Bounds and low-rank approximation for controlled Markov processes

by

Flemming Holtorf

Stochastic processes have captivated scientific interest by balancing conceptual simplicity with the ability to model complex, poorly understood, or even entirely unknown phenomena. Still, the deployment of stochastic process models remains hampered in practice by their intrinsically uncertain nature, complicating the computation and interpretation of model predictions. This thesis addresses two distinct challenges for the study of a rich class of Markov processes and associated control problems: certification and scale.

In the first part of this thesis, we present a framework for conservatively answering the question: What is the best performance a controlled jump-diffusion process can attain? Answers to this question, even if conservative, shed light on fundamental limits, allowing us to distinguishing situations where intrinsic noise masks poor decisions from situations where any attempt of improvement is futile. We connect infinite-dimensional linear programming over cones of occupation measures to techniques for approximating solutions to Hamilton-Jacobi-Bellman and Kolmogorov backward equations. The result is a hierarchy of structured sum-of-squares programs that furnishes a sequence of hard, yet computable bounds for common control performance indicators encoding, for instance, operating cost or the probability of failure. These bounds in turn serve as witnesses of fundamental limitations or certificates of optimality and safety. We apply the developed framework to the analysis of stochastic chemical systems, where it unifies and extends a range of recently proposed schemes for moment bounding and finite state projection. Furthermore, we leverage it to map out the performance limits of controlled quantum devices and, by extension, quantum information technologies at large. In the context of open-loop quantum control, our framework extends quantum speed limits, such as the Mandelstam-Tamm bound, and provably allows characterization of performance limits to arbitrary precision. For closed-loop controlled quantum systems, it constitutes the first approach to rigorously bound performance losses induced by continuous measurement.

The second part of the thesis is devoted to the challenge of scale as commonly encountered when studying Markov process models in high-dimensional spaces. Here, the complexity of computing and representing model predictions routinely exceeds available resources and renders interpretation challenging. We develop computational tools based on dynamical low-rank approximation that allow us to extract the dominant characteristic features of processes described by vast, nonlinear matrix-valued differential equations and track their evolution over time.

The methods developed in this thesis are accompanied by software solutions exploiting the features of the Julia programming language to enable deployment of Markov process models in the context of scientific inquiry and engineering advancement more widely, with greater ease, and rigorous guarantees.

Thesis supervisor: Alan Edelman
Title: Professor of Applied Mathematics, Department of Mathematics

Thesis supervisor: Richard D. Braatz
Title: Edwin R. Gilliland Professor, Department of Chemical Engineering