstrapped estimates. Q-learning sits at one extreme: it estimates the value by combining a single-step simulation with the current value of the next steps via the Bellman equation. At the other extreme, Monte-Carlo control estimates values from complete rollouts of the current decision making function. Between these extremes are methods that simulate over some finite horizon and then use the current value estimate as a terminal cost.



It has been known since 1992 that Q-learning converges to the optimal solution under very general conditions. By contrast, the convergence of Monte-Carlo control remains an open problem, except under highly specific uniformity conditions that are generally impossible to satisfy in practice.

As Richard Sutton wrote in 1999 "It is hard to imagine any RL method simpler or more likely to converge than this, yet there remains no proof of asymptotic convergence to [optimality]. While this simplest case remains open we are unlikely to make progress on any control method for [intermediate horizon lengths]"

In this talk we provide practical conditions under which Monte-Carlo control is guaranteed to converge to optimality, along with surprising examples of non-convergence to optimality for common implementations. We also propose a precise conjecture that guarantees convergence under a broad class of "sensible" conditions, and explain why we believe the conjecture is true. This is joint work with Octave Oliviers

Professor Rodolphe Sepulchre
Control Group, Department of Engineering, Cambridge University

A neuromorphic architecture for hierarchical event-based control design



We present a neuromorphic architecture for hierarchical event-based control design. The architecture has the classical graphical representation of an electrical circuit. The circuit is called neuromorphic because the nodes (voltage potentials) are excitable and the currents (branches) are synaptic. The excitability of the nodes shapes the events (spikes or bursts) whereas the topology of the circuit shapes the sequences of events (spatio-temporal patterns). The hierarchy of the design stems from the definition of an event at a coarser scale as a rhythm of events at a finer scale. We discuss the design of the nervous system of a snake machine to illustrate the unique role of

excitability in enabling an architecture that combines the adaptation of the continuous and the reliability of the discrete.

Professor Mustafa Khammash
Department of Biosystems Science and Engineering, ETH Zurich

Cybergenetic Cellular Control Systems: From Theory to Therapy

Cybergenetics unites control theory and engineering biology to endow cells with regulatory functions once limited to engineered manmade systems. In this lecture, I will describe how control theoretic insights are being translated into the design and realization of genetic control circuits that operate reliably within the noisy, nonlinear environment of living cells.

I will introduce a universal internal model principle (IMP) for living systems, showing how integral feedback and its generalizations emerge naturally as the only architectures capable of achieving kinetics-independent robust tracking and disturbance rejection.

I will present examples of genetic feedback architectures implemented in living cells, and will discuss how these same principles are being applied to create therapeutic cells capable of autonomously sensing disease cues and restoring physiological balance. Together, these advances chart a path from the mathematics of feedback to the engineering of adaptive, intelligent cell-based therapies.



Professor John Doyle
Control and Dynamical Systems, Caltech

The road to DGKF and MuTools



I've had fantastic collaborators and students, but Keith is a singularity. Over 30% of my citations are to publications we coauthored. We both escaped the US academic modern control purgatory in 1976, to Honeywell and Cambridge. He was then central to the 1980's consolidation of robust control by uniquely spanning math and engineering; from the 1984 Mpls workshop, to my yearly pilgrimages to Cambridge, culminating in DGKF and ultimately MuTools. Unfortunately, many of our wonderful students and collaborators are now gone but not forgotten.

Bio: I don't play golf, except with Keith.