Model-based Design and Qualification of Complex Systems

John C. Doyle Richard M. Murray Michael Cross Gil Refael
Control and Dynamical Systems / Physics
California Institute of Technology

30 April 2007

This report summarizes the work supported by Boeing over the period of April 2006 to March 2007.

1 Project Overview

The broad goal of this project is to develop new theory, algorithms and demonstrations of model-based design strategies for complex systems. This activity is broken up into three broad themes:

1. *Robust Yet Fragile Behavior:* Study the robust-yet-fragile (RYF) nature of complex systems, and specifically to identify the common structures contributing to the RYF behavior, and develop both simple explanatory and detailed predictive models with associated analysis tools.

2. *Multi-scale Modeling:* Systems modeling theory and practice with emphasis on multi-resolution modeling, and managing multiple distinct product representations that must be mapped to each other.

3. *Engineering Implementation:* Apply analysis and methods in robust-yet-fragile behavior and multi-scale modeling to specific engineering systems of systems that will provide an evaluation of the efficacy of both the framework and the tools toward applications. Two specific testbeds are being used for this purpose: the Caltech multi-vehicle wireless testbed (MVWT) and the Caltech autonomous vehicle testbed ("Alice").

In addition to core theory relevant to model-based design and qualification of complex systems, we have to date focused on on two primary applications to test these: communications protocols for packet-based networks and complex and autonomous vehicles (Alice). These two applications represent different extremes of complex systems. The Internet consists of millions of nodes and is a true system of systems. The protocols that control traffic across this network are models for designing provably correct interfaces for systems in which there is no central design of the individual modules that comprise the system. Alice, on the other hand, is an example of a system with large complexity, but with a traditional engineering design approach: specifications and interfaces can be defined and managed centrally, with design teams responsible for building components and modules.

2 Core Theory

Our approach to the development of core theory has been to focus on case studies that allow us to build appropriate conceptual frameworks, methods, and mathematical tools. For example, the current Internet serves as perhaps an ideal case study in complex network systems in that it is an infrastructure of critical importance, yet our existing knowledge in communications, computing, and control does not provide a theoretical basis for why it works as well as it does. This apparent contradiction has invited attention to study of system architecture in a way that is not only intellectually challenging but of practical importance for the design of next-generation networks. Our use of optimization-based methods as a means of reverse-engineering system architecture has provided new insight into problems of system verification, and we continue to focus on robustness issues related to model specification.

2.1 Analysis of nonlinear dynamical systems using Sum of Squares (Dennice Gayme)

This research is concerned with the analysis of nonlinear dynamical systems using tools from Robust control theory and Sum of Squares methods. Much of the early worked focused on set invariance for various systems; examples include determining set membership for chaotic maps and computing regions that guarantee minimum signal to interference ratios for cellular networks. More recently we have begun to study turbulent shear flows, specifically concerning the mechanism that facilitates transition from laminar to turbulent flow.

Hydrodynamic stability theory is concerned with how laminar flows become unstable which is a precursor to turbulence [12]. Traditionally the stability of the flow is determined by linearizing the equations of flow (the Navier Stokes equations) around a laminar base flow and then studying the eigenvalues. This analysis is carried out at various Reynolds numbers and the first Reynolds number that yields an unstable eigenvalue is said to be the critical Reynolds number. However the results of such analysis tends to agree poorly with experiments in wall bounded shear flows in that they tend to transition to turbulence at Reynolds numbers far below the critical Reynolds number.

There is strong experimental evidence that external disturbances play a role in the early transition to turbulence. Further the fact that the linearization of the equations governing wall bounded shear flows result in non-normal operators leads to the potential for extraction of energy from any disturbances in the background flow, even without exponential instabilities. This has led to the widely held belief that linear mechanisms are crucial for the transition to turbulence and that non-normality of the linearized Navier Stokes operator is also a necessary condition.

In the case of Couette flow it has been shown that although linear analysis predicts stability for all Reynolds numbers, the perturbation energy actually grows as a function of Re^3. This amplification is believed to produce a turbulent profile consisting of nearly stream-wise constant vortices and streaks as shown in figure . There has been a lot of work suggesting that these stream-wise constant structures contain most of the turbulent energy.

Although linear mechanisms play an important role there are important questions remaining about the behavior in nonlinear regimes, given that the high amplification means that we are far away from the laminar behavior. Our work is aimed at gaining a greater understanding of the global nonlinear picture. As a first step we study a stream-wise constant version of the Navier Stokes equation proposed by Bobba [1]. The goal is to show that noise forcing of this so-called 2D/3C model along with the linearized full Navier Stokes equations can be used to explain most of

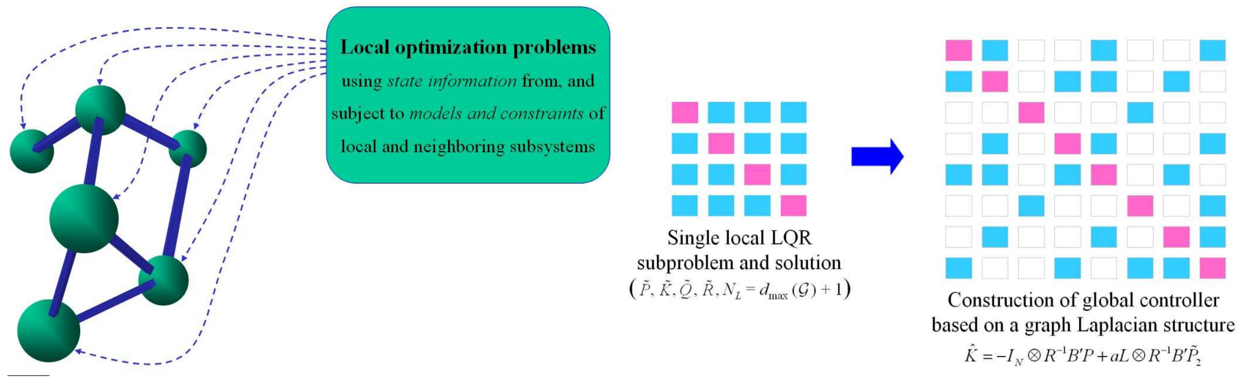
Local optimization problems
using state information from, and subject to models and constraints of local and neighboring subsystems
Single local LQR subproblem and solution
$(\tilde{P}, \tilde{K}, \tilde{Q}, \tilde{R}, N_L = d_{\max}(\mathcal{G}) + 1)$
Construction of global controller based on a graph Laplacian structure
$\hat{K} = -I_N \otimes R^{-1}B'P + aL \otimes R^{-1}B'\tilde{P}_2$

Figure 1: Illustration of distributed predictive control and LQR design schemes.

the turbulent behavior except for some small norm bounded uncertainty. The overall idea is that the full turbulent picture can be viewed as shown in figure

2.2 (Distributed Control Design using LQR and Predictive Control Techniques (Tamás Keviczky))

The size and complexity of many large-scale technological systems (such as power networks, complex chemical processes, highway traffic flows, the Internet), necessitate a decentralized, distributed approach to understanding and controlling them, due to restrictions on computational power and information exchange. Thus it is not surprising that interest and research in the field of distributed control spans over several decades. Distributed control techniques today can be found in a broad spectrum of applications ranging from robotics and formation flight to civil engineering and large-scale industrial processes. The renewed interest in the field arises from technological advances such as autonomous vehicles, cheap wireless sensors and the abundance of networks of independently actuated systems. Networks of vehicles in formation, production units in a power plant, a network of cameras at an airport, and mechanical actuators for deforming surface are just a few examples. At the same time, new methodologies and advances in computing provide new ways of approaching such problems, which were mostly treated in an ad-hoc manner before.

Special problem structure is often a crucial element in such endeavors, which enables progress and allows the construction of algorithms and methods that are applicable in practice. We seek to rely on such problem structure by focusing on a class of large-scale systems composed of dynamically decoupled and independently actuated subsystems. In order to represent control objectives in various application domains, we formulate an optimal control problem, where the cost function and constraints couple the dynamical behavior of the subsystems. The coupling is described through a graph where each system is a node and the control action at each node is based only on local and neighboring state information (see top half of Figure 1). Recently, we applied this design philosophy in a distributed model predictive control framework for various applications, including cooperative control of multiple vehicles [9, 6]. A review of our contributions and future work directions in this area is presented in [10].

Our current work is focused on studying large-scale systems composed of unconstrained identical decoupled subsystems and devising powerful control strategies based on insights gained from our previous results using distributed predictive controllers. We employ the same design philosophy of constructing a globally stabilizing controller from local subproblem solutions, which respect the global objective and rely only on local and neighboring subsystem state information. We developed a distributed control design method, which requires the solution of a single local LQR problem, whose size depends only on the maximum vertex degree of the interconnection graph [2]. A simple sketch of this idea is depicted in the lower half of Figure 1. The proposed design procedure illustrates how stability of the large-scale system is related to the robustness of local controllers and the spectrum of a matrix representing the sparsity pattern of the distributed controller design problem. If the sparsity pattern is chosen to be a weighted graph Laplacian of the interconnection, then stability can be linked to important properties, such as the algebraic connectivity of the graph. The problem structure enables the construction of stabilizing distributed controllers independently of the choice of weighting matrices, which leads to more freedom in tuning. This feature provides a very powerful modularity to our approach, since the system architecture can be modified by adding or removing elements without the need for controller redesign, as long as the maximum vertex degree is not increased.

Besides gaining insights into the interplay between interconnection structure, local controller robustness and stability of the overall system, we are currently investigating the application of such distributed LQR-based controllers in stability analysis and design of distributed predictive schemes such as [9].

2.3 Distributed gradient systems and dynamic coordination (Demetri Spanos)

We have studied formal models of coordination in heterogeneous networks of dynamic agents. Unlike previous work in computer science, we have approached this question using tools from dynamics, feedback, and control theory. We hope that this approach will yield techniques and design architectures that integrate seamlessly with the design of mobile dynamic agents, such as unmanned vehicles.

Among the aims of this work is a general-purpose modeling methodology for studying coordination phenomena in network agents, without a priori knowledge of network structure and without any guarantees about the time-evolution of the network; we consider groups that split and rejoin, lose nodes due to failure, and undergo other topology-changing events. Despite this, we are able to recover quantifiable performance of coordinated behavior across a wide range of network uncertainties. Figure 2 illustrates some of the results.

2.4 A Game-Theoretic Approach to Contention Control (Lijun Chen)

The goal of network design and control is to design a system in which individual agents (e.g., end users such as TCP sources and network components such as links or autonomous systems) interact in a way that achieves an equilibrium (or stable operating point) with some desired systemwide properties. The most widely adopted property or design desiderata is probably the social optimality, where the system optimizes some global objective function such as the aggregate user utility or the aggregate cost. However, optimization framework is not universal and has its limitations. The computational and informational constraints may prevent the system from achieving global optimality, and because of informational structure or incentive issue, some design problems intrinsically do not fit in an optimization paradigm. A more general modeling framework is (noncooperative) game theory. Game-theoretic models are inherently distributed. They directly model/specify the

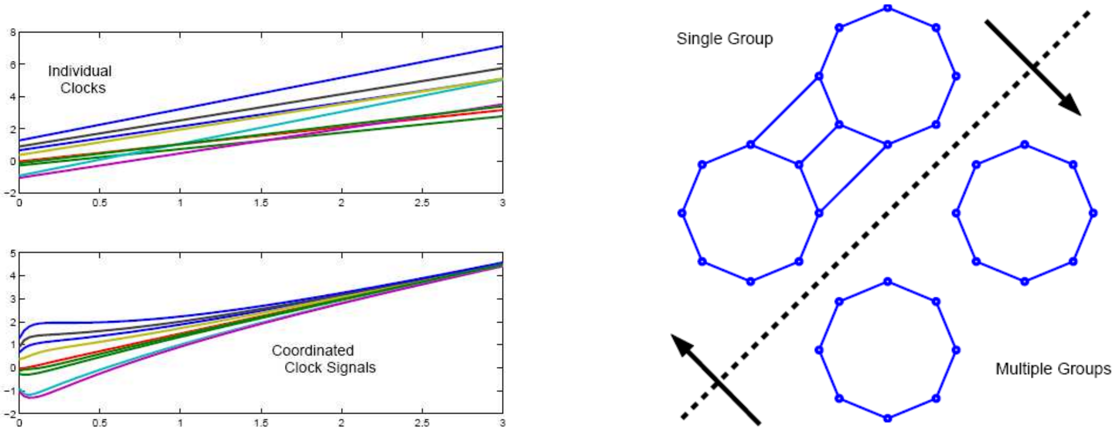

Figure 2: (a) Sample trajectories of the clock synchronizing dynamics. Each node has a local time that is independent of all the other clocks, and only constrained to grow linearly in time; the coordinating dynamics drives all of the members of the network to a single global time, independent of the underlying network topology. (b) Reconfigurable coordination refers to the ability to adapt to dynamic changes in group membership. In the case of the averaging system, we wish each connected component to track its own average quantity, and for the dynamics to naturally adapt to merging and splitting of groups.

behaviors of individual agents, so system- wide properties are emergent behaviors. Game theory also provides a series of equilibrium solution concepts such as the Nash equilibrium to predict the outcome of agent interaction. We will design network protocols according to distributed algorithms (strategy update mechanisms) achieving various kinds of game-theoretic equilibria.

Specifically, we apply this design methodology to medium access control, and develop a game-theoretic framework for contention control [3, 4]. We define a general game-theoretic model, called random access game, to capture the contention/interaction among wireless nodes in wireless networks with contention-based medium access. We characterize Nash equilibria of random access games, study their dynamics and propose distributed algorithms (strategy evolutions) to achieve Nash equilibria. This provides a general analytical framework that is capable of modeling a large class of systemwide quality of service models via the specification of per-node utility functions, in which systemwide fairness or service differentiation can be achieved in a distributed manner as long as each node executes a contention resolution algorithm that is designed to achieve the Nash equilibrium. We thus propose a novel medium access method derived from CSMA/CA according to distributed strategy update mechanism achieving the Nash equilibrium of random access game. Our access method adapts to a continuous contention measure – conditional collision probability, can stabilize the network into a steady state with a target fairness (or service differentiation) and high efficiency, and can decouple contention control from handling failed transmissions. As a case study of medium access control design in game-theoretic framework, we present a concrete medium access method and show that it achieves superior performance over the standard 802.11 distributed coordination function (DCF), and can provide flexible service differentiations among wireless nodes. In addition to guiding medium access control design, the random access game model also provides an analytical framework to understand equilibrium properties such as throughput, loss and fairness, and dynamic property of different medium access protocols and their interactions.

3 Applications

3.1 Sensing, Navigation and Reasoning Technologies for the DARPA Urban Challenge

(Note: this work is funded by DARPA, with some partial support of one student [Melvin Flores] through Boeing funding. Since the project serves as motivation for much of the theory and tools we are developing, we provide a description of the overall project here.)

Team Caltech is participating in the 2007 DARPA Urban Challenge, an autonomous vehicle race that will take place on 3 November 2007. Our primary technical thrusts are are three areas: (1) mission and contingency management for autonomous systems; (2) distributed sensor fusion, mapping and situational awareness; and (3) optimization-based guidance, navigation and control. Preliminary work in each of these areas, combined with a systematic approach to overall systems engineering and field testing, has demonstrated the ability for our system to navigate in traffic situations consistent with the DARPA Technical Criteria. The development of an optimization-based, dynamic planner has proceeded more slowly than expected and has held back our progress in accomplishing basic traffic and advanced navigation tasks. System level tests have been used to identify additional areas for continued work over the remaining days until the competition.

A key element of our system is the use of a networked control systems (NCS) architecture that we developed in the first two grand challenge competitions. Building on the open source *Spread* group communications protocol, we have developed a modular software architecture that provides inter-computer communications between sets of linked processes [5]. This approach allows the use of significant amounts of distributed computing for sensor processing and optimization-based planning, as well as providing a very flexible backbone for building autonomous systems and fault tolerant computing systems. This architecture also allows us to include new components in a flexible way, including modules that make use of planning and sensing modules from the Jet Propulsion Laboratory (JPL) and the OTGX software from Northrop Grumman, described in more detail below.

A schematic of the high-level system architecture that we are developing for the Urban Challenge is shown in Figure 3. This architecture shares the same underlying approach as the software used for the 2005 Grand Challenge, but with three new elements:

Canonical Software Architecture for mission and contingency management. The complexity and dynamic nature of the urban driving problem make centralized goal and contingency management impractical. For the navigation functions of our system, we have developed a decentralized approach where each module only communicates with the modules directly above and below it in the hierarchy. Each module is capable of handling the faults in its own domain, and anything the module is unable to handle would be propagated "up the chain" until the correct level had been reached to resolve the fault or conflict. This architecture builds on previous work at JPL [7, 8, 11].

Mapping and Situational Awareness. The sensing subsystem is responsible for maintaining both a detailed geometric model of the vehicle's environment, as well as a higher level representation of the environment around the vehicle, including knowledge of moving obstacles and road features. It associates sensed data with prior information and broadcasts the structure and uncertainty in the environment to the navigation subsystem. The mapping module maintains a vectorized

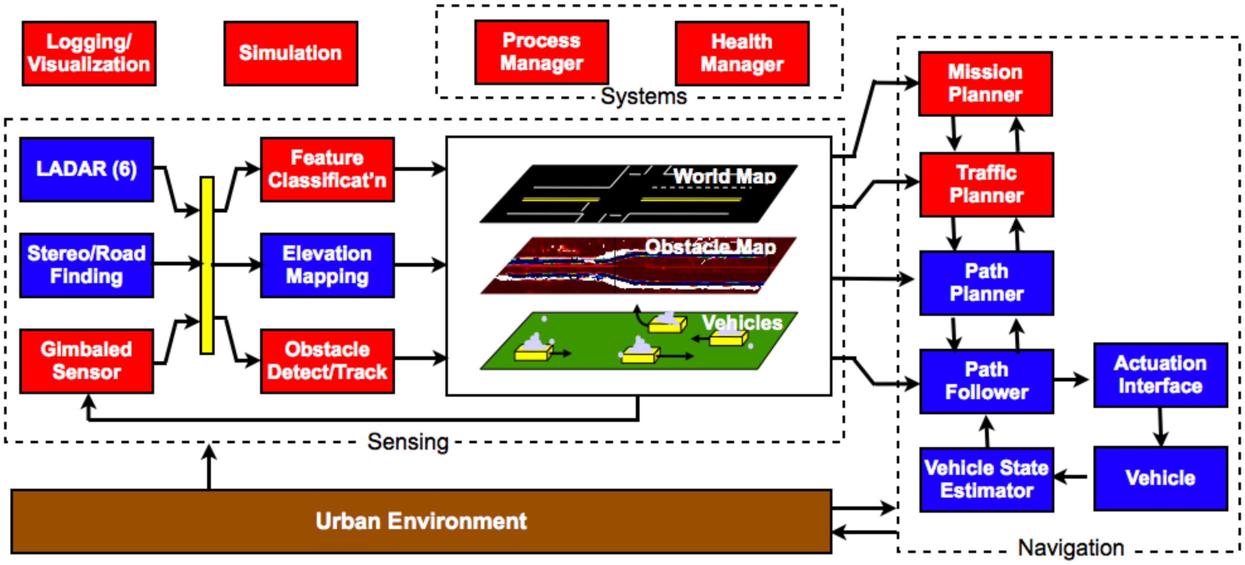

Figure 3: Systems architecture for operation of Alice in the 2007 Challenge. The sensing subsystem is responsible for building a representation of the local environment and passing this to the navigation subsystems, which computes and commands the motion of the vehicle. Additional functionality is provided for process and health management, along with data logging and simulation. The modules in blue (dark shading) were present in the 2005 architecture, the modules in red (lighter shading) are new modules that are currently being tested. Dashed boxes indicate functional subsystems and align with organizational teams.

representation of static and dynamic sensed obstacles, as well as detected lane lines, stop lines and waypoints. The map uses a 2.5 dimensional representation where the world is projected into a flat 2D plane, but individual elements may have some non-zero height. Each sensed element is tracked over time and when multiple sensors overlap in field of view, the elements will be fused to improve robustness to false positives as well as overall accuracy.

Route, Traffic and Path Planning. The planning subsystem determines desired motion of the system, taking into account the current route network and mission goals, traffic patterns and driving rules, and terrain features (including static obstacles). This subsystem is also responsible for predicting motion of moving obstacles, based on models of driving behavior and traffic rules, and for implementing defensive driving techniques. The planning problem is divided into three sub-problems (route, traffic, and path planning) and implemented in separate modules. This decomposition is well-suited to implementation by a large development team and modules can be developed and tested using earlier revisions of the code base as well as using simulation environments described in more detail below. This approach also allows easy access to the different layers of environment representation that are needed by different planning modules.

3.2 Fault-tolerant navigation and sensing (Julia Braman, NSF fellow)

Alice exited the 2005 DARPA Grand Challenge when it bounced over a concrete K-rail toward a group of journalists covering the race. Post-race analysis revealed that the system had been successfully operating with two of four LADAR units failed, but had failed to fully recover from a GPS dropout that occurred after passing under a power line, resulting in a misplaced confidence in uncertain heading knowledge [5]. Alice's premature exit from the race was partly the result of a failure of the control system to maintain adequate knowledge of its surroundings after a sensor failure. This example illustrates a significant problem associated with the design of large control systems, commonly referred to as "sensor fusion": the availability of large amounts of raw data from which a few key, and often complex, states must be determined. The Mission Data System (MDS), developed at JPL, provides an architectural solution to this problem through the use of knowledge goals, explicit constraints on the accuracy and precision of state knowledge required to achieve control objectives within a mission context. Specifically, this approach would enable a control system to deterministically manage the trade-offs between what it can accomplish given the knowledge is has available.

In recent work, ew have development of a method for converting a goal network control program into a hybrid system is given and a process for converting logic associated with the goal network into transition conditions for the hybrid automata is developed. The resulting hybrid system can then be verified for safety in the presence of failures using existing symbolic model checkers. An example task and goal network is designed, converted to hybrid automata, and verified using symbolic model checking software for hybrid systems.

3.3 Prediction & Probabilistic Planning on Autonomous Vehicles (Pete Trautman)

The point of this research is to explore the process of fully integrating sensing, planning and prediction on autonomous vehiclesthat is, we develop technology to reason over probability distribution function's (pdf s), rather than the statistics associated with those pdf's. Of course, any autonomous machine has to eventually make a decision (at which point we have to extract statistics, such as the mean and covariance, or the maximum a-posteriori) but the following methodology delays statistic extraction until absolutely necessary (e.g., at the PID controller level). We develop the technology both theoretically and experimentally. As a testbed, we utilize Caltechs DARPA Urban Challenge entry, Alice. The research explores three areas: (1) where we think the prediction component can best insert in Alices current software architecture; (2) the theory behind the predictive planner, which includes: a posterior over possible tra jectories, conditioned on the current pdf of the map, the predictive models for both the map (including dynamic obstacles) and Alice, and the cost (or weight) function which applies values to each point in the map; and (3) An implementation of the above theory, utilizing a particle filter based approach. We illustrate the process in the Figure 4 with a solution trajectory for the simple case of Alice passing another slower moving vehicle; in particular, all quantities in the map are corrupted with noise: the road boundaries, the vehicle in front of Alice, and Alice herself are all represented probabilistically. This forces the algorithm to weigh pdf's against pdf's, rather than weighing statistics against statistics.

4 Related Activities

Connections II: Mathematical Foundations of Network Science We have recently organized a workshop on the theory of network science, motiviated by applications in a broad cross-section of disciplines.

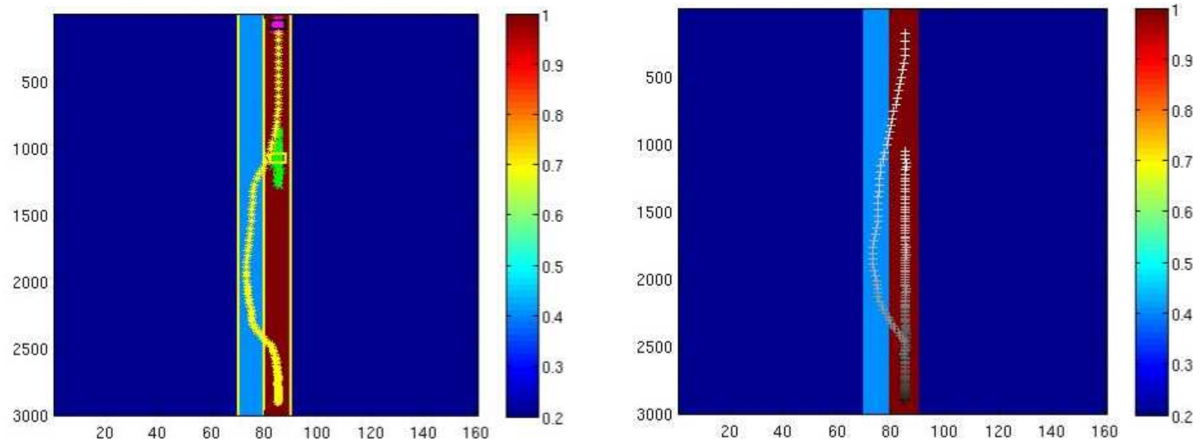

Figure 4: (a) Yellow asterisks = the mean of the particle distribution of the tra jectory of Alice at each time step; purple dots are Alice particles at simulations end; green dots are the dynamic obstacle particles at simulations end; black and yellow boxes surround the mean of the Alice and dy- namic obstacle particle distributions, re- spectively. (b) Crosses represent tra jectories of Alice (passing vehicle) and the dynamic obsta- cle; crosses fade from black to white; same colors represent same times in the simu- lation (i.e. white Alice cross corresponds temporally with a white dynamic obstacle cross).

MURI on V&V for Distributed Embedded Systems We have recently been awarded a Multidisciplinary University Research Initiative (MURI) award to investigate the specification, design and verification of distributed systems that combine communications, computation and control in dynamic, uncertain and adversarial environments. These systems consist of autonomous components (vehicles, sensors, communications nodes and command and control elements) that cooperate with each other and operate in environments with adversarial and random elements.

5 New Activities (Michael Cross)

We have recently proposed to study novel approaches to multiscale analysis of complex systems, based on methodologies and techniques developed for physical dynamical systems, and investigate the role of software in complex systems by exploring dynamical systems whose interconnection structure can be programmed. To develop the required methodologies and techniques we will study an archetype complex system which may contain a multitude of scales, namely collections of disparate coupled nonlinear oscillators where each oscillator may have its own natural frequency. These ensembles display a rich set of collective behaviors. We will use renormalization group methods inspired by the physics of phase transitions, and extend these methods going beyond the standard concepts of self-similarity. Many of the tools of dynamical systems and control can yield insights into the dynamics of software systems (and software controlled systems). Specifically, the dynamics of information flow in multi-agent/multi-component agreement protocols shares many characteristics with problems in coupled oscillators related to the role of the information topology in stability and robustness of the system. The collaboration between Engineering and Physics will work to understand the potential for appropriately integrating this line of research into the multiscale analysis of physical systems.

References

[1] K. M. Bobba. *Robust flow stability: Theory, computations and experiments in near wall turbulence*. PhD thesis, California Institute of Technology, 2004.

[2] F. Borrelli and T. Keviczky. Distributed LQR design for identical dynamically decoupled systems. *IEEE Trans. Automatic Control*, 2007, submitted.

[3] L. Chen, S. H. Low, and J. C. Doyle. Random access game and medium access control design. 2006.

[4] L. Chen, S. H. Low, and J. C. Doyle. Contention control: A game-theoretic approach. 2007.

[5] L. B. Cremean, T. B. Foote, J. H. Gillula, G. H. Hines, D. Kogan, K. L. Kriechbaum, J. C. Lamb, J. Leibs, L. Lindzey, C. E. Rasmussen, A. D. Stewart, J. W. Burdick, and R. M. Murray. Alice: An information-rich autonomous vehicle for high-speed desert navigation. *Journal of Field Robotics*, 2006. Submitted.

[6] W. B. Dunbar and R. M. Murray. Distributed receding horizon control for multi-vehicle formation stabilization. *Automatica*, 42(4):549–558, April 2006.

[7] D. Dvorak, R. D. Rasmussen, G. Reeves, and A. Sacks. Software architecture themes in jpl's mission data system. In *Proceedings of 2000 IEEE Aerospace Conference*, 2000.

[8] M. Ingham, R. Rasmussen, M. Bennett, and A. Moncada. Engineering complex embedded systems with state analysis and the mission data system. *J. Aerospace Computing, Information and Communication*, 2, 2005.

[9] T. Keviczky, F. Borrelli, and G. J. Balas. Decentralized receding horizon control for large scale dynamically decoupled systems. *Automatica*, 42(12):2105–2115, December 2006.

[10] T. Keviczky, F. Borrelli, and G. J. Balas. Distributed predictive control: Synthesis, stability and feasibility. In J. Shamma, editor, *Cooperative Control of Distributed Multi-Agent Systems.* John Wiley & Sons, 2007. In press.

[11] R. D. Rasmussen. Goal based fault tolerance for space systems using the mission data system. In *Proceedings of the 2001 IEEE Aerospace Conference*, 2001.

[12] L. N. Trefethen, A. E. Trefethen, S. C. Reddy, and T. A. Driscoll. Hydrodynamic stability without eigenvalues. *Science*, 261, 1993.