Model-based Design and Qualification of Complex Systems

John C. Doyle Richard M. Murray
Control and Dynamical Systems
California Institute of Technology

March 2006

This report summarizes the work supported by Boeing over the period of January 2005 to present.

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. One signature of RYF systems is power law statistics in event sizes. Power laws are ubiquitous in natural and human systems, and are heavily studied, yet remain a source of tremendous confusion in the scientific literature. Part of this work is to resolve this confusion and broadly educate the technical community about the relevance and rigorous treatment of power law statistics.

2. *Multi-scale Modeling* – Systems modeling theory and practice with emphasis on multiresolution modeling, and managing multiple distinct product representation that must be mapped to each other. One critical issue involving computational intractability is connecting multiple scales in models. One familiar example is that in many manufacturing processes the final product's macroscopic functional characteristics are determined by microscopic material properties, which are in turn determined by macroscopic process parameters. This passage of macro to micro to macro makes the prediction of final systems properties from the process design extremely difficult. Another critical issue is that in order for humans to understand and design complex systems, artificial and somewhat arbitrary decompositions (from a physical point of view) must be introduced so that different aspects of the design may be addressed.

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"). The MVWT consists of a collection of up to 24 robotic vehicles with onboard sensing (IR, vision, sonar), actuation and computation that perform cooperative tasks in a distributed, networked environment. Algorithms for

cooperative control will be designed and then analyzed using model-based tools. The Caltech autonomous vehicle testbed consists of a Ford E-350 outfitted for offroad racing as part of the DARPA Grand Challenge. Autonomous motion is performed using a collection of sophisticated perception and decision-making algorithms, all running on a high speed cluster of computers that form an advanced systems of systems architecture. Correct operation in highly uncertain environments with fault tolerance and reconfiguration will be analyzed by breaking out subproblems that are amenable to analysis through SoS, RYF and other tools. Both systems of systems are being modeled using public domain tools such as Gazebo and ODE.

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

2.1 Design Principles for Systems of Systems Engineering (David Alderson)

Modern engineering systems are the simultaneous victims of conflicting trends in requirements and capabilities. On the one hand, the scope of what goes into system design is now enormousas evidenced by the growing list of system specifications, typified by "x-ities": flexibility, scalability, evolvability, sustainability, extensibility, etc. On the other hand, the way in which large numbers of diverse components are integrated into a functioning whole with robust collective behavior is often viewed as more art than science. All too frequently, we are reminded that our ability to mass produce network-enabled devices guarantees little about their collective behavior when deployed.

The existing Internet provides an excellent case study in complex networks. A principle challenge in the study of many complex systems is to understand the relationship between system structure and large-scale system function, as illustrated the left panel of Figure 1. Despite general familiarity with the Internet, there has been considerable confusion within the broad scientific community about some of the Internets most basic features, including its large-scale topological structure and the corresponding implications for network robustness in the presence of router failures. Because the Internet evolved in a largely ad hoc manner without central planning or design, it is sometimes perceived that its most important robustness features are the result emergent self-organization. Yet our work over the past year suggests that what often appears as "self-organized emergence" can be explained by well-conceived (albeit perhaps heuristic) design, with explanations that are mathematically rigorous, in agreement with engineering reality, and fully consistent with network measurements. With John Doyle, Lun Li, and Walter Willinger (AT&T LabsResearch) I have continued to develop a methodological approach to reverse-engineering system structure using an optimization-based framework, and we have used it along with a first principles understanding of

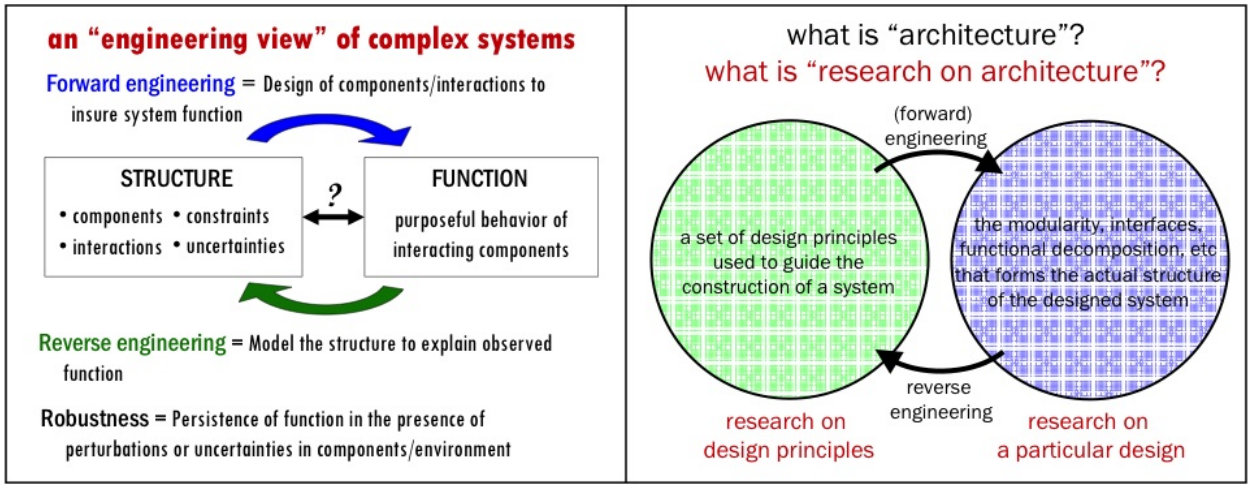

Figure 1: **Dave: please send me a short caption**

network design to explain key structural features of the current Internet. This work has shown that perhaps the most popular notion of complex network vulnerability (i.e., the "Achilles' heel" of scale-free networks) is based on faulty reasoning and does not exist for the Internet. Furthermore, by reconciling the differences in the assumptions and methodologies of different modeling approaches, our work has helped to resolve much of the confusion and controversy that has surrounded network modeling and evaluation.

A related topic of growing interest within the engineering systems community is What is architecture? In the context of the Internet, we have shown that one may conduct research on architecture at two levels (right figure). At an abstract level, one is concerned with architecture as a set of guiding design principles that can be used to do forward engineering. At a practical level, the study of an existing architecture (in our case, the wired Internet) serves to inform how the particular implementation of design principles can be used to solve specific engineering problems. Clarifying the role of architecture is increasingly important in the design of next-generation networks, including the newly launched NSF Future Internet Network Design (FIND) Initiative that looks to use a "clean slate" design approach to rebuild the Internet from scratch.

2.2 Optimization-Based Methods for System Verification (Stephen Prajna)

A body of techniques based on convex optimization and sum of squares programming has been developed for verification of a large class of dynamical systems, including those with nonlinear dynamics, uncertainties, hybrid (mixed discrete-continuous) dynamics, stochasticity, and time-delay [11, 13, 16, 15, 12]. These techniques verify temporal properties such as safety (something bad never happens), reachability (something good can happen), eventuality (something good will surely happen), and their simple combinations, using certain functions of states called barrier certificates and density functions [11, 16].

For a simple illustration, consider a continuous system $\dot{x} = f(x, d)$ where x is the state of the system taking its value in the state space $\mathcal{X}$ and d is a disturbance input taking its value in $\mathcal{D}$. In addition, consider $\mathcal{X}_0 \subset \mathcal{X}$ as the set of possible initial states, and $\mathcal{X}_u \subset \mathcal{X}$ as the set of unsafe

states. Suppose there exists a barrier certificate, i.e., a differentiable function $B : \mathcal{X} \to \mathbb{R}$ satisfying

$$B(x) \leq 0 \quad \forall x \in \mathcal{X}_0, \tag{1}$$

$$B(x) > 0 \quad \forall x \in \mathcal{X}_u, \tag{2}$$

$$\frac{\partial B}{\partial x}(x) f(x, d) \leq 0 \quad \forall x \in \mathcal{X} \times \mathcal{D}. \tag{3}$$

Then it is easy to see that the safety property holds, i.e., that for all possible initial state $x_0 \in \mathcal{X}_0$ and for all possible disturbance input there exists no trajectory of the system that goes from the initial set to the unsafe set. Systems with hybrid dynamics can be treated in an analogous manner, by asking that during the discrete transition the value of $B(x)$ also satisfies certain non-increasing conditions similar to (3).

It is obvious that simulation is of limited use to address the verification of safety property stated above. Since the state space of the system is uncountable, verifying by simulation that the property holds in all cases is never exact, simply because it is impossible to test all system behaviors. In fact, simulation alone may fail to uncover the existence of bad behaviors. Using barrier certificates and density functions to prove safety, reachability, and eventuality is analogous to using Lyapunov functions to prove stability. It eliminates the needs to run simulations, to explicitly compute the flow of the system, or to propagate sets of states.

For stochastic systems, such as those described by stochastic differential equations, safety verification can also be handled by computing an appropriate barrier certificate which upper-bounds the probability of reaching the unsafe set [13]. In this case, a barrier certificate $B : \mathcal{X} \to \mathbb{R}$ which generates a stochastic process $b(t) := B(x(t))$ that is a supermartingale, i.e., whose evolution along time is non-increasing on the *average*, is used. We also ask that the value of the barrier certificate at the initial states be lower than its value at the unsafe states. The probability of reaching the unsafe region can then bounded from above using a Chebyshev-like inequality for supermartingales.

There are other classes of systems that can be handled using this methodology. One example is given by time-delay systems. For verification of a time-delay system, a functional of states is used as a barrier certificate [12]. The forms of the functionals resemble the Lyapunov-Razumikhin functions or the Lyapunov-Krasovskii functionals used in stability analysis of time-delay systems. In [12], a hierarchy of functional structures is proposed to prove safety with decreasing levels of conservatism.

The conditions that must be satisfied by barrier certificates and density functions are formulated as convex programming problems. In addition to benefits in terms of computation, the duality structure inherent because of their formulation as convex programs also gives theoretical advantages [16, 15]. For example, the dual of safety verification, i.e., reachability verification, concerns proving the existence of a trajectory starting from the initial set that reaches another given set. Using insights from the linear programming duality appearing in the discrete shortest path problem, it is shown in [16] that reachability of continuous systems can also be verified through convex programming. Several convex programs for verifying safety, reachability, and other temporal properties are formulated. As another example, a completeness statement in safety verification using barrier certificates is obtained by exploiting the strong duality between safety verification and reachability verification [15], stating that under reasonable technical conditions, the existence of a barrier certificate satisfying (1)–(3) is both sufficient and *necessary* for safety.

For systems and sets whose descriptions are in terms of polynomials, sum of squares programming [14] provides a hierarchy of scalable algorithmic methods for computing barrier certificates

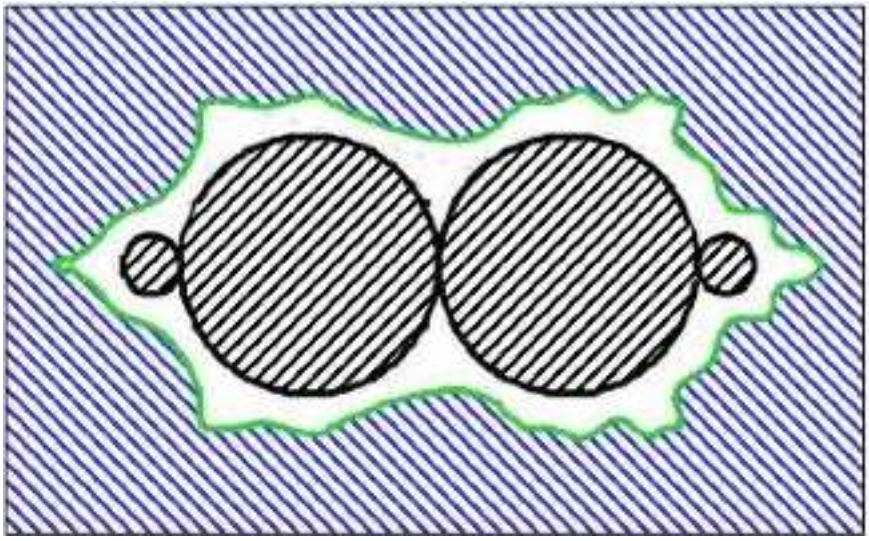

Figure 2: **Dennice:** Please send me a caption

and density functions, where at each level the computational cost grows polynomially with respect to the system size. Because of this, the methodology seems to be more scalable than many other existing methods that can handle nonlinear continuous and hybrid systems. Successful application of the method for verifying the safety property of a NASA life support system, which is a nonlinear hybrid systems with 6 discrete modes and 10 continuous states, has been reported in [10].

2.3 Invariance and Set Membership using SOS (Dennice Gayme)

The most widely used technique for determining set membership for any of the quadratic maps is to grid the space into some large number of points and exhaustively simulate the system dynamics. Then after some number of iterations the points that remain bounded are in the set. This type of simulation is limited by two main problems; first the problem growth is of the form N n , which is badly exponential in n. The more relevant problem, in regards to determination of membership for the Mandelbrot set is that as one gets near the boundary simulations can take arbitrarily long to diverge or go to a limit cycle, these cases result in infinite computational cost.

However using SOS techniques one can determine set membership for certain parameter values as well as establish that other values are outside the set in a computationally tractable manner. By bounding the set in both directions one can isolate the number of points that are indeterminate. Further the methods allow one to make definitive statements for the largest regions in the Mandelbrot set as well as to describe the initial conditions in the dynamic system, (z plane), that allow one to remain in the set. In this manner we are able to obtain a four dimensional hyperspace in z and λ that describes the region of attraction. One can also determine a polynomial equation defining the points that are outside of the set after 5 iterations. The resulting inner and outer bounds are shown in Figure 2.

Current work focuses on how different forms of Positivstellensatz can be applied to logistic map and how that affects the order of SOS coefficients needed to obtain a proof. One of the open questions regarding SOS techniques has to do with how best to formulate a problem to obtain the lowest order proof and in turn how to interpret the proof order with respect to the problem one

is analyzing. By studying the different forms of Positivstellensatz and how they affect proof order in a variety of systems I hope to gain insight into which part of the proof the constraints are best applied in minimizing the order of the proof.

We are also interested in how reduction theory in Geometric Mechanics and exploitation of problem structure (both symmetry and sparsity) in SOS tools are related. To this end we are reviewing reduction theory in mechanics and the next step would be to go through the theory and see how things change if things are restricted to the real polynomial ring. We plan to begin looking at this problem particularly applied to problems in Discrete Mechanics and Optimal Control (DMOC) and start with the examples of a satellite with rotors and the fish swimming model developed by Eva Kanso [**?**].

2.4 Data observation windows for forecasting in stochastic systems (Alfred Martinez)

For the estimation problem, it is well known that model misspecification can lead to seriously biased parameter estimates. Likewise, one can expect model misspecification to have severe consequences for the forecasting problem. For example, in the case of non-stationary processes with temporal structural breaks, using data previous to the break will increase error bias but decrease error variance. Such trade-off suggests the non-trivial nature of the problem of selecting data observation windows. The two common techniques used for data selection are a rolling window and expanding window but both are ad hoc methods with little or no basis for their application. We develop algorithms to estimate the size of the data observation window which results in optimal forecasts. Our algorithms base the determination of the optimal data window on the nature of the processes in question (stationarity, dependence structure etc). Due to the ubiquitous nature of misspecification, our work has application in engineering, the natural sciences and economics.

We analyze forecasting of a data generating process (DGP) by considering a stochastic process Z Zt : Rm+1 , m N, t = 1, . . . , n + 1 defined on a complete probability space (, F , P) where F = Ft , t = 1, . . . , n + 1 and Ft is the -field Ft Zs , s t. We denote by Yt the component of interest of the observed vector Zt , Yt R, and interpret the remaining components, denoted Wt , as being an m 1 vector of other variables. In other words, we let Zt (Yt , W t) .

The forecasting problem considered involves forecasting the variable Yt+s , where s is the prediction horizon of interest, s 1. Xt is a k 1 vector of Ft measurable variables that are used to forecast Yt+s . In practical applications, Xt can contain (1) various lags of the variable of interest Yt , (2) realizations of the other variables Wt , as well as (3) any function of the previous two. As such, our setup allows for applications involving both time-series and cross-section data.

The forecasters focus is on the mean of Yt+1 conditional on the entire information set Ft which we denote Et (Yt+1). Not knowing the true DGP, the forecaster uses a model for Et (Yt+1) which is linear in Xt , Xt , in which is a k 1 parameter vector, B, B compact in Rk . The forecasting model is misspecified whenever Et (Yt+1) is not a linear function of Xt . Depending on Et (Yt+1), misspecification of the forecasting model can be dynamic, functional, distributional or a combination of the aforementioned.

With a linear model, the forecast of Yt+1 has the form of the regression, Yt+1 = Xt + Vt+1 , with an error term Vt+1 such that Et (Vt+1) = 0. The parameter is estimated by an ordinary least squares (OLS) estimator. The OLS estimator of , t,n , has a non-trivial dependence on the observation window sizes, n, and is used to construct the forecast Yt+1,n of Yt+1 as follows Yt+1,n = t,nXt .

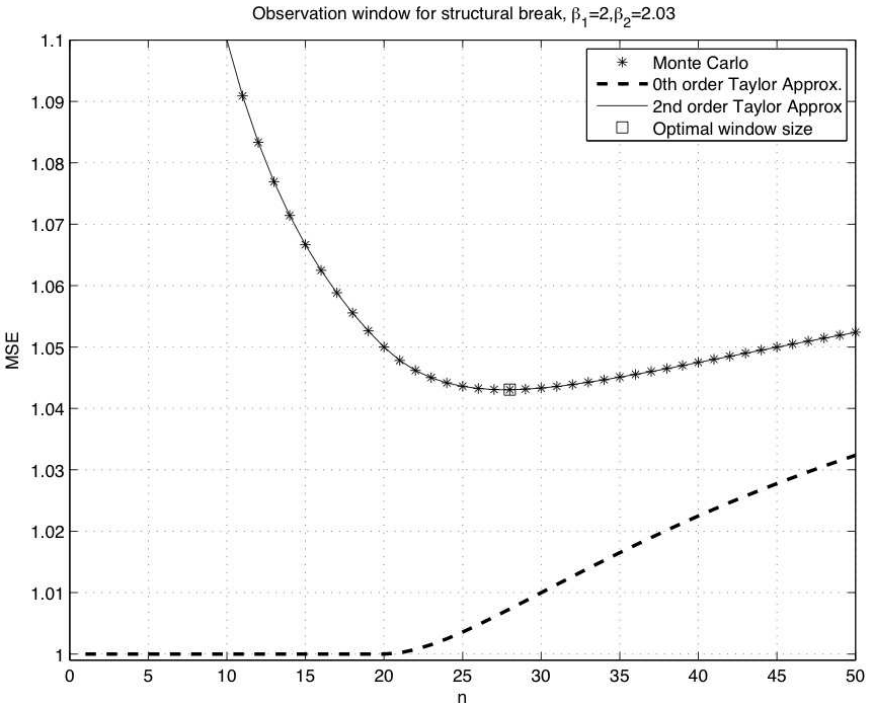

Figure 3: MSE for a linear structural break process and a linear model.

The criteria for optimality in our evaluation of the accuracy of the forecasts Yt+1,n is the mean square error (MSE), M S En = E[(Yt+1 Yt+1,n)2]. We examine the dependence of the MSE on the window size n by constructing an algorithm based on Taylor polynomials to approximate the MSE and the optimal window size n which minimizes M S En . To illustrate, the figure below presents the MSE obtained with Monte Carlo simulations for a linear DGP which undergoes a structural break at n = 20. The break involves the linear parameter which changes from 1 = 2 to 2 = 2.03. The figure also shows two approximations of the MSE obtained with our algorithm. While the zeroth order approximation performs poorly, the second order 1 approximation predicts the optimal data window size to be n = 28 which coincides with the true Monte Carlo results.

3 Communication Networks

3.1 Cross layer optimization in TCP/IP (Lun Li)

Next-generation network-enabled system will focus on control over networks, as opposed to simple control of networks. Nonetheless, classic TCP/AQM over the wired Internet serves as an ideal case study for understanding issues related to cross-layer optimization and decentralized control.

TCP/AQM can be interpreted as distributed primal dual algorithms to maximize aggregate utility over source rates. Previous work has assumed that the routing is fixed during the time of interest. We study the effect of routing changes by investigating the joint optimization of utility over both source rates and their routes. From the figure we can see that the dual of this problem suggests using shortest path routes with congestion prices as the cost. We therefore define a TCP/IP process as follows: assume routing update operates at a much lower time scale while TCP converges

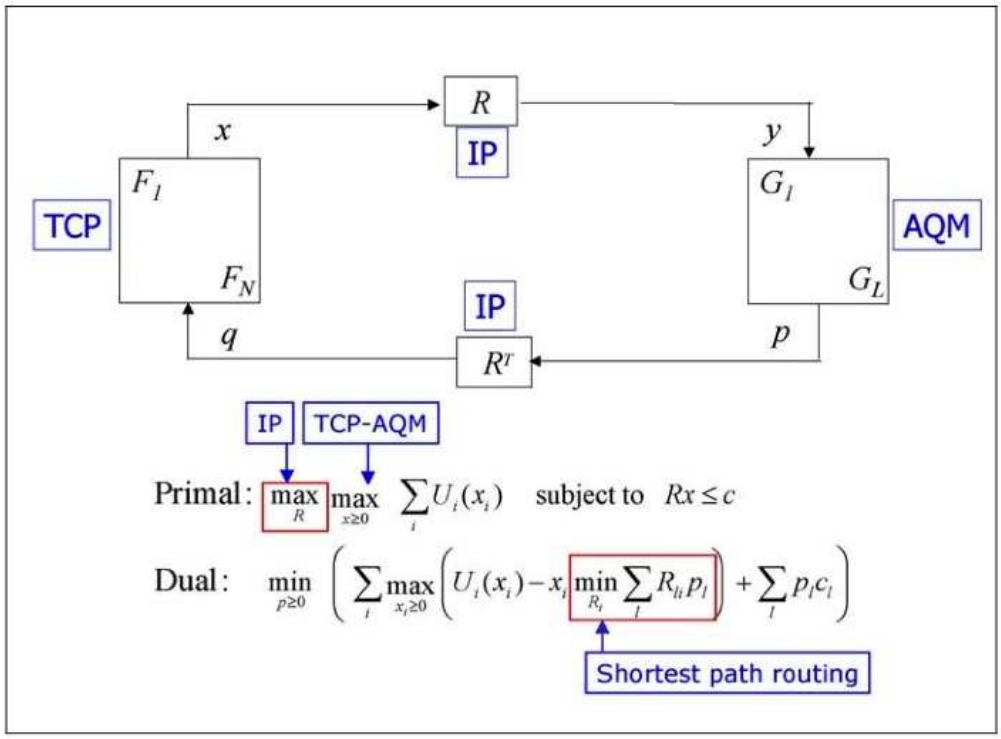

Figure 4: **Lun: please send me a good caption**

instantly and produces a congestion price. Use this price as link cost to generate the minimal cost routes for a new iteration of TCP. Thus TCP/IP form a feedback system where routing interacts with congestion control in an iterative process. We are interested in the equilibrium and stability properties of this iterative process.

Our results show that in the case of pure dynamic routing as describe above, equilibrium of TCP/IP system exists if only if there is no duality gap of the joint utility maximization problem. In this case, TCP/IP equilibrium solves both primal and dual problem. Moreover, it incurs no penalty for not splitting traffic across multiple paths: optimal single path routing achieves the same aggregate utility as optimal multipath routing. Multipath routing can achieve a strictly higher utility only when there is a duality gap between the single path optimization problems, but in this case, the TCP/IP system does not even have an equilibrium. Even when the single-path problem has no duality gap and TCP/IP has an equilibrium, the equilibrium is generally unstable but it can be stabilized by adding a sufficiently large static component to the link cost but at the expense of a reduced utility in equilibrium. With a theoretical study in a ring network and numerical simulation in randomly generated networks, we demonstrate that there is an inevitable tradeoff between the routing stability and achievable utility.

3.2 Congestion control in wireless networks (Lijun Chen)

Communication networks are complex systems consisting of intelligent units such as PCs that have computing and communicating capabilities. The functioning of the network as a whole is made possible by all kinds of protocols that integrate these individual units together. As networks

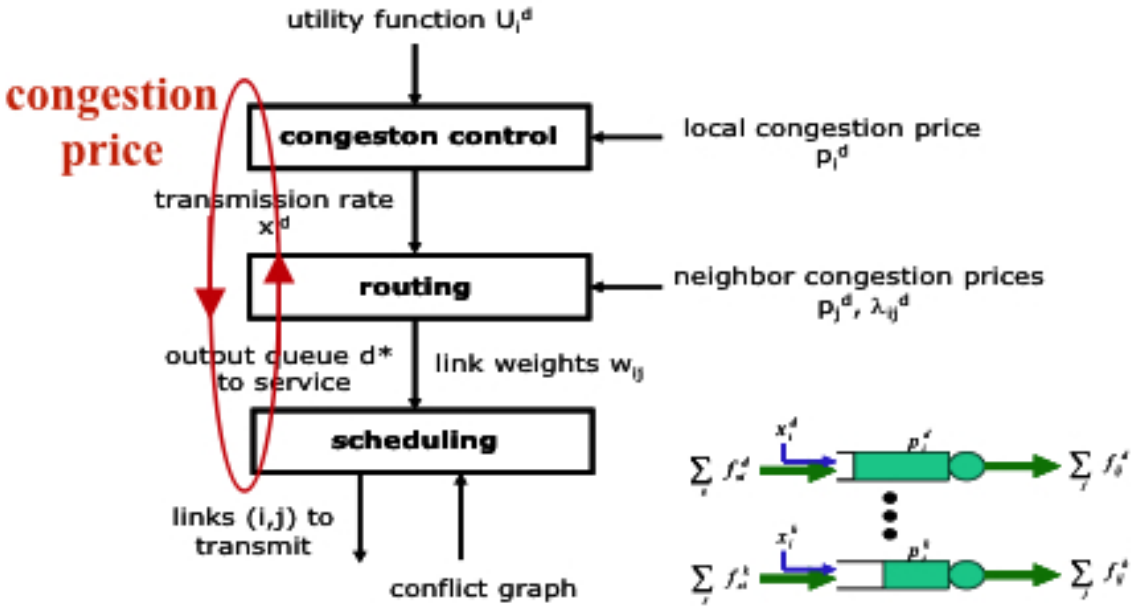

Figure 5: **Lijun: please send me a short caption**

adopt a layered structure, protocols for different layers are usually optimized and implemented separately, and then interconnected, often in ad hoc manner. Though justified by the success of many communication networks, this "layered" design methodology does not perform well for wireless networks because of time-varying channel, contention based channel access and mobility. In order for wireless networks to provide better performance, we must re-think the protocol stack as a whole, and exploit the interactions among various layers to do cross-layer design.

The approach of protocol as distributed solution to some global optimization problem through dual decomposition has been successfully applied to TCP congestion control. The key innovation from this line of work is to view network as an optimization solver and congestion control protocol as distributed algorithms solving a network utility maximization (NUM) problem. This approach has recently been substantially extended from an analytic tool of reverse-engineering TCP congestion control to a general approach to understand interactions across layers. Application needs form the objective function, i.e., network utility to be maximized, and the restrictions in the communication infrastructure are translated into many constraints of a generalized network utility maximization problem. Such problems in general may be very difficult nonlinear, nonconvex optimization with integer constraints. There are many different ways to decompose a given problem, each of which corresponds to a different layering scheme. These decomposition (i.e., layering) schemes have different trade-offs in efficiency, robustness, and asymmetry of information and control, thus some are better than others depending on the criteria set by the network users and operators.

We apply this approach to design an overall framework for the protocol architecture in ad hoc wireless networks, with the goal of achieving efficient resource allocation through cross-layer design [5, 4, 3]. As illustrated in the above picture, our current theory integrates three functions congestion control, routing, and scheduling in transport, network, and link layers into a coherent framework and makes transparent their interactions not only vertically across the protocol stack,

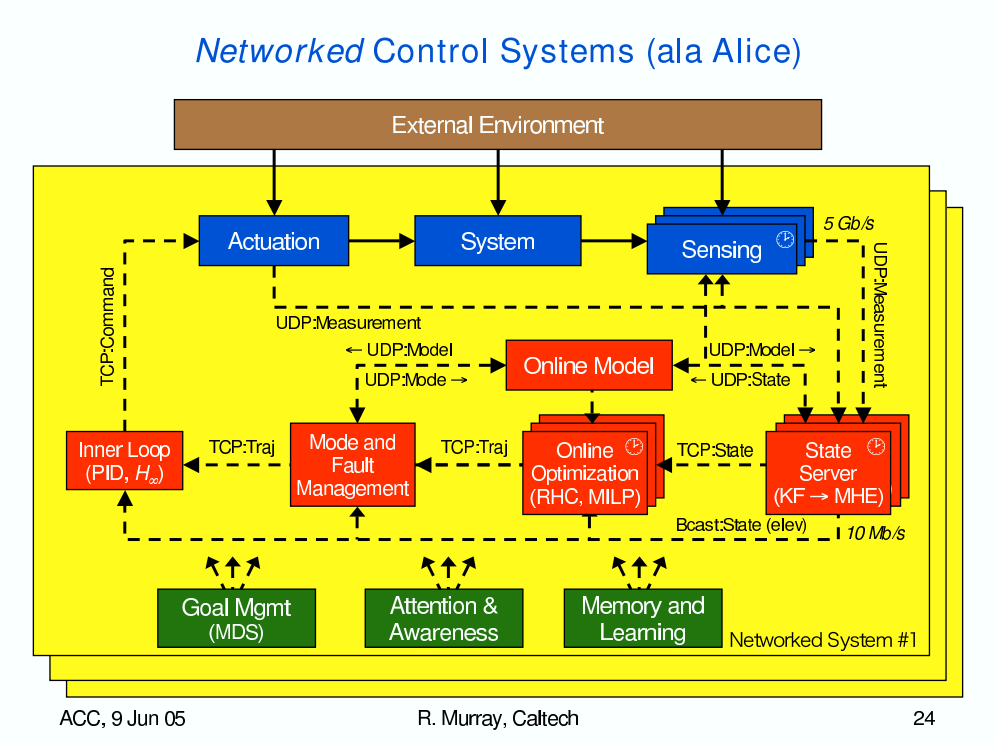

but also horizontally across multiple network nodes. These three functions interact through and are regulated by congestion price so as to achieve a global optimality, even in a time-varying environment. The structural simplicity of the underlying optimization problem also leads to simple and robust equilibrium and dynamic behaviors. Even though this framework does not provide all the design and implementation details, it helps us understand issues, clarify ideas, and suggests directions, leading to better and more robust designs for ad hoc wireless networks.

4 Autonomous Vehicles

4.1 Gazebo models for Alice and MVWT (SURF)

Gazebo is an open source, multi-robot simulator for outdoor environments. It is capable of simulating a population of robots, sensors and objects in a three-dimensional world. It generates both realistic sensor feedback and physically plausible interactions between objects, including accurate simulation of rigid-body physics. Caltech is using Gazebo as its primary modeling environment for Alice and for the MVWT. In both cases, Gazebo was chosen because of its ability to include vision and LADAR sensors, a feature which many other simulators (eg, SIMULINK and Modelica) lack.

The simulator for Alice includes

The simulator for the multi-vehicle wireless testbed is still under development, but is being used to replace a customized simulation that had been used previously (as part of Cornell's RoboFlag simulation).

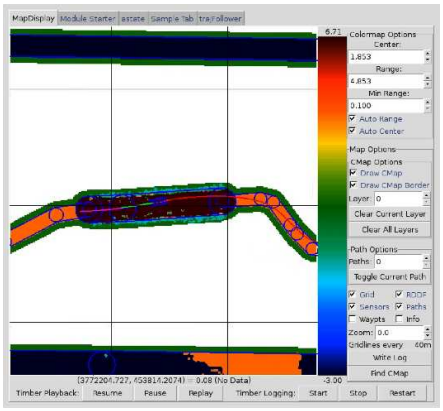

Figure 6: Caltech's autonomous vehicle, Alice. The left figure shows the vehicle as configured for the 2005 DARPA Grand Challenge. The right figure shows the graphical user interface. The colors represent allowable speeds and the magenta line is the currently planned path.

4.2 Real-time trajectory generation (Dima Kogan)

The DARPA Grand Challenge (DGC) was an off-road autonomous vehicle race in the Mo jave desert. One of the open problems involved in the design and development of an entry to the DGC is the navigation of the vehicle. The vehicle has to be able to drive autonomously in unstructured and previously-unknown terrain, sensing its environment as it moves. This research addresses the planning problem with a non- linear optimization method running in real time. The vehicles on-board computers continually solve an optimization problem to find a time-optimal, dynamically feasible tra jectory from the vehicles position to some receding horizon ahead (20m-70m forward). The optimization is performed in two stages, one seeding the other. First, a rough, globally optimal spatial path is found by evaluating sets of piecewise linear curves through the map. Then the locally optimal nonlinear optimizer is run, optimizing both the spatial and temporal components of the trajectory simultaneously.

This method has been implemented and tested on a modified Ford E350 van. Using four LIDAR units as terrain sensors, the vehicle was able to consistently traverse a 2 mile obstacle course at the DGC qualifying event. At the main DGC event, the vehicle drove 8 autonomous miles through the Nevada desert before experiencing issues with its state estimator.

During the race, the vehicles 2.2 GHz Opteron CPU produced 4.28 plans/second on average. During one of the qualification runs, 5713 successful planning computation were completed during an 1187 second run, for an average rate of 4.8 plans/second. There were 47 planning attempts that did not converge in the time allotted, for a success rate of 99.18%. Of those 47 failed attempts, 22 occurred in a rapid succession in a situation made infeasible due to state drift.

Shown above are non-consecutive planner iterations, illustrating the vehicle avoiding a parked car during a qualifying run. The vehicle is at the West (left) edge of each snapshot, and is traveling East (to the right). The result of the first planning stage is shown in green and the final solution in blue. Grayscale represents terrain speed limits (obstacles are represented by a very slow region). Red represents regions of no-data. The speed profile of each final plan is shown below each spatial snapshot. Cars south and straight ahead of the vehicle are visible, along with the sensor shadow of no-data for the vehicle straight ahead. To keep up the vehicle speed, distant no-data cells are

treated favorably, which can be clearly seen in the snapshots. Additionally, the snapshots clearly show a newly detected second car and new plans to avoid it.

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

Caltechs entry in the 2005 DARPA Grand Challenge race, named Alice, exited the race 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 [6]. 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.

We are currently using MDS to develop and evaluate supervisory controllers for systems such as Alice. Our goal is to provide rigorous analysis of goal networks that allow us to reason about the possible failure modes and then redesign the supervisory control logic to better manage goals and handle contingencies. We plan to apply and extend the MDS methodology to design "adaptive" estimation algorithms and control mechanisms for reconfiguring the system in response to specified knowledge goals, and variations in the quality of available knowledge due to sensor loss, or degradation. Alice provides an excellent platform for developing and demonstrating this type of robust sensor fusion capability because it has a large number of sensors, each of which have their own unique abilities and weaknesses.

5 Future Activities

Autonomous driving in dynamic environments

Connections II: Mathematical Foundations of Network Science

MURI on V&V for Distributed Embedded Systems

References

[1] D. Alderson, L. Li, W. Willinger, and J. C. Doyle. Understanding internet topology: Principles, models, and validation. *IEEE/ACM Transactions on Networking*, 13(6), 2005.

[2] D. Alderson and W. Willinger. A contrasting look at self-organization in the Internet and next-generation communication networks. *IEEE Communications Magazine*, July 2005.

[3] L. Chen, T. Ho, S. H. Low, and J. C. Doyle. Cross-layer rate control and scheduling in wireless networks with network coding. Technical report, In preparation, 2006.

[4] L. Chen, S. H. Low, M. Chiang, and J. C. Doyle. Cross-layer congestion control, routing and scheduling design in ad hoc wireless networks. In *Proceedings of IEEE Infocom*, 2006.

[5] L. Chen, S. H. Low, and J. C. Doyle. Joint congestion control and media access control design for wireless ad hoc networks. In *Proceedings of IEEE Infocom*, 2005.

[6] 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.

[7] J. C. Doyle, D. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger. The "robust yet fragile" nature of the internet. *Proceedings of the National Academy of Sciences*, 102(41):14497–14502, 2005.

[8] D. M. Gayme, J. C Doyle, S. Prajna, A. Papachristodoulou, and Maryam Fazel. Optimization based methods for determining basins of attraction in the logistic map and set membership in the mandelbrot set. Preprint, 2006.

[9] D. M. Gayme, M. Fazel, and J. C. Doyle. Sos proofs of invariant regions in the logistic map. In *Proc. IEEE Control and Decision Conference*, 2006. Submitted.

[10] S. Glavaski, A. Papachristodoulou, and K. Ariyur. Safety verification of controlled advanced life support system using barrier certificates. In *Hybrid Systems: Computation and Control*, 2005.

[11] S. Prajna and A. Jadbabaie. Safety verification of hybrid systems using barrier certificates. In *Hybrid Systems: Computation and Control*, 2004.

[12] S. Prajna and A. Jadbabaie. Methods for safety verification of time-delay systems. In *Proceedings of the IEEE Conference on Decision and Control*, 2005.

[13] S. Prajna, A. Jadbabaie, and G. J. Pappas. Stochastic safety verification using barrier certificates. In *Proceedings of the IEEE Conference on Decision and Control*, 2004.

[14] S. Prajna, A. Papachristodoulou, P. J. Seiler, and P. A. Parrilo. SOSTOOLS – Sum of Squares Optimization Toolbox, User's Guide. Available at `http://www.cds.caltech.edu/sostools` and `http://www.mit.edu/ parrilo/sostools`, 2002, 2004.

[15] S. Prajna and A. Rantzer. On the necessity of barrier certificates. In *Proceedings of the IFAC World Congress*, 2005.

[16] S. Prajna and A. Rantzer. Primal-dual tests for safety and reachability. In *Hybrid Systems: Computation and Control*. Springer-Verlag, 2005.

[17] J. Wang, L. Li, S. H. Low, and J. C. Doyle. Cross-layer optimization in tcp/ip networks. *IEEE/ACM Transactions on Networking*, 13(3), 2006.