Statement of Research

on New Frontiers for Computation in Control

Applications in Energy and Biology

Matthew M. Peet

Over the last 20 years, we have gotten very good at using computation to control linear state-space models. My research question is: Where do we go from here?

Intractable Problems Most control problems in the engineering world are complicated; they have multiple features such as delay, nonlinearity, uncertainty and switches. There might be simple states, such as position and velocity, or there might be distributed states, such as temperature in a rod. In my research I work with many types of models. For example (funded by the French government) I work on control of plasma in a fusion reactor - a problem which is defined by the spatial distribution of temperature and magnetic field. This system is nonlinear, uncertain and modeled using partial-differential equations. I have worked on the problem of modeling internet congestion control. This model has nonlinearity, delay and logical switching. Additionally, I work on understanding how the immune system controls diseases such as cancer. These models are nonlinear and uncertain and are based on highly suspect conjectures of how the immune system operates. In each of these cases, the underlying control problem is, by the standard definition of NP-hard, intractable. Yet despite this fact, the problems remain important and are too complex to solve by hand. So my question is: How do we use computation to solve these control problems?

While I don't have an easy fix that will make intractable problems tractable, I do advocate a standardized computational approach to solving these problems (funded by NSF CAREER).

A Changing Model for Computation The speed of a single-core CPU has not increased significantly since 2005. Instead of building faster chips, manufacturers are focusing on multi-core platforms, combining hundreds or thousands of computing cores with relatively low latency. Currently, 8-core chips are commercially available for desktop computers, with 16 core chips expected by the end of the year and 48 core chips in testing. As computing architectures become exponentially multi-core, sequential algorithms are becoming increasingly obsolete. This is bad news for control, as vector and matrix optimization - the foundation for many control algorithms - is known to be an inherently sequential problem. This means that unless $P = NC$ (NC is the class of problems with parallel solutions), there will never be an efficient parallelization of linear or semidefinite programming. This leaves us with two choices: we can abandon optimization as a foundation of control (move the field backward 30 years); or we can constrain our work to have a structure which can be exploited in a massively-parallel-computing environment.

In my work, (supported by the NSF), I have built massively-parallel algorithms to create robust controllers for systems with uncertainty and nonlinearity.

Polynomial Computing and Sum-of-Squares Intractable problems by definition admit no polynomial-time solution. Control technology has advanced quickly over the past 20 years due primarily to the use of polynomial-time convex-optimization-based algorithms such as LMI solvers and semidefinite programming. To solve "intractable" control problems, I still use convex optimization, but I relax the polynomial-time constraint. In particular, I look for a sequence of approximation algorithms, each of which is polynomial time and has bounded error ϵ_i with $\lim_{i\to\infty} \epsilon_i = 0$. See Figure 1 for an example of using an asymptotically exact algorithm for stability analysis of a nonlinear PDE [1]. The figure shows ϵ_i vs. i. To create these approximation algorithms, I use polynomial computing and Sum-of-Squares (SOS).

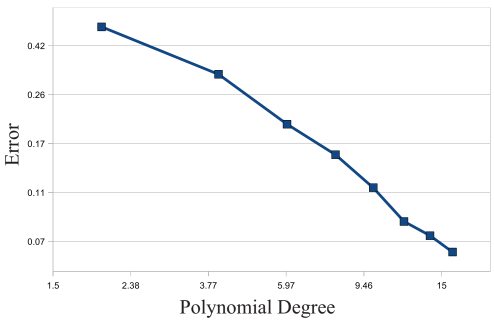

Figure 1: Log-Log plot of accuracy vs. polynomial degree for analysis of a PDE.

Polynomial computing is the optimization of polynomial variables subject to convex constraints. A simple example is

$$\max c^T x : \quad x^T f(y) \geq 0 \qquad \text{for all } y \in \mathbb{R}^n$$

where f is a vector of polynomials. Polynomial computing is convex, yet NP-hard. To solve these problems, we use asymptotic methods such as Sum-of-Squares. Sum-of-Squares (SOS) is a parametrization of positive polynomials of fixed degree using positive matrices. By optimization of the positive matrices, we can solve polynomial computing problems for a fixed polynomial degree. As we increase the degree of the polynomial, the accuracy of the solution increases.

Areas of Applied Research

In my work, I consider two kinds of problems: 1) How to solve polynomial optimization problems; 2) How to use polynomial optimization to solve applied engineering problems. In this document, I first describe two important applied engineering problems on which I work. Following these applications, I describe some of the more fundamental aspects of my research.

Tolerance and Control in the Immune System In October, 2011, the Nobel prize in Medicine was given jointly to Bruce Beutler, Jules Hoffman and Ralph Steinman for contributions to understanding activation of the immune system via dendritic cells. Unfortunately, one of these three, Dr. Steinman, died from pancreatic cancer three days before the prize was announced. It seems that Dr. Steinman had spent the previous three years, in collaboration with scientists from around the world, trying to design a customized treatment strategy based on the activation of dendritic cells. Sadly, it was the failure of this activation to trigger a full immune response that led to the failure of the treatment and ultimately, his death.

Understanding how the immune system can be used to fight disease has been a singular focus of medical research since vaccines were first developed by Jenner in the 1770s (prior to vaccination, innoculation has been used since the 1500s), with continued research leading to recent Nobel prizes in 1972, 1977, 1980, 1984, 1987, 1996, 1997, 2008 and 2011. As we understand it, the immune system is a network of cells which identifies and eliminates harmful infections or toxins while maintaining tolerance to proteins which it identifies as self. Failure of the immune system to recognize and tolerate self is what leads to auto-immune disease such as Type I diabetes, organ rejection and allergies. On the other hand, failure to recognize cancer cells as harmful is what results in malignant tumours. The mechanism by which the immune system makes a self-nonself determination and communicates this information is unknown.

In my research, in collaboration with Dr. Lee at Stanford, Doron Levy at U. Maryland and Peter Kim at U. Sydney, we view the immune system as a decentralized decision and control system which has been optimized by evolution to minimize the dual threats of pathological infection and autoimmune disease. Using dendritic cells, macropahages and antigen-presenting-cells (APCs) as sensors, we model the decision process through the interaction of countervailing effector and regulatory T cell populations and we model the communication process through a combination of mutual recognition and cytokine signalling. Some results include the following.

- At INRIA, I initiated a 1-year study of Myelogenous Leukemia funded by the Research Council of the National Center for Research in Computer Science and Control and conducted in collaboration with hematologists and other members of the ARC ModLMC group. We view the chronic state of Leukemia as a stable equilibrium between cell growth and immune response. By contrast, the blast phase of Leukemia, resulting in Acute Myelogenous Leukemia, is due to a change in patient parameters which yield a transition to unstable equilibria [2, 3].

- In [4], we presented a model for self-tolerance in the immune system. In this model, we proposed that the immune system reacts primarily to the growth rate of antigen, as opposed to simply the concentration. Using delay in regulatory response we proposed a biological model with a delay-induced differential controller acting on a cytokine switch. Additionally, adaptive regulatory cells induce integral contraction of the response. We used polynomial optimization techniques to map the parameter region for which the contractive response was stable.

Control of Fusion Energy In 2008, a panel convened by the National Academy of Engineering at the request of the National Science Foundation identified "Provide Energy from Fusion" as one of the 14 engineering Grand Challenges for the 21st Century. Fusion energy refers to the difference in potential energy between particles in free state and particles bound together by the strong nuclear force. The energy released from fusion reactions per unit mass is substantially greater than that from fission reactions. For example, the ^{2}H + ^{3}H to ^{4}He + ^{1}n fusion reaction results in a release of 3.5 MeV/nucleon energy, while the decay of ^{235}U results in an approximately .85 MeV/nucleon release of energy. Unfortunately, while it is relatively easy to initiate fission in ^{235}U, creating and controlling plasma at the temperatures needed for fusion is a monumental engineering challenge.

In this project, supported by the French government through the Chateaubriand program, the primary goal is stabilization of high-temperature plasma instabilities known as magnetohydrodynamic (MHD) modes. We have a collaboration between Emmanuel Witrant (U. Grenoble), Eugenio Schuster (Lehigh), and Didier Moreau (Tore Supra). Prof. Witrant, in particular, has spent more than a month at IIT working on plasma control problems. Our close collaboration allows us access to real-time data, advanced models and the opportunity to implement controllers in functioning Tokamak reactors. The following are contributions to elimination of MHD modes.

- As discussed in more detail below, we have created a new and effective framework for computational design of controllers for PDE systems.
- We have designed a method for synthesizing observers-based controllers for PDE models [5].
- We have designed observer-based controllers for stabilization of the safety factor profile in Tokamaks, resulting in a >30% improvement in performance in simulation [6].

My goal for this research is to design and implement controllers capable of increasing the efficiency of fusion reactors.

Contributions to Basic Research in Control

In the following, I discuss contributions to control outside the context of specific applications.

Solving Large-Scale Problems using Parallel Structure Even if our effort to quantify the complexity of NP-hard problems is successful, for most practical problems, the computational burden will still be too high for current desktop computers. The problem is that single-core processing power is no longer increasing. Instead, chip manufacturers have focused on the distributed computing paradigm best articulated by Mark Weiser in 1991, wherein he described a world where computation wasn't limited to large powerful machines. In this world, computation was everywhere distributed and networked among countless small, specialized devices. In the intervening 20 years, reality has caught up with this vision. Now we have small, specialized computational devices which have become ubiquitous and are connected via LAN, WWAN, WiFi and Bluetooth - Indeed the venerable PC desktop itself has started to break down from a single CPU powerhouse into a collection of core-processors distributed from the CPU to the motherboard to the graphics card to the hard drive.

In this new scenario, we must create algorithms which decentralize efficiently. Unfortunately, interior-point algorithms (the workhorse of optimization and control), do not have a decentralized structure. The goal of our lab , supported by the NSF, is to create massively parallel algorithms to solve polynomial computing problems. These algorithms must have both distributed computation and distributed memory structures.

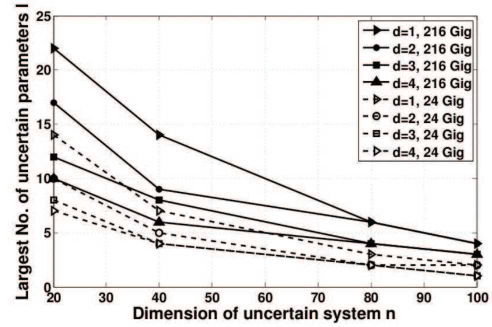

Figure 2: Maximum Uncertain parameters for an n-Dimensional System with Accuracy Level d using an 80-node Cosmea Linux cluster with either 216GB or 24GB of RAM

- We have developed an efficient parallel framework for solving polynomial computing problems on the simplex by using Polya's theorem instead of the Sum-of-Squares methodology [7, 8].
- In collaboration with Argonne National Labs, we have created an efficient MPI-based parallel implementation of Polya's theorem to set up and solve polynomial computing problems [9].
- These algorithms have been adapted to and tested on multi-core computers, cluster computers, and the IBM Blue Gene supercomputer for large-scale (100+ states) robust control problems. The algorithm demonstrates linear speedup using up to 200 nodes. The potential of the algorithm is shown in Figure 2 which shows the maximum number of uncertain parameters vs. the number of states and desired accuracy for an 80-node Linux cluster with either 24GB or 216GB of RAM.

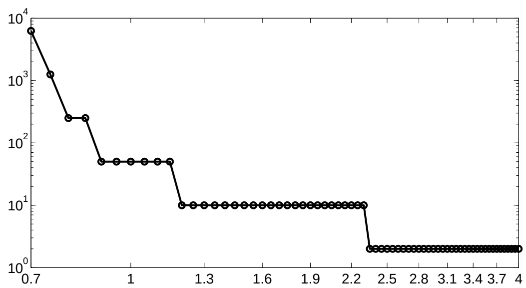

Figure 3: Degree bound vs. Exponential Convergence Rate

Converse Lyapunov Theory Lyapunov functions are a well-known tool for stability analysis of nonlinear and infinite-dimensional systems.

To quantify the tradeoff between accuracy and complexity when using polynomial Lyapunov functions, I have developed some results on converse Lyapunov theory. For a nonlinear system

$$\dot{x}(t) = f(x(t)),$$

I have shown that if f is thrice differentiable, then exponential stability implies the existence of a polynomial Lyapunov function [10–12]. In [13], I gave a quantitative bound on the Lipschitz continuity factor of converse Lyapunov functions. In [14, 15], I have shown that if f is polynomial, there exists a sum-of-squares Lyapunov function and I have given a *bound on the degree*. This bound is shown in Figure 3. Note that this bound decays to two which means that a sufficiently stable system will always have a quadratic Lyapunov function. We have also shown that exponentially stable systems with delay also have Lyapunov functions defined by polynomials [16–18]. This is discussed in more detail below.

My continuing goal for this research is to develop bounds on accuracy of the polynomial computing test as a function of compexity.

Control of Systems with Delay Delays are a common feature in control systems which involve communication, discretization, or transport. However, analysis of systems with delays is NP-hard. We have used polynomial computing to greatly enhance understanding of these systems. The following are our contributions.

- We have shown how to recast stability of linear time-delay systems as polynomial computing [19–24] with no conservatism [16–18].
- We have extended our result to analysis of nonlinear delayed systems [25, 26], systems with time-varying delays [27] and singular systems with delay [28–31].

- We have proposed a theory of duality for linear time-delay systems [32].
- We have reduced the complexity of the computation by several orders of magnitude [28–30, 33].
- We have done some analysis of delayed problems in networking and biology [3].
- Polynomial computing for frequency-domain methods in delay can be found in [31, 34–36].
- We considered control of sampled-data systems in [37].

Control of Partial-Differential Equations Partial-differential models are used to describe elasticity, fluid motion and cellular interaction, among other things. Although PDE models are ubiquitous in engineering, our ability to compute controllers for these systems relies almost exclusively on discretization. Funded by the NSF CAREER program, our lab has made it possible to design controllers for PDEs in a systematic way by converting the problem to polynomial computing. The following are some contributions.

- We have used SOS to design observers for PDE models of heat transfer and for plasma in Tokamaks [5].
- We have used SOS to design observer-based controllers for linear PDE models of plasma in Tokamaks [6].
- We have proposed a method for using polynomial computing for analysis of nonlinear PDEs and demonstrated asymptotic accuracy [1].

My goal for this research area is to reproduce or exceed our success with delayed systems by reducing optimal control of partial-differential systems to a polynomial computing problem.

Control of Decentralized Networks The need for efficient exchange of information has led to the creation of large communication networks. Examples of such networks include the Internet, telephone, postal and power distribution services. In each case, the goal is to transmit information using decentralized control mechanisms. For many systems, ad hoc transmission controls have been adopted which often work well in the short term, but usually fail as the systems expand. Decentralized control systems can be seen as solving large-scale optimization problems using distributed resources.

$$\max \quad \sum_{i=1}^{k} f_i(x_i), \quad \text{subject to:} \quad Bx \geq c.$$

In large systems, centralized computation becomes infeasible. Furthermore, the presence of information constraints makes centralized implementation difficult. The goal, then, is to develop decentralized controls which, when combined, cause the system to converge to the global optimum.

For Internet congestion control, I considered decentralized controllers based on the application of the gradient-projection algorithm to the dual of the optimization problem. My contributions include the following.

- We proved global convergence of the Internet to optimality given a certain bound on the feedback gain at the source and using a hybrid, nonlinear, time-delayed model of the Internet [19, 38]. This bound was shown to be both necessary and sufficient.
- We considered an approach to Internet congestion control based on optimization algorithms applied to the primal problem with dynamics at the source and not the link [39]. This is the first proof of convergence of this model of Internet congestion control for an arbitrary topology.
- In [37], we showed how to use Sum-of-Squares to control continuous systems by feedback over an unreliable network (e.g. with packet dropping).

My long-term goal in this area is to improve understanding of the limits of decentralized control protocols and to apply the lessons from Internet congestion control to other types of decentralized networks.

References

[1] A. Papachristodoulou and M. M. Peet. On the analysis of systems described by classes of partial differential equations. In *Proceedings of the IEEE Conference on Decision and Control*, pages 747–752, Dec. 2006.

[2] M. M. Peet, P. Kim, S.-I. Niculescu, and D. Levy. New computational tools for modeling chronic myelogenous leukemia. *Mathematical Modeling of Natural Phenomena*, 4(2), 2009.

[3] A. Papachristodoulou and M. Peet. *Topics in Time Delay Systems*, volume 388 of *Lecture Notes in Control and Information Sciences*, chapter SOS Methods for Nonlinear Delayed Models in Biology and Networking. Springer, October 2009.

[4] M. Peet, P. Kim, and P. Lee. Biological circuit models of immune regulatory response: A decentralized control system. In *Proceedings of the IEEE Conference on Decision and Control*, 2011.

[5] A. Gahlawat and M. M. Peet. Designing observer-based controllers for PDE systems: A heat-conducting rod with point observation and boundary control. In *Proceedings of the IEEE Conference on Decision and Control*, 2011.

[6] A. Gahlawat, M. Peet, and E. Witrant. Control and verification of the safety-factor profile in tokamaks using Sum-of-Squares polynomials. In *IFAC World Congress*, 2011.

[7] R. Kamyar and M. Peet. Decentralized computation for robust stability analysis of large state-space systems using polya's theorem. In *Proceedings of the American Control Conference*, 2012. Submitted.

[8] R. Kamyar, M. Peet, and Y. Peet. Solving large-scale robust control problems by exploiting the parallel structure of Polya's theorem. *IEEE Transactions on Automatic Control*, 2011. Submitted.

[9] M. M. Peet and Y. V. Peet. A parallel-computing solution for optimization of polynomials. In *Proceedings of the American Control Conference*, 2010.

[10] M. M. Peet. Exponentially stable nonlinear systems have polynomial Lyapunov functions on bounded regions. In *45th annual Allerton Conference on Communication, Control and Computing*, Sept. 2007.

[11] M. M. Peet and P.-A. Bliman. Polynomial Lyapunov functions for exponential stability of nonlinear systems on bounded regions. In *Proceedings of the IFAC World Congress*, 2008.

[12] M. M. Peet. Exponentially stable nonlinear systems have polynomial lyapunov functions on bounded regions. *IEEE Transactions on Automatic Control*, 52(5), 2009.

[13] M. M. Peet. A bound on the continuity of solutions and converse Lyapunov functions. In *Proceedings of the IEEE Conference on Decision and Control*, 2009.

[14] M. Peet and A. Papachristodoulou. Sum-of-squares Lyapunov functions for exponential stability: An existence proof based on the picard iteration. In *Proceedings of the IEEE Conference on Decision and Control*, 2010.

[15] M. Peet and A. Papachristodoulou. A converse Sum of Squares lyapunov result with a degree bound. *IEEE Transactions on Automatic Control*, 2011. To Appear as Full Paper.

[16] M. M. Peet and P.-A. Bliman. An extension of the Weierstrass approximation theorem to linear varieties: Application to delay systems. In *Seventh IFAC Workshop on Time-Delay Systems*, Sept. 2007.

[17] M. M. Peet and P.-A. Bliman. The Weierstrass approximation theorem on linear varieties: Polynomial lyapunov functionals for delayed systems. In *Proceedings of Mathematical Theory of Networks and Systems*, 2008.

[18] M. Peet and P.-A. Bliman. On the conservatism of the sum-of-squares method for analysis of time-delayed systems. *Automatica*, 47(11), 2011.

[19] M. Peet and S. Lall. On global stability of internet congestion control. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1035–1041, Dec. 2004.

[20] M. M. Peet and A. Papachristodoulou. On positive forms and the stability of linear time-delay systems. In *Proceedings of the IEEE Conference on Decision and Control*, pages 187–193, Dec. 2006.

[21] M. M. Peet and C. Bonnet. On positive forms and the stability of linear systems. In *Conference de la SMAI sur lóptimisation et la decision*, 2006.

[22] M. M. Peet and A. Papachristodoulou. Joint positivity of Lyapunov operators for linear time-delay systems. In *Proceedings of the IEEE Conference on Decision and Control*, Dec. 2007.

[23] M. M. Peet and A. Papachristodoulou. Using polynomial semi-separable kernels to construct infinite-dimensional Lyapunov functions. In *Proceedings of the IEEE Conference on Decision and Control*, 2008.

[24] M. M. Peet, A. Papachristodoulou, and S. Lall. Positive forms and stability of linear time-delay systems. *SIAM Journal on Control and Optimization*, 47(6), 2009.

[25] M. Peet and S. Lall. Constructing Lyapunov functions for nonlinear delay-differential equations using semidefinite programming. In *Proceedings of the 6th IFAC Symposium on Nonlinear Control Systems(NOLCOS)*, pages 381–385, Aug. 2004.

[26] A. Papachristodoulou, M. M. Peet, and S. Lall. Stability analysis of nonlinear time-delay systems. *IEEE Transactions on Automatic Control*, 52(5), 2009.

[27] A. Papachristodoulou, M. M. Peet, and S.-I. Niculescu. Stability analysis of linear systems with time-varying delays: Delay uncertainty and quenching. In *Proceedings of the IEEE Conference on Decision and Control*, Dec. 2007.

[28] Y. Zhang, M. Peet, and K. Gu. Reducing the complexity of the sum-of-squares test for stability of delayed linear systems. *IEEE Transactions on Automatic Control*, 56(1), 2011.

[29] Y. Zhang, K. Gu, and M. Peet. Reducing the complexity of the sum-of-squares stability test for time-delayed systems. In *Proceedings of the American Control Conference*, 2010.

[30] M. M. Peet Y. Zhang and K. Gu. *Time Delay Systems - Methods, Applications and New Trends*, volume 423 of *Lecture Notes in Control and Information Sciences*, chapter Accelerating Convergence of Sum-of-Square Stability Analysis of Coupled Differential-Difference Equations. Springer, 2012.

[31] M. M. Peet, C. Bonnet, and H. Ozbay. *Topics in Time Delay Systems*, volume 388 of *Lecture Notes in Control and Information Sciences*, chapter SOS Methods for Stability Analysis of Neutral Differential Systems. Springer, October 2009.

[32] M. M. Peet and A. Papachristodoulou. Inverses of positive linear operators and state feedback design for time-delay systems. In *8th IFAC Workshop on Time-Delay Systems*, 2009.

[33] K. Gu, Y. Zhang, and M. Peet. Accelerating convergence of sum-of-square stability analysis of coupled differential-difference equations. In *9th IFAC Workshop on Time-Delay Systems*, 2010.

[34] M. M. Peet and C. Bonnet. Stability and computation of roots in delayed systems of the neutral type. In *Conference on Distributed Parameter Systems*, page 49, 2006.

[35] C. Bonnet and M. M. Peet. Using the positivstellensatz for stability analysis of neutral delay systems in the frequency domain. In *Seventh IFAC Workshop on Time-Delay Systems*, Sept. 2007.

[36] M. M. Peet, H. Özbay, and C. Bonnet. Sos methods for delay-dependent stability of neutral differential systems. In *Proceedings of Mathematical Theory of Networks and Systems*, 2008.

[37] A. Seuret and M. M. Peet. SOS for sampled data systems. In *IFAC World Congress*, 2011.

[38] M. M. Peet and S. Lall. Stability analysis of a nonlinear model of internet congestion control with delay. *IEEE Transactions on Automatic Control*, 52(3):553–559, March 2007.

[39] A. Papachristodoulou and M. M. Peet. Global stability analysis of primal internet congestion control schemes with heterogeneous delays. In *Proceedings of the IFAC World Congress*, 2008.