C O M P L E X   S Y S T E M S

An ITL Program

Networks intellectual Technology competencies capabilities predictive program based virtually community non-existent fostering build fundamental metrics Systems Complex Network engineered sophisticated framework experts Science move provides creation solutions need core

ACTIVITIES REPORT

FISCAL YEARS 2007-2008

complex systems

Information Technology Laboratory

NISTIR - 7569

NIST

National Institute of Standards and Technology

U.S. Department of Commerce

Complex Systems Program Activities Report Fiscal Years 2008-2009

Sandy Ressler

Information Technology Laboratory

Table of Contents

Abstract

This report summarizes the work of the Complex Systems Program within NIST's Information Technology Laboratory. The Overview presents some background material and a high-level overview of the Program's activities including significant technical accomplishments. The Projects chapter provides short descriptions of each of the projects which comprise the program. Programmatic Activities describes the activities taking place to manage the Program and enable inter-project interactions. Finally the Goals, Plans & Strategy Chapter presents strategic goals and objectives of the program.

For further information contact Sandy Ressler, NIST, 100 Bureau Drive, STOP 8940, Gaithersburg, MD 20899; phone (301) 975-3549, email: sressler@nist.gov or see the Program web site at: http://www.itl.nist.gov/ITLPrograms/ComplexSystems/.
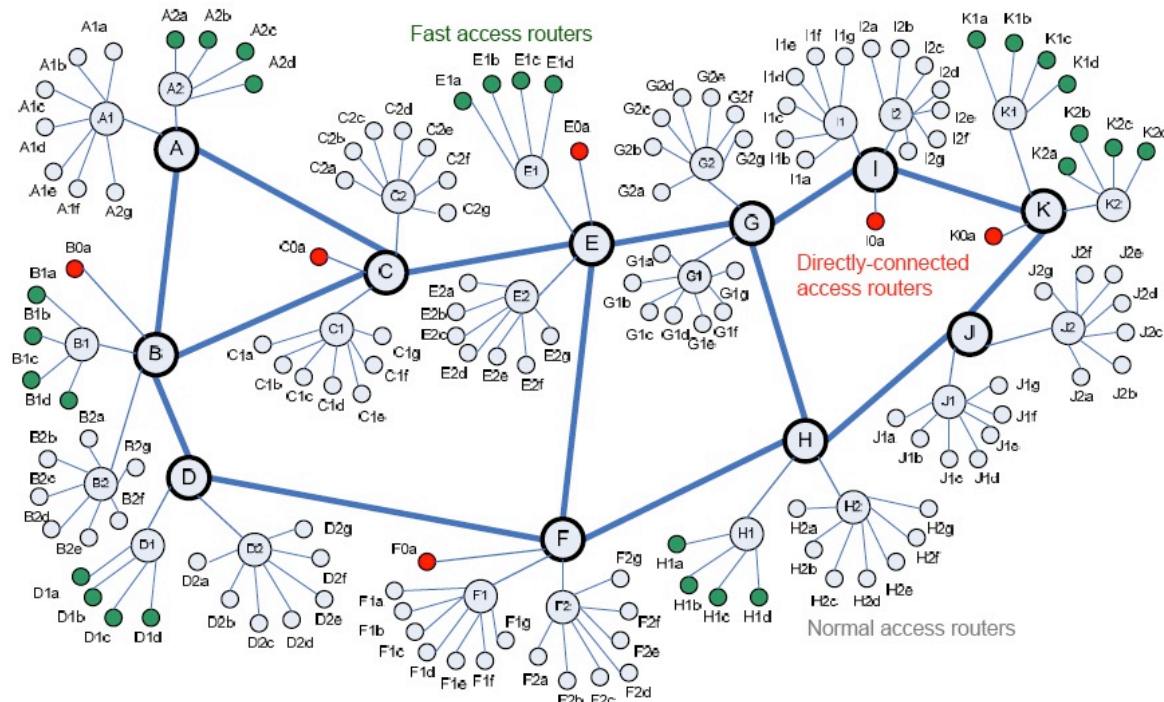
Network Toplogy Used in Congestion Control Simulation Studies in the Measurement Science for Complex Systems

Overview

"Every now and then I like to lean out my window, look up and smile for a satellite picture." - Steven Wright

Complex Systems are composed of large interrelated, interacting entities which taken together, exhibit a macroscopic behavior which is not predictable by examination of the individual entities. The Complex Systems Program seeks to understand the fundamental science of these systems and develop rigorous descriptions (analytic, statistical, or semantic) that enable prediction and control of their behavior. Initially focused on the Internet and Grid Computing, this Program will facilitate predictability and reliability in these areas. In the future we hope to apply these complex systems capabilities to other domains such as biotechnology, nanotechnology, semiconductors, and engineering which we would work with other laboratories at NIST.

The Complex Systems Program is one of eight programs in the Information Technology Laboratory. These programs represent the strategic direction for research and activities of the lab. In brief the other programs are:

Cyber and Network Security

Cybersecurity is focused on ensuring three IT system objectives: confidentiality, integrity, and availability. This Program addresses long-term scientific issues in some of the building blocks of IT security - cryptography, security testing and evaluation, security metrics, and security properties - providing a more scientific foundation for computer and network security, while maintaining a focus on near-term issues.

Enabling Scientific Discovery

Modern scientific research has become more and more dependent on mathematical, statistical, and computational tools. This Program promotes the use of these tools to dramatically advance our ability to predict the behavior of a broad range of complex scientific and engineering systems and processes. In particular, this Program focuses on inter-disciplinary projects that involve novel computational statistics and simulation methods and software.

Identity Management Systems

Identity management systems are responsible for the creation, use, and termination of electronic identities used to access logical and physical resources. This Program is pursuing the development of common models and metrics for identity management, critical standards, and interoperability of electronic identities. These efforts will improve the quality, usability, and consistency of identity management systems while protecting privacy.

Information Discovery, Use, and Sharing

This Program fosters innovation throughout the information life cycle by developing the measurement infrastructure to enhance knowledge discovery, information exchange, and information usability. The Program enables novel computational approaches to data collection and analysis to be combined with improved interoperability techniques to effectively extract needed information from the wealth of available data.

Pervasive Information Technologies

Pervasive information technology is the trend towards increasingly ubiquitous connected computing sensors, devices, and networks that monitor and respond transparently to human needs. This Program facilitates the creation of standards for sensor communication, networking interoperability, and sensor information security.

Trustworthy Information Systems

Trustworthy information systems perform as intended for a specific purpose, with operational resiliency and without unwanted side effects or exploitable vulnerabilities. This Program will improve the ability to model, produce, measure, and assess trustworthiness through innovative technologies, models, measurements, and software tools.

Virtual Measurement Systems

A virtual measurement is a quantitative result and its uncertainty, obtained by a computer simulation or computer-assisted measurements. This Program introduces metrology constructs, standard references, uncertainty characterization, and traceability into scientific computation and computer-assisted measurements to enable predictive computing with quantified reliability.

Context and National Needs

One important question we must answer is who cares about this "complex systems" work? In the context of IT networks there is a nearly 30 year history of unexpected failures:

- ARPAnet congestion collapse, 1980
- Internet congestion collapse, Oct 1986
- Cascading failure of AT&T long-distance network, Jan 1990
- Collapse of AT&T frame-relay network, April 1998
- Slammer Worm near collapse of several countries Internet access (Korea), Jan 2003
- Mass use of backup routes, 4000 Cisco routers cause 7hr downtime in NE Japan, May 2007

As society becomes increasingly dependent on IT networks these types of failures become increasingly expensive to recover from, and prediction is increasingly valuable.

No one understands how to measure, predict or control macroscopic behavior in complex information systems

"[Despite] society's profound dependence on networks, fundamental knowledge about them is primitive. [G]lobal communication … networks have quite advanced technological implementations but their behavior under stress still cannot be predicted reliably."(pp. 1)

"There is no science today that offers the fundamental knowledge necessary to design large, complex networks in such a way that their behaviors can be predicted prior to building them." (pp. 3)

— Network Science, National Research Council, 2006
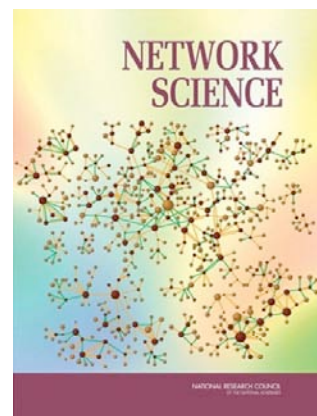
Network problems remain such that cascading failures of the power grid, Internet, phone and cable systems continue to occur. While the Internet was formed organically we really don't understand its topology, therefore management is problematic and the impact of new services and protocols is unpredictable. There is no measurement science for these systems.

This program is an attempt to develop a measurement science for a broad collection of network systems.
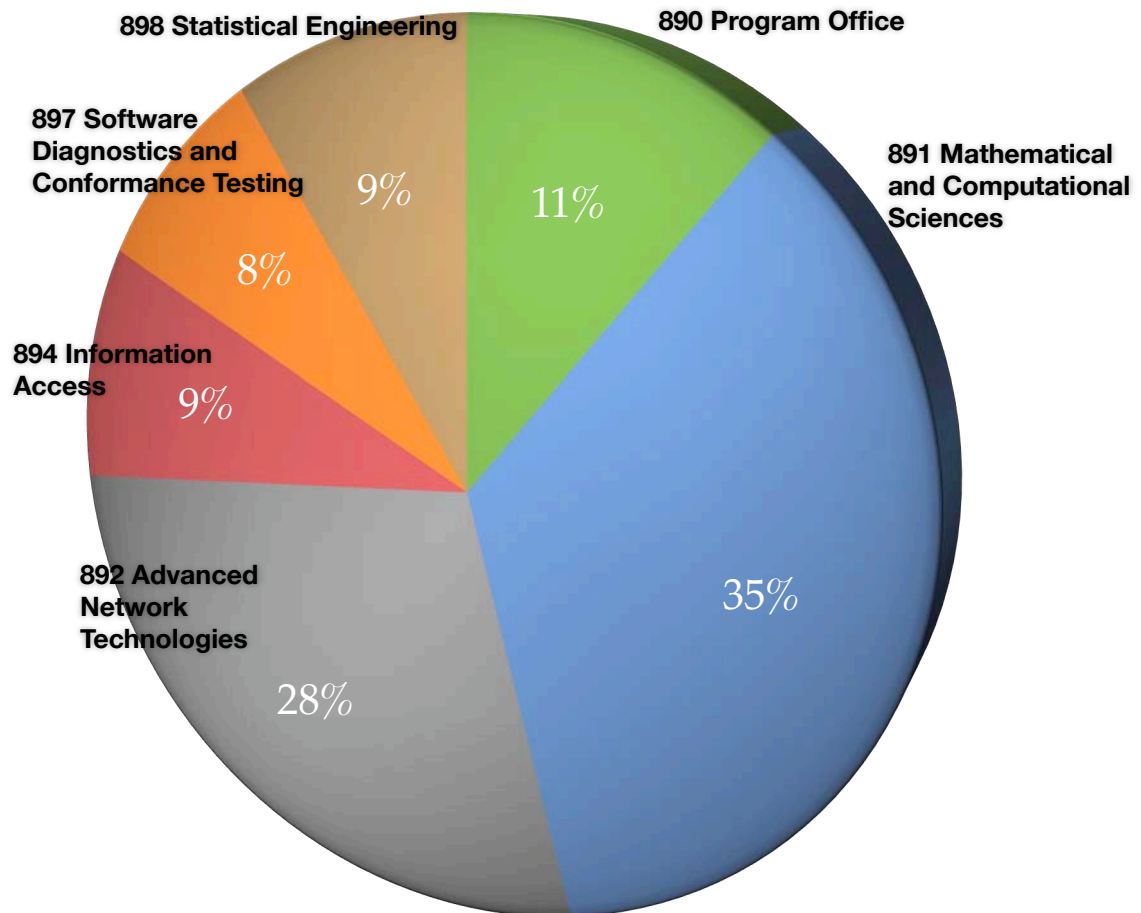
Portfolio of Projects

The projects which comprise the Complex Systems Program interact and enhance each other in some way. The projects are meant to address a gap in research at NIST or an enhancement to a pre-existing project. Clearly, the IMS Measurement Science for Complex Information Systems forms the core of the program. (IMS, Innovations in Measurement Science projects are internally funded basic search projects.) In addition, the need to develop new core competencies is being addressd by three other small projects. Finally, an existing core competence (statistics) project was begun to explore new complex systems problems.

	Project	Gaps
IMS or closely related	Measurement Science for Complex Information Systems	fundamentals, scale
	Visualizing Networks	infoviz
Applying core statistical competence	Discrete Experimental Design Statistics	fundamentals
Building core competence	Multi-robot Performance	performance
	Methods for Characterizing Complex Networks	attributes/scale
Building research community	Complex Sys Fundamentals	fundamentals
	Testbed	testing metrics
	Complex Systems Study Group	fundamentals

Projects

"How does a project get to be a year behind schedule? One day at a time." - Fred Brooks

ITLDivision Participation in Complex Systems Program

898 Statistical Engineering

890 Program Office

897 Software Diagnostics and Conformance Testing

891 Mathematical and Computational Sciences

9%

11%

8%

894 Information Access

9%

892 Advanced Network Technologies

35%

28%

Complex Systems Fundamentals

Roldan Pozo

The study of complex or "self-organizing" networks is a relatively young, yet active research area. Much of the focus is to determine how graphs from "real" phenomena (such as protein interactions, social networks, structure of the World Wide Web, etc.) differ from synthetic or "engineered" graphs that are constructed by design. Clearly they are different, but how? Three key metrics (clustering coefficient, graph diameter, and degree distribution) are commonly used in the literature, and these techniques have served as encouraging first steps, but further research is needed to yield more meaningful distinctions and truly understand the nature of these complex networks.

A fundamental question, is what are the precise "real" graphs that are being used to drive these theories? Clearly, the conclusions being reached can be only as good as the quality and availability of the original data. Most papers provide little information about this, making any kind of comparison or verification of approaches difficult.

For example, a commonly studied collaboration network in the literature is the "Movie Actor" graph, where two movie actors are connected if they have worked together on the same movie. This information is usually culled from the Internet Movie Database (imdb.com). However, there is no single "Movie Actor" graph because this database is changing snapshot of the entertainment industry at any given time. Furthermore, what one chooses to count in their particular search (foreign films, documentaries, TV movies, specific epochs or genres) will affect the resulting graph. In fact, there are several of these "Movie Actor" graphs floating around in the literature, so a citation is not really meaningful unless it contains a timestamp and documenation of the specific queries used in its generation. A similar argument can be made for other association networks, such as co-author or citation graphs.

Another set of examples are networks based on the structure of the World Wide Web. In this case, each web page is a node and two pages are connected if there a hypertext link from one to the other. These are typically generated by webcrawlers and can vary greatly depending on where they crawl and what decisions they make about what constitutes a "link" (e.g. do different anchors to the same page count as separate links?) how non-static pages are handled, what file types are processed, how robust it is with broken links and sloppy HTML pages, and so on. This list can be quite long, and each decision path will yield a different graph. On top of this, the Web is, of course, not a static entity. Many of the Web graphs out there provide few of these details and usually present only the topological structure (i.e. URLs removed) so it is difficult to confirm or validate them independently.

Similar concerns exist for graphs taken from other application domains. Often the data is incomplete, poorly documented, or missing pieces crucial to its description. There are ongoing debates about whether particular networks are really scale-free or not, or whether the conclusions of some papers are actually valid. Usually these concerns lie at the heart of how data was collected and analyzed.

Our motivation for developing a standard reference data set for complex networks is to provide a collection of public and well-documented "real-world"graphs that serve as a test bed of rigorously validated data which researchers can utilize to (1) verify and compare the algorithms and analysis of various approaches in the literature, (2) to ensure that everyone is using the same reference data in their studies, (3) to have one convenient location where one can browse and search for network graphs of various characteristics, (4) to provide a focal point for the research community to contribute and exchange network graphs from various application domains, and (5) to provide a testbed for development of software analysis tools (e.g., graphical

viewers, graph partitioners, clustering algorithms) that will aid in the further development and research of complex networks.

In further aiding the understanding of such network systems, we have developed a set of software tools for working with mathematical graphs (Ngraph++) and processing and transforming data collections into various formats (e.g., Matrix Market, GraphViz, Matlab) for output and analysis. Ngraph++ 64 computes clustering coefficients and degree distributions for directed/undirected graphs. Although a small and simple library, it provides intuitive graph operations (subgraphs, unions, differences) and complexnetwork features not commonly found in more complex packages. Together with various Python and scripting tools, and a reference dataset of network graphs, this forms a kernel package for analyzing complex network data in a portable and universal fashion.

Outputs
Software

NGraph++: a simplified graph library for the analysis of complex networks based on ANSI/ISO C++. Includes basic graph operations (construction, intersections, subgraphs) and algorithms for computing degree distribution and clustering coefficients. R. Pozo

Data Reference Web Site

"Complex Networks Data Sets," a site referencing several collections and data repositories of complex networks used in research. Includes graphs & data used by Barabasi, the Stanford WebBase projects, the Pajek collection of network graphs, email databases from Enron (social networks), and internet maps of Bill Cheswick.

Presentations

Feb 9, 2007 - Roldan Pozo, review of paper Statistical Mechanics of Complex Networks by Albert & Barabasi to NIST Complex Systems Study Group

# Measurement Science for Complex Information Systems

*Kevin Mills(**proj lead**) Daniel Genin, Christoph Dabrowski, Fern Hunt, Vladimir Marbukh, Don Cho Yeon, Dennis Leber, James Filliban*

The Innovations in Measurement Science (IMS) project is the largest project in the Complex Systems Program and indeed forms the heart of the program. It consists of several highly interrelated sub-projects.

No one understands how to measure, predict or control macroscopic behavior in complex information systems. This threatens our national security and costs billions of dollars. Our purpose is to formulate and study selected models and problems involving dynamical systems that are of potential relevance to understanding the behavior of large and complex information systems. Unlike the physical sciences where there is a large body of measurable phenomena caused by natu-

rally occurring forces, computer networks are human made and with haphazard development so salient measurement and metrics for explaining and predicting their behavior are not obvious. We seek to use the tools of dynamical systems and stochastic models to analyze networks of various topologies in an effort to understand how network topology and dynamics among and between the various layers of the network interact. A product of our efforts would be the theoretical foundation of tools for predicting and controlling network operation.
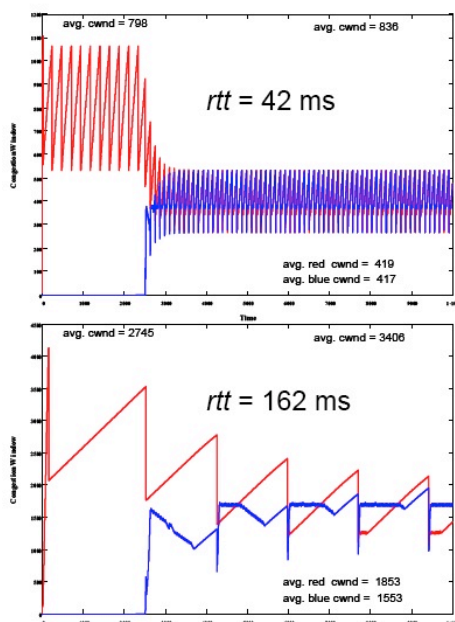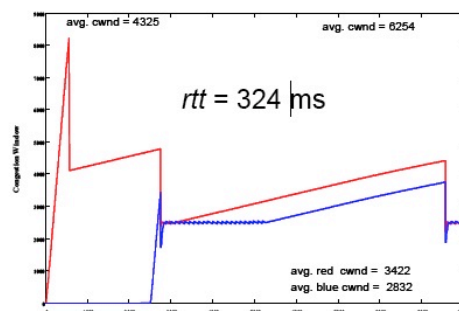
Output from MesoNet a Mesoscopic Simulation Model of a Router-Level Internet-like Network

We are working to fill a gap that exists between detailed simulations and abstract mathematical models. Detailed simulations are difficult to parameterize and infeasible to use when simulating large topologies with many simultaneous TCP flows. Abstract mathematical models (such as fluid approximations), lose the iden-

tity of individual flows and packets and sacrifice the ability to model a network under changing conditions.

One other research group (at UC Santa Barbara) is also working to fill this gap. Their approach is to produce a hybrid model that computes probabilities of packet

losses and then uses those probabilities to assign losses to specific simulated TCP flows as time progresses.

The goal of these researchers is to produce simulations that can be verified against ns2 for small topologies. As described in the literature, their approach has been applied to a single-tier (Abilene backbone) network with up to 100 simultaneous TCP flows. The hybrid model includes sub-models for several proposed replacement congestion-control algorithms. The Santa Barbara research team considers investigation of alternate congestion-control algorithms to be a feasible application of their hybrid model.

Outputs
Papers
K. Mills, A Brief Survey of Self-Organization in Wireless Sensor Networks, Wireless Communications and Mobile Computing, Wiley Interscience, Vol. 7, No. 7, (pages 823-834) October 2007.

V. Marbukh, S. Klink, Decentralized control of large-scale networks as a game with local interactions: cross-layer TCP/IP optimization, ValueTools, 2nd International Conference on Performance Evaluation Methodologies and Tools, Nance France, Oct 2007.

V. Marbukh, Utility Maximization for Resolving Throughput/Reliability Trade-offs in an Unreliable Network with Multi-path Routing, ValueTools, 2nd International Conference on Performance Evaluation Methodologies and Tools, Nance France, Oct 2007.

V. Marbukh, Fair bandwidth sharing under flow arrivals/departures:effect of retransmissions on stability and performance, in ACM SIGMETRICS Performance Evaluation Review, Volume 35 Issue 2, September 2007.

Software
Flexi-Cluster: A Simulator for a Single Compute Cluster, V. Marbukh and K. Mills.

MesoNet: A Mesoscopic Simulation Model of a Router-Level Internet-like Network, K. Mills

MesoNetHS adds six congestion-control algorithms for networks with high bandwidth-delay products; algorithms include: binary-increase control (BIC), compound TCP (CTCP), fast AQM scalable (FAST) TCP, H-TCP, high-speed TCP (HS TCP) and Scalable TCP (STCP), K. Mills, E. Schwartz and J. Yuan.

EconoGrid: A detailed Simulation Model of a Standards-based Grid Compute Economy, C. Dabrowski and K. Mills.

MesoGrid: a mesoscopic scale simulation model of a computational grid economy, C. Dabrowski and K. Mills.

Presentations
Vladimir Marbukh, Towards Price Based Network Management and Provisioning , May 6, 2008, NST Complex Systems Study Group

J. Filliben, Sensitivity Analysis Methodology for a Complex System Computational Model, 39th Symposium on the Inter-

face: Computing Science and Statistics, Philadelphia, PA, May 26, 2007

D. Genin, Percolation: Theory and Applications, NIST Complex Systems Study Group, Oct 17, 2007.

Grid Simulation using Piece-wise Homogeneous Markov Chains

Christopher Dabrowski and Fern Hunt

This past year, C. Dabrowski and Fern Hunt have collaborated on the development of a Markov chain model of the operation of a large-scale computational grid system. Typically, large-scale discrete event simulation can be used to model the interaction of service providers and consumer clients as they enter into service agreements for the use of grid computing resources. Such simulations, particularly those that seek to accurately reproduce system behavior in detail, are computationally expensive; even simulating a single day takes many hours of CPU time. Nevertheless, there is a critical need for models that can provide rapid, and potentially scalable, simulations of large-scale grid systems. By varying key parameters and decision processes, such models can be systematically perturbed to produce alternative system execution paths and identify scenarios that lead to system failures or other anomalous behavior.

By representing the behavior of the large-scale grid model in seven system states, Dabrowski was able to create a piecewise homogeneous Markov chain that captured the evolution of a grid system over an 8-hour period. The resulting Markov chain representation served as a basis for the development of a perturbation scheme that could be used to model the impact on system performance caused by failure to meet key grid system quality of service guaranties. The Markov chain simulations compared very favorably with the results of large-scale grid simulations for the same scenarios. The range of system behaviors produced by the Markov chain simulation provided reasonable approximations for the results of the large-scale simulation (as measured by the job completion rate). Moreover, the perturbation method made it possible to identify changes to probabilities of state transitions that caused severe declines in system performance and to identify possible reasons for system failures. The Markov

chain simulations proved to be good predictors of the results of large-scale simulations and required less than two orders of magnitude compute time than the large-scale simulation.

Outputs
Papers

K. Mills and C. Dabrowski, Can Economics-based Resource Allocation Prove Effective in a Computation Marketplace? in Journal of Grid Computing Special Issue on Grid Economics and Business Models. DOI (10.1007/s10723-007-9094-4)

C. Dabrowski, Reliability in Grid Computing Systems. submitted to the journal Concurrency and Computation: Practice and Experience as part of a special issue on OGF work

V. Marbukh and K. Mills, On Maximizing Provider Revenue in Market-based Compute Grids in Proceedings of ICNS07, June 2007.

Analysis of a Distributed Protocol for Network Control

Fern Hunt, Vladimir Marbukh

Congestion control is critical to the (smooth) functioning of modern communication networks in general and the Internet in particular. The goal of congestion control protocols is to allow many users to share network resources without causing congestion collapse. In addition, it is desirable to use as much of the available network capacity as possible and ensure some fairness among users. It was shown by Frank Kelly of the University of Cambridge, and later by many others, that achievement of these objectives can be represented as a problem of global optimization of system utility where some notion of fairness is expressed by a choice of utility function. Congestion control can then be seen as a distributed iterative solution to this problem (Lin and Shroff).

Recently Vladimir Marbukh considered the problem of appropriately routing traffic networks within the TCP/AQM protocol and proposed the introduction of controls that randomly assign routes to traffic in a way that minimize the mean link load. A utility optimization problem was formulated that incorporates a random route allocation scheme, defined by a probability distribution whose degree of randomness is controlled. The distribution entropy was used to modulate between the two extremes of path specification; minimum cost path selection, as used in the OSPF routing protocol on

one end and equiprobable path allocation on the other. Earlier, in simulations of a stochastic model, Marbukh and Klink showed that networks operating under a random allocation scheme displayed oscillatory behavior when the entropy approached zero — corresponding to minimum path cost routing, and non-oscillatory (called stable) behavior as the entropy increased passed a critical value.
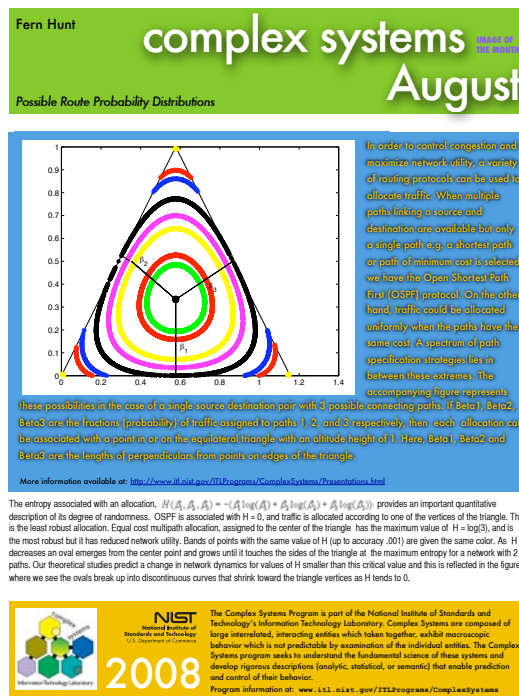
A general model of network operation was derived under this congestion/routing control protocol as a solution of our optimization problem. The behavior of this model was examined for two sample topologies. Here we demonstrated the existence of equilibria that are determined by the link capacities and the topology of the network. Given a topology and set of link capacities there is a distinguished critical route distribution where one sees stable equilibrium behavior. This is maintained for distributions with larger entropy but unbounded behavior appears for distributions with smaller entropy. This is analogous to the simulation observations of Markbukh and Klink. We show that under smoothness assumptions on our objective func-

tion, equilibrium is the solution of the original optimization problem.

One of our examples, a five-link, three-route network, as simple as it is, shows a new phenomenon that is not present in a two link network — the first and simplest possible network. Using the route distribution entropy as a parameter it can be shown that the problem changes character once the entropy of the three route network reaches log 2, the largest entropy achievable by a two route network. This is illustrated in the figure on this page.

For entropies less than log 2, the region defined by the optimization problem is no longer convex. Thus for the problem we consider the range of route distributions is restricted to those that are not achievable by a simpler topology with two links.

Each allocation can be associated with a point in or on the equilateral triangle with an altitude height of 1. Here β_1, the perpendicular distance of the point to the horizontal base of the triangle is the fraction (or probability) of traffic sent to route 1. β_2, the perpendicular distance from the point to the left side of the triangle is the fraction allocated to route 2 and β_3, the fraction allocated to route 3 is the analogous distance for the right side. The entropy associated with an allocation,

$$H(\beta_1, \beta_2, \beta_3) = -(\beta_1 \log(\beta_1) + \beta_2 \log(\beta_2) + \beta_3 \log(\beta_3)$$

provides an important quantitative description of its degree of randomness. In OSPF (H = 0), traffic is sent along a path defined by one of the vertices of the triangle. Equal cost multipath allocation is assigned to the center of the triangle and has the maximum value of H = log(3). Bands of points with the same value of H (up to accuracy .001) are given the same color. The black band is associated with the critical value H = log(2), the maximum entropy for a network with 2 paths.

Outputs
Papers
F. Hunt and V. Marbukh, Dynamic Routing and Congestion Control Through Random Assignment of Routes, in International Multi-Conference on Engineering and Technological Innovation, July 2008, pp. 161-164.
Presentations
Fern Hunt, Mathematical Modeling of TCP/IP protocols , May 17, 2008, NIST Complex Systems Study Group.

Metastability in Complex Communications Systems

Two projects in particular have addressed the concept of metastability. These are systems with multiple relatively stable equilibria, called metastable equilibria. A metastable system can move unpredictably and abruptly from one metastable state to another, between randomly distributed periods of stability. Often one of the metastable states is undesirable and can be viewed as a failure mode of the system. This was, the case with the Internet congestion collapses of the early 1990's. Prediction of conditions under which metastability can arise and design of mechanisms for its control are topics of great practical interest.

TCP Metastability and Cascading Failures in the Internet

Vladimir Marbukh

Computer networks that provide the foundation for modern commerce are astounding in their complexity. Hundreds of millions of nodes interact in real time controlled by lean communication protocols. Unfortunately, unexpected macroscopic behavior, such as waves of network congestion and failure, have been observed in this large complex system. Such events have already led to the loss of opportunity and productivity valued in the billions of dollars. Our fundamental understanding of such behavior remains weak. It remains impossible to reliably predict the emergent collective behavior of the large complex interconnected systems that we can now readily build and deploy. The goal of this project is to develop and investigate realistic models of the Internet subject to flow control protocols. A preliminary analysis of a mean-field model with retransmissions induced by control feedback indicates the possibility of instability in cases thought to be stable from the analysis of simpler models.

The Transmission Control Protocol (TCP) controls data rates of the majority of Internet users. While a large number of TCP versions have been implemented, and even more versions have been proposed, all rely on

Explicit Congestion Notifications (ECN) as the indication of congestion on the end-to-end path. Under this protocol, each controlled TCP source increases its data rate until it encounters congestion as indicated by an ECN. Once an ECN is received, the source reduces its rate to reduce congestion. In this way the Internet operates as a distributed control system with feedback provided by Explicit Congestion Notifications.

High dimension, non-linearity and delays in the feedback loop make stability and performance evaluation of the Internet under TCP control a very challenging problem. Observations of the Internet, simulations of Internet fragments, and analytical modeling under simplifying assumptions, suggest that the TCP controlled Internet is stable under "normal" operating conditions.

Given the number of flows $N = N_r$ carried on all feasible routes $r \in R$, the equilibrium source transmission rates x_r can be approximated by the weighted (α,w)-fair rate assignment [1],

$$x_r = \left(w_r \Big/ \sum_{l \in r} p_l \right)^{1/\alpha} \qquad (1)$$

where the parameters α, $w_r > 0$ characterize TCP version-specific fairness, and the parameters p_l represent the average rate of Explicit Congestion Notifications generated by links l on route r. The average transmission rate of each TCP controlled Internet user is determined by the aggregate ECN rate from all links on the user's route r. For most widely used TCP versions, ECNs indicate packet losses and thus the parameters p_l in (1) represent the packet loss probabilities on links l. When $w_r = 1$, the cases $\alpha \to 0$, $\alpha \to 1$, and $\alpha \to \infty$ in (1) correspond, respectively, to a rate assignment which achieves maximum throughput, is proportionally fair, or is max-min fair. The case $\alpha = 2$ corresponds to the so-called TCP-Reno rate assignment in the case of small packet losses $p_l \ll 1$.

The number of active Internet users fluctuates over time, with each user alternating between active and silent periods, causing arrival and departures of flows

carrying files, streaming voice and video. Markov models of fair bandwidth sharing (1) under arriving/departing flows have been proposed and analyzed in [2,3,4]. This analysis suggests that the Internet is stable under arriving/departing flows if the average exogenous file transfer load is less than the link capacity for all network links, that is,

$$\rho_l \overset{def}{=} \frac{1}{C_l} \sum_{r:l \in r} \Lambda_r b_r < 1, \ \ \forall l \in L . \qquad (2)$$

In (2) the average arrival rate of files to be transferred on route r is Λ_r, the average size of these files is b_r, the capacity of link l is C_l, and hence the utilization on link l is ρ_l.

Nevertheless, simulations [5] have demonstrated that in the case when ECNs indicate lost packets, and thus parameters p_l in (1) indicate probability of packet loss, the network may be unstable even when conditions (2) are satisfied for all Internet links l. The reason for the inconsistency between analytical modeling [2,3,4] and simulations [5] is that models [2,3,4] do not account for retransmissions of lost packets, which TCP uses to ensure reliable transport.

The goal of our project is to develop and investigate realistic models of TCP controlled Internet under arriving/departing flows. A mean-field model which incorporates packet retransmissions has been proposed in [6,7]. The model accounts for additional link load due to packet retransmissions on links carrying packets dropped downstream and then retransmitted. This additional load increases with increase in the number of flows in progress, creating a positive feedback: sufficient increase in the number of flows in progress due to temporary load increase of fluctuation further increases the number of flows in progress, etc. Consistent with previous simulation results [5], our preliminary stability analysis indicates that this positive feedback may cause model instability even when conditions (2) are satisfied. Model instability implies that the number of flows in progress grows unbounded in time resulting in congestion collapse. Of course, infinite growth of the number of flows in progress is an abstraction since in

the Internet demand drops as user dissatisfaction grows with the increase in the number of flows in progress.

Preliminary analysis [6,7] has also produced new findings. A large-scale TCP controlled network with exogenous load satisfying (2) has a desirable metastable, i.e., persistent but slowly changing, network state with a finite number of flows in progress. Despite the fact that eventually the network may become unstable resulting in infinitely growing number of flows in progress, on time scales of practical importance the network converges to the desirable metastable state if the initial number of flows in progress is sufficiently small. By analogy with phase transitions of the first kind, one may expect that eventual transition from desirable metastable state to congestion collapse occurs through process of cascading failures: supercritical local overload may spread over a major portion of the network causing global congestion collapse. Our preliminary results indicate that appropriately designed flow admission control can transform the desirable metastable state into a stable state, eliminating instability at the cost of a small flow rejection probability.

Current and future research involves the validation of model results, estimation of the stability margin of the desirable metastable state, and developing practical techniques for preventing cascading failures and the possibility of global congestion collapse. More long-term plans include investigation of the effect of specific TCP parameters, including fairness, on TCP stability and performance. Evaluation and balancing of the numerous trade-offs between TCP fairness, stability, performance and robustness under fixed and fluctuating sets of flows in progress are matters of significant and urgent practical importance for optimization and design of current and future versions of the TCP protocol.

References

[1] J. Mo and J. Walrand, Fair end-to-end window-based admission control, IEEE/ACM Transactions on Networking 8 (2000), pp. 556-567.
[2] S. Ben Fredj, T. Bonald, A. Proutiere, G. Regnie, and J. Roberts, Statistical bandwidth sharing: a study of congestion at flow level, SIGCOMM 2001.
[3] G. de Veciana, T.J. Lee, and T. Konstantopoulos, "Stability and performance analysis of networks supporting elastic services," IEEE/ACM Transactions on Networking 9 (2001), pp. 2-14.
[4] P. Key, L Massoulie, A. Bain, and F.P. Kelly, Fair Internet traffic integration: network flow models and analysis, Annales des Telecommunications 59 (2004), pp. 1338-1352.
[5] L. Massoulié and J.W. Roberts, Arguments in favor of admission control for TCP flows, ITC 16, Edinbourg, June 1999.
[6] V. Marbukh, Metastability of fair bandwidth sharing under fluctuating demand and necessity of flow admission control, Electronics Letters 43 (13 September 2007).
[7] V. Marbukh, Can TCP Metastability Explain Cascading Failures and Justify Necessity of Flow Admission Control in the Internet? Proceedings of the 15th International Conference on Telecommunications, St. Petersburg, Russia, June 2008.

Outputs
Papers
V. Marbukh, Metastability of fair bandwidth sharing under fluctuating demand and necessity of flow admission control in Electronics Letters, 13th September 2007, Vol 43, No 19.
Presentations
D. Genin, Metastability and networks , Feb 26, 2008, NIST Complex Systems Study Group

Metastability in Cellular Networks

Vladimir Marbukh, Daniel Genin

Another system with potential for metastable behavior is the cellular communications network. Here metastability arises as the result of interaction between users with different resource demands (e.g. voice and video) and their migration in the network. Our models show that depending on the initial network state and exogenous load, the network may settle into one of several metastable states or wander among them erratically. In practice, neither situation is acceptable since the network operator needs to be able to guarantee a certain quality of service to her customers. Mean-field approximation allows us to predict the metastable states of the network, which may further be made globally unique and stable with admission control. In the future, it may be possible to put metastability to good use by allowing the network operator to select the operating network regime from a set of metastable states with the aid of admission control.
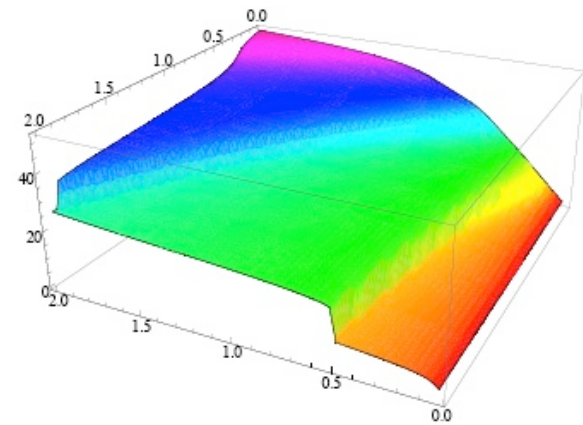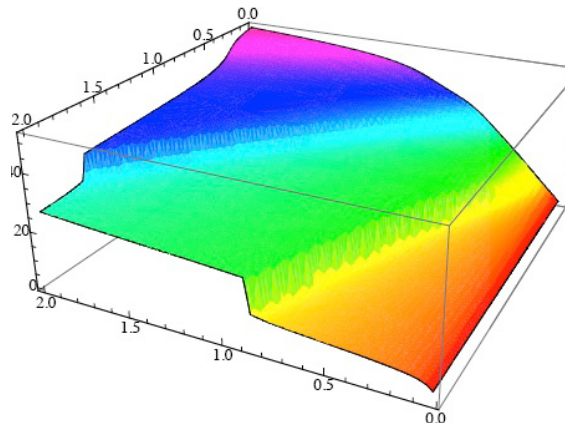
Convex hulls of the lower and upper branches respectively of the equilibrium manifold for service class 1. Phase diagram (below) constructed by projecting the equilibrium manifold by averaging. Regions outlined in black represent areas of metastability (coexistence).
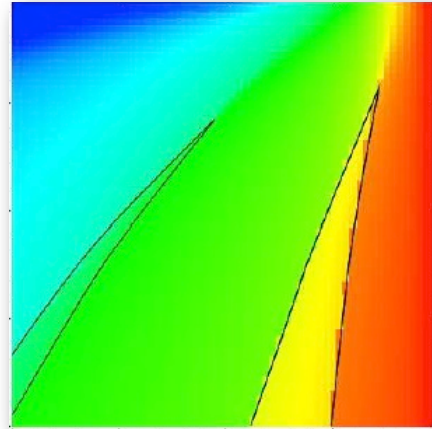
Outputs
Papers

D. Genin and V. Marbukh, Toward Understanding of Metastability in Large-Scale Loss Networks with Mobile Users: Emergence and Implications for Performance Proceedings of 12th World Multiconference on Systems, Orlando FL, June 2008.

.

Methods for Characterizing Complex Networks

*Isabel Beichl(**proj lead***) and Brian Cloteaux*

A type of data of increasing importance in information-technology is one not based on continuum models but rather on connections between objects which can be-modeled as graphs or networks. Complex systems of this type arise in many applications, such as the power grid, the Internet, communication systems, and trans-portation systems, situations where resources are pro-vided through restricted channels.

In order to quantify the security and reliability of com-plex IT systems, it is important to develop an underly-ing measurement science based on sound mathematical methods. Unfortunately, the size of real networks often makes exact measurement impossible. Thus, we are building efficient tools to approximate properties of real networks. An important aspect is the creation of realistic models to support research in this field. In particular, it is necessary to understand how real net-works differ from random models.

Towards these goals, we have been developing new methods to generate random network models con-strained to have a given degree sequence and s-metric. In addition, we are developing Monte Carlo methods and software based on sequential importance sampling for measuring properties of these networks. Applica-tions considered have been counting the number of spanning trees of a graph and approximating the coef-ficients of the reliability polynomial.

A *graph* is a collection of nodes (or vertices) and a set of connections between them, called edges. The *degree* of a node is the number of edges that are connected to it. A graph is termed *connected* if there is a path be-tween any two nodes. Graphs are very useful in model-ing of interconnected systems in a wide variety of ap-plications from communications to biology.

Network Model Generation. The theory of random graphs is well developed. However, many real world systems cannot be approximated using simple random graphs (the Erdős-Renyi model). One reason for this is that the distribution of nodal degrees of random graphs

follows a Poisson distribution, whereas it has been shown that the degree distribution for many real-world systems follows a power law distribution. In other words, the probability p of a node having degree k is $p \sim k^\alpha$ for some α. Nevertheless, generating random graphs with a power law degree distribution is still not sufficient to model real networks. For any given degree sequence, there can be a large number of non-isomorphic graphs with that sequence.

We have shown that real networks can have very dif-ferent properties from random networks with the same degree distribution. We demonstrated this by examin-ing the structure of the autonomous systems (AS) to-pology of the Internet as measured at UCLA [6]. We compared this with a series of random graphs with the same degree sequence generated using an algorithm of Blitzstein and Diaconis [3]. The conclusion was that simple degree-based models do not capture important details contained in real world data.

Thus, network models need more structure to be realis-tic. To be able to further distinguish between networks with identical degree sequences, Li *et al.* [5] intro-duced the *s-metric*. This measures the connectedness of high degree nodes in networks having a power law distribution. Formally, the s-metric value for a graph G is defined as

$$s(G) = \sum_{(i,j) \in E(G)} \omega_i \times \omega_j$$

where $E(G)$ is the set of edges and ω_i is the degree of vertex i.

We have developed a method to generate graphs with a given degree sequence and an *s*-metric close to a given target [1]. It uses a Monte Carlo technique called threshold acceptance. If we have a graph with a given degree sequence, we use random edge switches that do not change the degree sequence but can increase or reduce the *s*-metric. While the models generated this way are a definite improvement, they still do not cap-

ture certain critical structural details of the AS topology [2]. We are continuing to develop new methods for more accurate network modeling.

Computing the Reliability Polynomial. A simplified measure of the reliability of a network is its reliability polynomial. This function gives the probability that a network will remain connected given that each edge in it can fail independently with a probability p. One definition for the reliability polynomial R of a connected graph G with n vertices and m edges and for the edge failure probability p is

$$R(G; p) = \sum_{j=1}^{m} a_j (1-p)^j \, p^{m-j}$$

where a_j is the number of distinct connected subgraphs of G with exactly j edges.

The problem in using the reliability polynomial for network analysis has been the fact that determining the coefficients of the polynomial itself is computationally intractable. In fact, it has been shown that, in general, this problem is #P-complete. Thus the best we can expect to do in general is to approximate the polynomial. Our current research has led to a new method to do this. Our approach has two advantages. First, while there are several algorithms that provide pointwise approximations of the polynomial, our approach approximates the polynomial coefficients themselves. Secondly, when compared to the known approximation method of Colbourn *et al.* [*Estimating the Coefficients of the Reliability Polynomial*], our method has shown much faster rate of convergence.

Counting Spanning Trees. A *spanning tree* in a graph is a set of edges such that each node is connected and there is no cycle in the edges (see Figure 10). The number of spanning trees in a graph can be a measure of reliability. As the number of spanning trees increases, the number of edges that need to be cut to disconnect the graph also increases (i.e., more spanning trees means many alternate paths between any two nodes).

In order to efficiently count the number of spanning trees in a network, we have developed a Monte Carlo method based on sequential importance sampling. The reason for creating an approximation algorithm to count spanning trees is slightly different than for estimating the reliability polynomial. There is an exact method to count all spanning trees of a graph in polynomial time. Unfortunately, this method becomes impractical for the size of the many of networks we measure. An example of a larger network that our new algorithm has let us analyze is the AS topology of the Internet. One instance of the AS topology, which we obtained from the UCLA Internet topology database, contained 24,567 nodes and 102,946 edges. Using our algorithm, we were able to compare random graphs to this AS topology in terms of number of spanning trees.
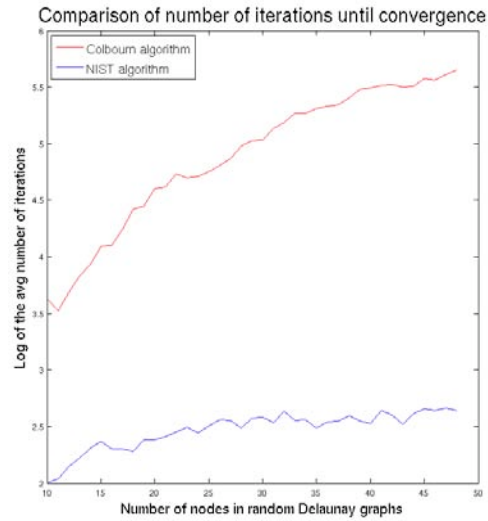
Comparison of Colbourn and NIST algorithms

Figure 10. *Illustration of the number of spanning trees in a graph.*

References

[1] Isabel Beichl and Brian Cloteaux, Generating Network Models using the S-metric," in Proceedings of the International

Conference on Modeling, Simulation and Visualization Methods, 2008, pp. 159–164.

[2] Isabel Beichl and Brian Cloteaux, Measuring the Effectiveness of the s-metric to Produce Better Network Models," in Proceedings of the Winter Simulation Conference, 2008, pp. 1020–1028.

[3] Joseph Blitzstein and Persi Diaconis, A Sequential Importance Sampling Algorithm for Generating Random Graphs with Prescribed Degrees," submitted. 2006.

[4] Charles J. Colbourn, Bradley M. Debroni, and Wendy J. Myrvold, Estimating the Coefficients of the Reliability Polynomial, Congressus Numerantium 62 (1998), pp. 217–223.

[5] Lun Li, David Alderson, John C. Doyle, and Walter Willinger, Towards a Theory of Scale-free Graphs: Definition, Properties, and Implications, Internet Mathematics 2 (2005) pp. 431–523.

[6] UCLA. UCLA internet topology collection. http://irl.cs.ucla.edu/topology, 2008, accessed 1/14/2008.

Outputs

Papers

I. Beichl and B. Cloteaux, Generating Network Models Using the S-Metric, in Proceedings of the International Conference on Modeling, Simulation and Visualization Methods, Las Vegas, NV, July 2008, pp. 159-164.

I. Beichl and B. Cloteaux, Measuring the Effectiveness of the s-Metric to Produce Better Network Models, in Proceedings of the Winter Simulation Conference, Miami FL, Dec 2008, pp. 1020-1028.

Presentations

Isabel Beichl and Brian Cloteaux "Counting Spanning Trees"j, Oct 31, 2007, NIST Complex Systems Study Group.

B. Cloteaux, "How can we reliably model massive networks?" postdoc poster presentation for Sigma Xi, NIST, January 14, 2008. **Won best poster in mathematics and statistics.**

I. Beichl - "Using Sequential Importance Sampling to Speed Up MCMC", May 21, 2008, NIST Statistical Engineering Division.

B. Cloteaux - "Generating Network Models Using the S-Metric", July 15, 2008, International Conference on Modeling, Simulation and Visualization Methods, Las Vegas, NV.

B. Cloteaux - "Generating Network Models" August 27, 2008, IDA Center for Computing Sciences, Bowie, MD

Software

B. Cloteaux , A C++ implementation of a sequential importance sampling algorithm, created by us, for estimating the number of spanning trees in a graph, Sept. 2007.

B. Cloteaux , A C++ implementation of various network metrics including minimum vertex cover (MC). Uses several of the kernelization techniques of Langston et al. The kernelization method involves taking a graph G and producing a smaller graph G', the kernel, with a value k such that MC(G) = MC(G')+k. For the smaller G' we solve for minimum vertex cover exactly, Feb. 2008.

B. Cloteaux , A program for generating random graphs with a given degree sequence. Took code written in R by Joseph Blitzstein at Harvard, converting it to C++, and then optimizing it with the use of novel datastructures. These datastructures allowed the C++ program to run much larger cases than the original program. In particular, it could then handle the Autonomous Systems data from the UCLA database. Software has been released to J. Blitzstein (Harvard), C. Priebe (Johns Hopkins), J. Devinney (IDA Center for Computing Sciences), March 2008.

B. Cloteaux , Software to generate graphs with specified s values. These used the Monte Carlo technique of threshold acceptance, May 2008.

B. Cloteaux , Software to calculate the maximum s value for a set of graphs. This includes a new approximation technique we developed, namely a deterministic version of Blitzstein's algorithm to produce the upper bound, May 2008.

B. Cloteaux , Software to enumerate all graphs with a given degree sequence.Done for illustration purposes for small graphs, May 2008.

B. Cloteaux , C++ implementations of Wilson's algorithm to uniformly sample spanning trees in graphs and Welsh's Monte Carlo algorithm for counting bases of a frequent matroid, Aug. 2008.

B. Cloteaux , Software implementing our algorithm, based on sequential importance sampling, to approximate the coefficients of the reliability polynomial of a graph, Sept 2008.

B. Cloteaux , A C++ implementation of Colbourn's algorithm to estimate the coefficients of the reliability polynomial, Sept 2008.

B. Cloteaux , Software to estimate specific values of the reliability polynomial. Uses Chernoff bounds to determine convergence. It produces a graph of connected probability vs edge probability, Sept 2008.

Visualization of Network Dynamics

*John Hagedorn, Cedric Houard, Judith Terrill(***proj lead***), Dong Yeon Cho*

The Measurement Science for Complex Information Systems (IMS) project is producing vast amounts of data. This project will visualize the three classes of computer models of networks - micro, meso, and macro models specified in the IMS project. Sites on the network consist of provider sites, user sites, and directory sites. There are compute clusters at sites with multiple jobs and users. The micro model is very detailed with thousands of sites. It also has a spatial component with x,y,z values. The meso model is less detailed with multiple sites and no spatial component. The macro model consists of a single site with a single cluster with multiple jobs and users. The goal is to develop an understanding of the spatial and temporal behavior of the systems.

The goal of the visualization will be to assist in understanding the behavior of these systems.The visualization will combine in an informative way:direct visualization of the components of the network models and outputs of the models such as statistics and time series. The ability to interact with the models and their outputs will be included.
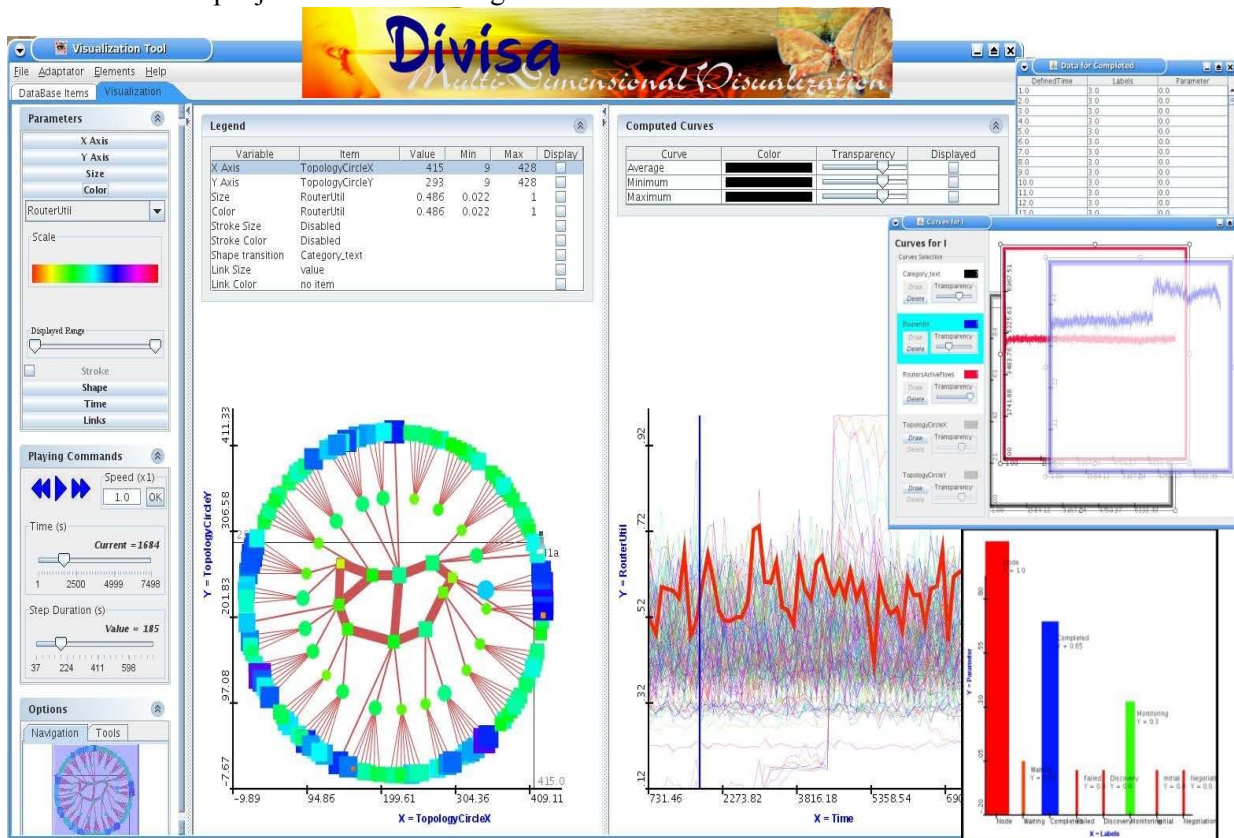
One frame from an animated visualization of a simulation of the "Abilene" network. On the left panel, the innermost nodes are the backbone routers, the next ring represents the subnet routers, and the outermost ring shows the leaf routers. On the right panel, the time series curves of the routers (one is currently selected) are displayed. On the right side, different options are available: numeric data display (top), superposed time series curves (middle); and diagram view (bottom).

This interaction will include the ability to develop other visualizations on the fly and the ability to conduct varying analyses on the data and display it along with the models. The visualization of the network model will be created, where possible, from a user specified set of descriptions of the model, outputs, and analyses. The visualizations will be created automatically from this file or set of files. This will provide a way to create visualizations from an arbitrary run of the models.

If it is possible to run the model on the visualization machine, additional visualization capability will be possible. For example, a set of runs that constitute a design of experiments could be specified and run, with the results summarized in the visualizations.
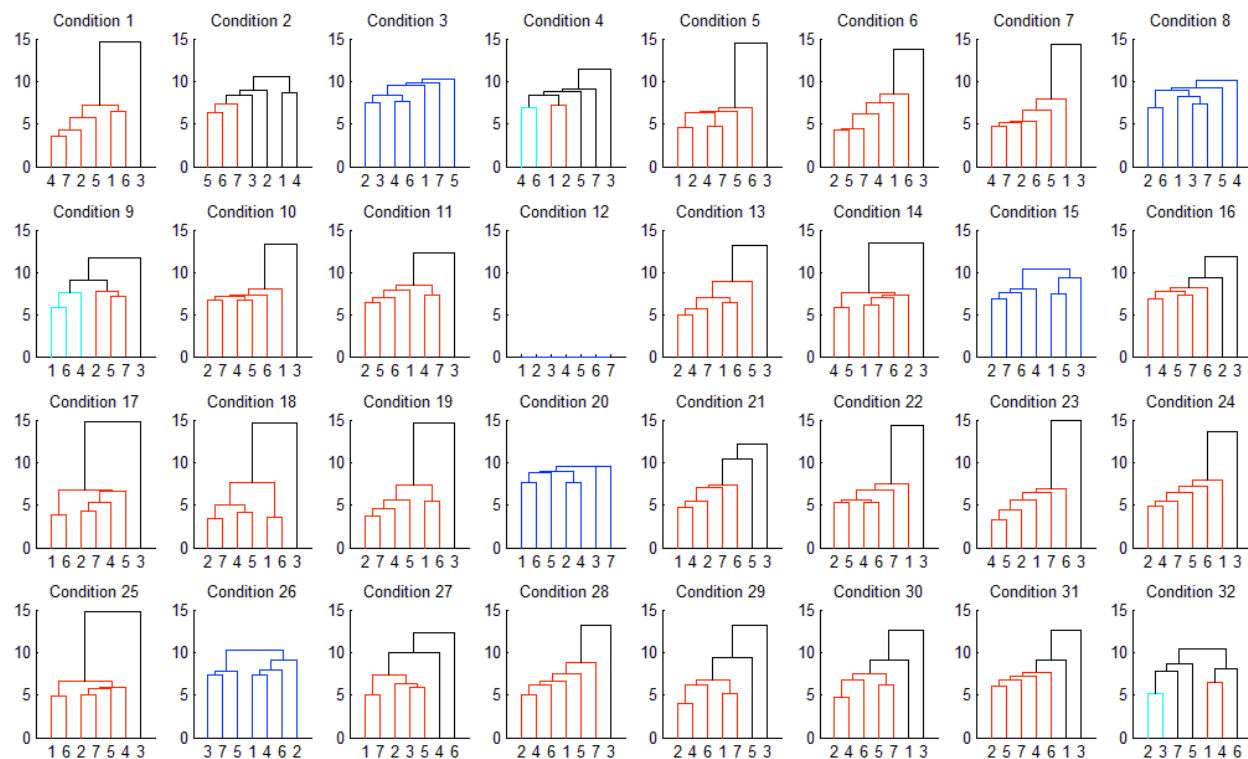
A cluster analysis of the differences in one time period from among 45 responses. The X axis represents y congestion control algorithms. Each "small" figure represents one of 32 conditions applied to the 7 algorithms. Note that in condition 12 we see an identical response from each algorithm (due to no congestion on the network). Also note that in many of the conditions the largest distance graphed is from algorithm number 3 and the rest of the algorithms. Number 3 is the "FAST" TCP congestion control algorithm which performs markedly different.

Outputs
Software
DiVisa: a multi-dimensional visualization tool. Can read many types of data (simulation, statistics, text or numeric, etc.), and converters have been implemented to read several data formats without need for reformatting.

Visualizations
Animation of simulation of Abilene network. C. Houard, J. Hagedorn. http://www.cam.nist.gov/~hagedorn/abilene/

Complex System Testbed

Vincent Stanford(proj lead), Lukas Diduch, Imad Hamchi, Mathieu Hoarau, Antoine Fillinger

NIST Data Flow (NDF) System II provides abstractions for distributed data acquisition and processing. The NDF system is computationally powerful and we are investigating its use as a distributed compute plafom for the computational needs of complex systems problems.

The NDF has been used successfully to conduct large scale scientific data collection and creation of research corpus for computer vision and spoken language research. It supports the acquisition and synchronization of multimodal input streams with hundreds of microphones and multiple high definition video streams distributed across many computers. It is a lightweight protocol capable of acquiring and transporting gigabit data streams on commodity hardware, and its computational metaphor can be adapted and extended to the needs of complex cystems simulations. complex system simulation, and analysis of complex datasets from sensor nets, will necessarily involve the generation and acquisition of large amounts of data which will be processed by data driven computational arrays or grids.

The existing NDF computational infrastructure will need certain extensions to function as a testbed for complex systems studies. Broadly we see complex systems applications as computational grids with various capabilities embedded into the nodes.
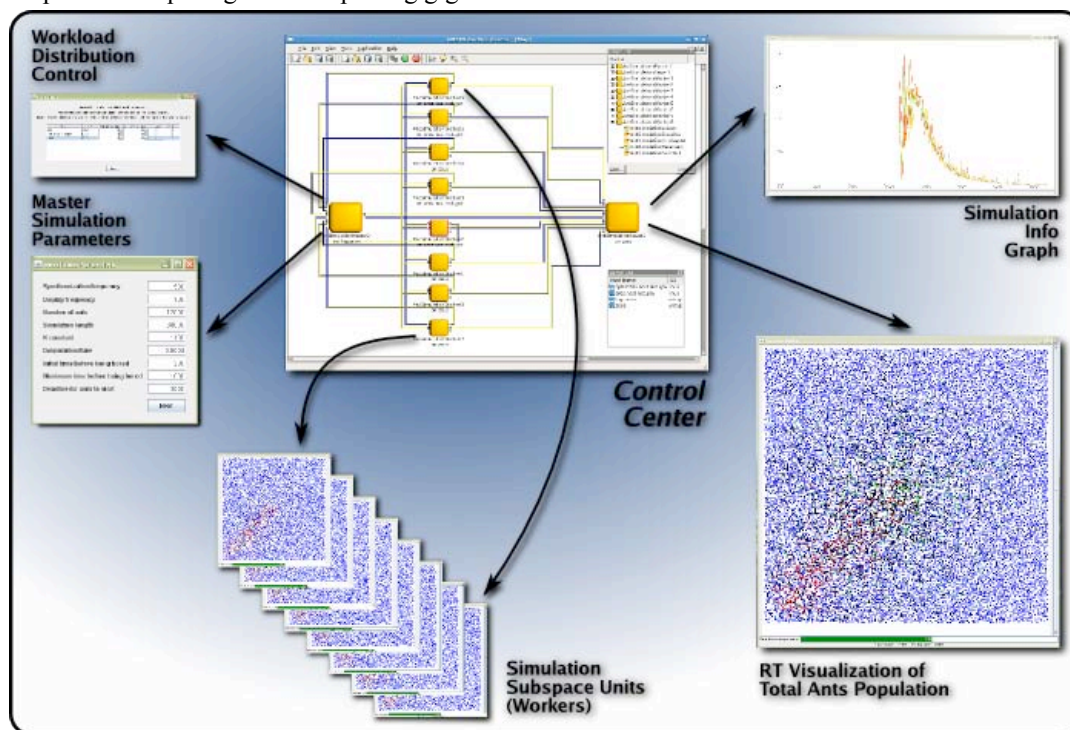
Agent-based simulations show emergent phenomena when numerous simple agents interact to solve large problems. For example, Ant Colony Optimization (ACO) is an agent-based graph search architecture in which agents operate randomly as individuals until the population increases to a critical value where phase change to collaborative search with optimality properties emerges. Solving large problems using these methods requires large populations of agents. We have developed an architecture allows distributed multi-core hosts to increase the agent populations which scales almost linearly with the number of processors applied to the problem.
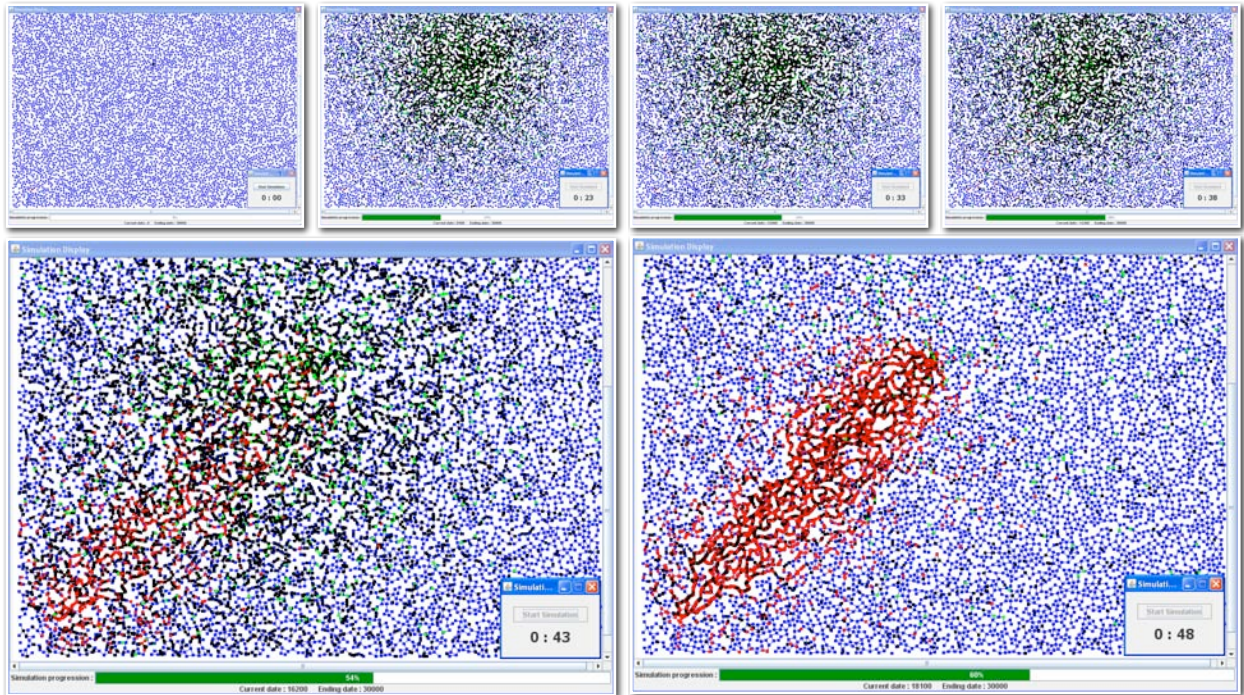
Large graphs can be searched using an agent-based technique called Ant Colony Optimization. Software agents simulate the stochastic trail marking used by biological ants for collaborative food foraging. The top images show random foraging; those at bottom show the transition to social foraging, which results in a near optimal path between the nest and the food source.

Outputs

Papers

Fillinger, I. Hamchi, S. Degré, L. Diduch, T. Rose, J. Fiscus and V. Stanford. Engineering the Pervasive Future: Common Middleware, Research Corpora, and Metrology. IEEE Pervasive Computing Mobile and Ubiquitous Systems. Accepted for Publication 2008.

L. Diduch, M. Hoarau, A. Fillinger, V. Stanford. Synchronization of Data Streams in Distributed Realtime Multimodal Signal Processing Environments on Commodity Hardware. Proceedings of the IEEE International Conference on Multimedia & Expo (ICME), Hannover, Germany, June 2008.

A. Fillinger, L. Diduch, I. Hamchi, M. Hoarau, S. Degré, V. Stanford. The NIST Data Flow System II: A Standardized Interface for Distributed Multimedia Applications. IEEE International Symposium on a World of Wireless, Mobile and MultiMedia Networks (WoWMoM), Newport Beach, California, June 2008 - **Awarded best demonstration**.

L. Diduch R. Mueller, G. Rigoll. A framework for modular signal processing systems with high-performance requirements. Proceedings of the IEEE International Conference on Multimedia & Expo (ICME), Beijing, China, July 2007.

A. Fillinger, L. Diduch, I. Hamchi, S. Degré and V. Stanford. NIST smart data flow system II: speaker localization. Proceedings of the 6th international conference on Information processing in sensor networks (IPSN), Cambridge, Massachusetts, USA, 2007.

Presentations

V. Stanford, Time Series Prediction Forecasting the Future and Understanding the past (PPT). Santa Fe Institute Proceedings on the Studies in the Sciences of Complexity Edited by Andreas Weingend and Neil Gershenfeld. NIST Complex System Program, Perspectives on Standard Benchmark Data, In Quantifying Complex Systems, NIST Complex Sysems Study Group, August 31, 2007.

Software

The NIST Data Flow System II, a generic cross-platform middleware for parallel process and distributed computing.

Distributed processing for multiple instances of Octave or Matlab computational engines for parallel matrix algebra needed in complex system simulations. Provided with the NDFS-II.

Distributed simulation architectures for the study of emergent behaviors in large populations of simple agents that form complex systems. The distributed processing was shown to support very large agent population capable of solving large combinatorial optimization problems. Provided with the NDFS-II.

The ant colony Engine that simulates the behavior that biological ant colonies show in nature when foraging for food and finding efficient paths from the nest to the food sources. Provided with the NDFS-II.

Visualizations

Animations of firefly and pacemaker experiments on NDFS-II testbed L. Diduch
http://www.itl.nist.gov/ITLPrograms/ComplexSystems/Visual.html

Distributed Ant Colony Optimization Simulation is a new method of parallelization based on the NDFS-II middleware. I. Hamchi

Bayesian Modeling Applications to Robot Simulation

Blaza Tomon

Probabilistic Graphical Models have become a popular tool in machine learning. They have been described (Jordan (1999) Learning in Graphical Models, MIT Press) as "a marriage between probability theory and graph theory" and as "a natural tool for dealing with two problems that occur throughout applied mathematics and engineering – uncertainty and complexity". A probabilistic graphical model is a graph whose nodes are random variables and whose arcs indicate conditional dependence assumptions. Directed Graphical Models, also called Bayesian networks, incorpo-

rate causal relationships via these arcs. There is now a rich literature and a set of software tools on which to draw when applying these models to real life problems. In this project we propose to study these models and apply them to model the behavior of complex systems. Specifically, we will use these models to represent the behavior of multiple robots in USARSim (Urban Search And Rescue Simulator) directly, at a high level but realistically,
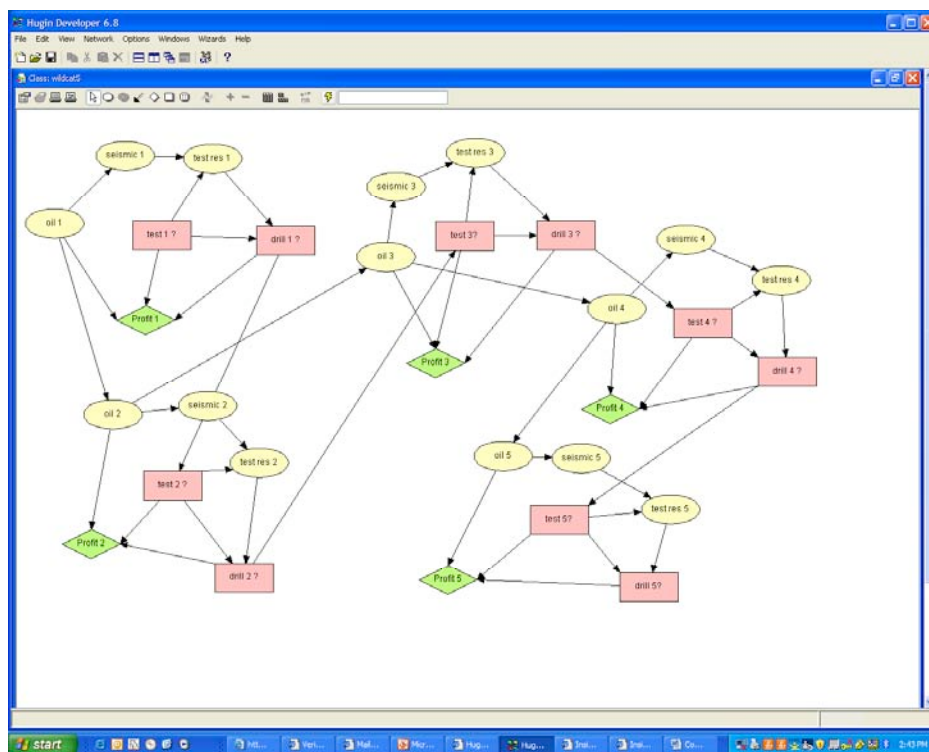
This is a graphical representation of a Bayesian network made up of five robots whose function is to perform one of two actions. They could either perform a seismic test to help assess the likelihood of oil being present or they could drill for oil. The machines are located in a large field where the a priori probability of oil is uniform throughout. It is likely that neighboring positions would have similar ground conditions and thus that the presence of oil at a particular location would increase the probability of oil at neighboring locations. This can be represented by the conditional probability tables of the Oil nodes. This is the origin of "cooperation" of the units of this complex system. An observation in the form of a test result at a single location is propagated through the system, updating via Bayes Theorem the marginal probabilities of oil at all of the locations. This has an effect on the choice of optimal action for many if not all of the units and thus is at the heart of the collective behavior of the system.

Programmatic Activities

"Report writing, like motor-car driving and love-making, is one of those activities which almost every Englishman thinks he can do well without instruction. The results are of course usually abominable." - Tom Margerison 'A Random Walk in Science'

This chapter describes the activities of the Complex Systems Program office. Primarily these are meant to be a "value add" to the technical conducted by the individual projects. In addition these activities are meant to foster and stimulate a complex systems research community within NIST. All information is available at http://www.itl.nist.gov/ITLPrograms/ComplexSystems/

Image of the Month

Each month the program office produces a one page flyer to highlight the work of one particular project. The illustrations and text are intended to convey technical information to the layman and are not meant to be overly technical. March 2009 marks the one year anniversary of these images. All are available at:

http://www.itl.nist.gov/ITLPrograms/ComplexSystems/image.html

Complex Systems Study Group

An approximately bi-weekly series of meetings, open to all, covering topics of interest and relevance to the Complex Systems research community.

Complex Systems Study Group Seminar Series

Every other week the Complex Systems Study Group (CSSG) meets to discuss issues and topics relevant to the wonderful world of Complex Systems. The meetings are deliberately informal and all are welcome. For particular meetings participants make semi-formal presentations, and these are generally made available via links (in the deliverables section of this page) here. In addition every once in a while a formal Complex Systems Lecture Series talk is presented by someone outside of NIST (ComplexTalkPosters).

Coming Events - All talks in Bldg 225/B111 unless otherwise noted
- Just imagine YOUR TALK HERE!

Past Events
- May 6, 2008; 2:00PM - Vladimir Marbukh, Towards Price Based Network Management and Provisioning
- Mar 17, 2008; 2:00PM (note this is a Monday) - Fern Hunt, Mathematical Modeling of TCP/IP protocols
- Feb 26, 2008; 2:00PM - Dan Genin, Metastability and networks
- Dec 12, 1:30PM (AML Lecture Room) - Tim Sauer, George Mason University, Analyzing Time Series from Chaotic System Part 1 and Part 2
- Nov 16, 10:30AM - Vince Stanford, Stephane Degre, Lukas Diduch, Complex Systems Testbed Project - demos of simulations
- Oct 31, 10:30AM - Isabel Beichl and Brian Cloteaux "Counting Spanning Trees"
- Oct 17, 10:30AM - Dan Genin Percolation: Theory and Applications
- Sept 26, 10:30AM - Vince Stanford, Time Series Prediction
- Sept 17, 1:30PM - William Cheswick, AT&T Labs, 40 Years of Internet Security, Are We There Yet?
- July 23, 1:00PM - Abhijit Deshmukh & Eduado Misawa, NSF, overview of NSF complex systems plans
- June 8, Lec Rm B, 1:00PM - Dave Touretzky, Carnegie Mellon Unversity, Cognitive Robotics
- Apr 17 & May 1 - K. Sriram, "Simulation Study of BGP Peering Session Attacks and Their Impacts on Internet Routing Performance"
- Feb 9, 2007 - Roldan Pozo, review of paper Statistical Mechanics of Complex Networks by Albert & Barabasi

Program Overview Brochure
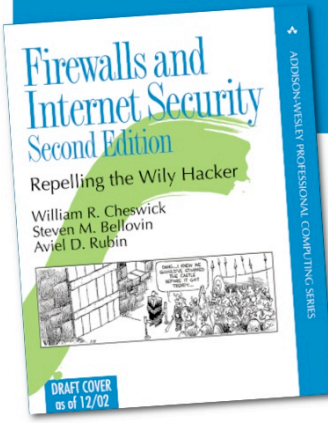
A trifold brochure was produced to be used as "giveaways" to visitors. A thumbnail view of it follows:
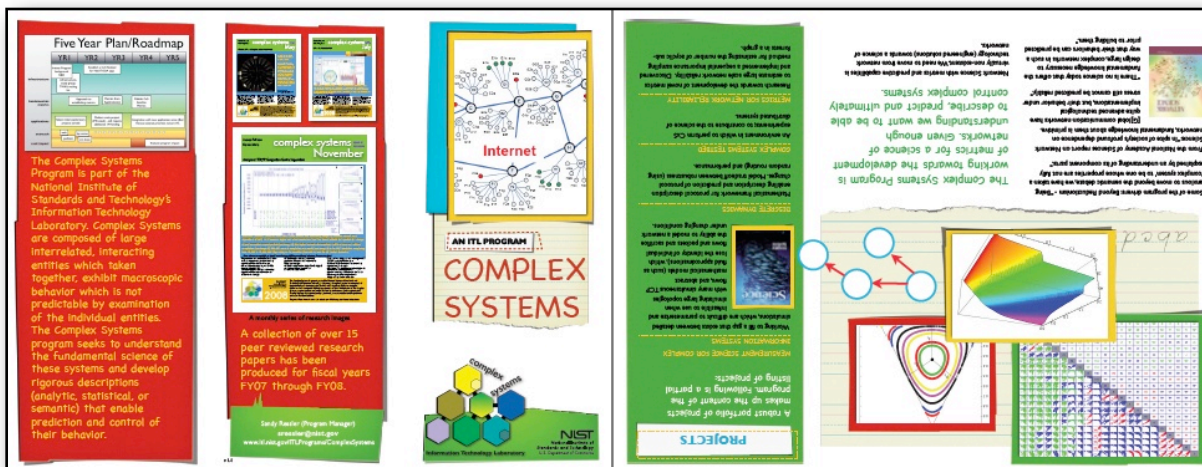
Complex Systems Lecture Series

Occasionally the Program supports outside and significant internal speakers in an effort to keep abreast of relevant research and foster a robust research community. Some of the posters from past lectures follow:

40 Years of Internet Security: Are We There Yet?

William Cheswick
AT&T Research

Monday Sept. 17, 2007
Red Auditorium
1:30PM

Firewalls and Internet Security, Second Edition — Repelling the Wily Hacker. William R. Cheswick, Steven M. Bellovin, Aviel D. Rubin. ADDISON-WESLEY PROFESSIONAL COMPUTING SERIES. DRAFT COVER as of 12/02

The third character ever transmitted on the Arpanet crashed the server. The Morris worm infected thousands of computers in 1988. Now infections are so widespread they are studied using the math of epidemics, or the economics of small countries. Where attacks used to be the domain of desperately anti-social teenagers seeking fame, the nuisances and very real threats today come from shady businesses and organized crime with distinct profit motives, and governments seeking information and advantage.

Saltzer and Schroeder gave classic good security advice in 1975, but Sun and SGI seemed to ignore it in the 1980s and 1990s, and Microsoft certainly ignored it until 2001. Most of the attacks we see today rhyme with those of twenty years ago. Most of these problems (aside from a few DOS viruses) never seem to go away.

We've built walled gardens so vast that it is sometimes unclear which side is outside. We've updated and patched countless billions of times, often replacing old bugs with new.

How bad is it, really? Are we making progress? Which technologies are likely to be technically and economically feasible?

Part of the Complex Systems Lecture Series of the Complex Systems Program of ITL
For more information, contact Sandy Ressler, x3549

Analyzing Time Series From Chaotic Systems

TIM SAUER
PROFESSOR OF MATHEMATICS
DEPT. OF MATHEMATICS, COLLEGE OF SCIENCE
GEORGE MASON UNIVERSITY

CHAOS
AN INTRODUCTION TO DYNAMICAL SYSTEMS

KATHLEEN T. ALLIGOOD
TIM D. SAUER
JAMES A. YORKE

December 12, 2007, 1:30 - 2:30p.m.
Bldg. 215-AML, Room C103-C106

The discovery of the prevalence of chaos in simple dynamical systems in the 1970's led to an array of new techniques for the analysis of data from complex systems. Methods for the approximation of fractal dimension, Lyapunov exponents, entropy, and other mathematical invariants were developed, and new approaches to the standard time series analysis goals of prediction, noise reduction, and control were proposed.

The talk will begin with a historical survey of these ideas, developed during the 1980's and 1990's, and then discuss new directions that have been explored since then. In particular, we will focus on various open problems in the inference of network dynamics from measured data.

(Part of the Lecture Series of the Complex Systems Program of ITL)
For more information, contact Sandy Ressler at x3549 or sandy.ressler@nist.gov

Dave Touretzky, Carnegie Mellon, Cognitive Robotics

TWIKI
A standard TWIKI is used for the Program as a repository of knowledge and method for announcements (internal only).

Milestones and Plans

"Anyone who says businessmen deal in facts, not fiction, has never read old five-year projections." -Malcom Forbes

Roadmap

Designed as a five year program we are following a development roadmap.

Potential Future Opportunities

• Intersection with predictive modeling
• Intersection with multi-scale modeling
• Specific metrics for particular National Needs
• Trusted publicly available data
• Mathematical descriptions of emergence and relationship to inverse and ill posed type problems.
• Are catastrophic failures the result of emergent behaviors?

Deliverables Data

"Underpromise; overdeliver" - Tom Peters

Papers

2008

1. F. Hunt and V. Marbukh, Dynamic Routing and Congestion Control Through Random Assignment of Routes, in International Multi-Conference on Engineering and Technological Innovation, July 2008, pp. 161-164.

2. I. Beichl and B. Cloteaux, Generating Network Models Using the S-Metric, in Proceedings of the International Conference on Modeling, Simulation and Visualization Methods, Las Vegas, NV, July 2008, pp. 159-164.

3. I. Beichl and B. Cloteaux, Measuring the Effectiveness of the s-Metric to Produce Better Network Models, in Proceedings of the Winter Simulation Conference, Miami FL, Dec 2008, pp. 1020-1028.

4. K. Mills and C. Dabrowski, Can Economics-based Resource Allocation Prove Effective in a Computation Marketplace? in Journal of Grid Computing Special Issue on Grid Economics and Business Models. DOI (10.1007/s10723-007-9094-4)

5. D. Genin and V. Marbukh, Toward Understanding of Metastability in Large-Scale Loss Networks with Mobile Users: Emergence and Implications for Performance, Proceedings of 12th World Multiconference on Systems, Orlando FL, June 2008.

6. V. Marbukh and K. Mills, Demand Pricing & Resource Allocation in Marketbased Compute Grids: A Model and Initial Results in Proceedings of ICN 2008, April 2008.

7. C. Dabrowski, Reliability in Grid Computing Systems. submitted to the journal Concurrency and Computation: Practice and Experience as part of a special issue on OGF work.

8. V. Marbukh, Can TCP Metastability Explain Cascading Failures and Justify Flow Admission Control in the Internet? in Proceedings of the 15th International Conference on Telecommunications (ICT'08).

9. I. Fillinger, S. Hamchi, L. Degré, L. Diduch, T. Rose, J. Fiscus and V. Stanford. Engineering the Pervasive Future: Common Middleware, Research Corpora, and Metrology. IEEE Pervasive Computing Mobile and Ubiquitous Systems. Accepted for Publication 2008.

10. L. Diduch, M. Hoarau, A. Fillinger, V. Stanford. Synchronization of Data Streams in Distributed Realtime Multimodal Signal Processing Environments on Commodity Hardware. Proceedings of the IEEE International Conference on Multimedia & Expo (ICME), Hannover, Germany, June 2008.

11. A. Fillinger, L. Diduch, I. Hamchi, M. Hoarau, S. Degré, V. Stanford. The NIST Data Flow System II: A Standardized Interface for Distributed Multimedia Applications. IEEE International Symposium on a World of Wireless, Mobile and MultiMedia Networks (WoWMoM), Newport Beach, California, June 2008 - **Awarded best demonstration**.

2007

1. V. Marbukh, Towards Understanding of Complex Communication Networks: Performance, Phase Transitions & Control to appear in a special issue of Sigmetrics "Performance Evaluation Review" DRAFT version MAMA07, June 13.

2. V. Marbukh and K. Mills, On Maximizing Provider Revenue in Market-based Compute Grids in Proceedings of ICNS07, June 2007.

3. K. Mills and C. Dabrowski, Investigating Global Behavior in Computing Grids, Self-Organizing Systems, Lecture Notes in Computer Science, Volume 4124 ISBN 978-3-540-37658-3, pp. 120-136, Oct 2007.

4. K. Mills, A Brief Survey of Self-Organization in Wireless Sensor Networks, Wireless Communications and Mobile Computing, Wiley Interscience, Vol. 7, No. 7, (pages 823-834) October 2007.

5. V. Marbukh, S. Klink, Decentralized control of large-scale networks as a game with local interactions: cross-layer TCP/IP optimization, ValueTools, 2nd International Conference on Performance Evaluation Methodologies and Tools, Nance France, Oct 2007.

6. V. Marbukh, Utility Maximization for Resolving Throughput/Reliability Trade-offs in an Unreliable Network with Multipath Routing, ValueTools, 2nd International Conference on Performance Evaluation Methodologies and Tools, Nance France, Oct 2007.

7. V. Marbukh, Metastability of fair bandwidth sharing under fluctuating demand and necessity of flow admission control in Electronics Letters, 13th September 2007, Vol 43, No 19.

8. V. Marbukh, Fair bandwidth sharing under flow arrivals/departures:effect of retransmissions on stability and performance, in ACM SIGMETRICS Performance Evaluation Review, Volume 35 Issue 2, September 2007.

9. C. Dabrowski, Investigating Resource Allocation in a Standards-Based Grid Compute Economy. National Institute of Standards and Technology: Gaithersburg, MD. Interagency Report 7463, November 2007.

10. L. Diduch R. Mueller, G. Rigoll. A framework for modular signal processing systems with high-performance requirements. Proceedings of the IEEE International Conference on Multimedia & Expo (ICME), Beijing, China, July 2007.

11. A. Fillinger, L. Diduch, I. Hamchi, S. Degré and V. Stanford. NIST smart data flow system II: speaker localization. Proceedings of the 6th international conference on Information processing in sensor networks (IPSN), Cambridge, Massachusetts, USA, 2007.

Presentations
1. R. Pozo, review of paper Statistical Mechanics of Complex Networks by Albert & Barabasi to Complex Systems Study Group, Feb 9, 2007 -

2. V. Stanford, Time Series Prediction Forecasting the Future and Understanding the past. Santa Fe Institute Proceedings on the Studies in the Sciences of Complexity. NIST Complex System Program, Perspectives on Standard Benchmark Data, In Quantifying Complex Systems, internal presentation to CxS Study Group. August 31, 2007.

3. V. Marbukh, Towards Price Based Network Management and Provisioning , May 6, 2008 to Comples Systems Study Group

4.   J. Filliben, Sensitivity Analysis Methodology for a Complex System Computational Model, 39th Symposium on the Interface: Computing Science and Statistics, Philadelphia, PA, May 26, 2007

5. D. Genin, Percolation: Theory and Applications, internal presentation to Complex Systems Study Group, Oct 17, 2007.

6.   F. Hunt, Mathematical Modeling of TCP/IP protocols , May 17, 2008 to Complex Systems Study Group.

7. D. Genin, Metastability and networks , Feb 26, 2008 to Complex Systems Study Group

Software
1. Flexi-Cluster: A Simulator for a Single Compute Cluster, V. Marbukh and K. Mills.

2. MesoNet: A Mesoscopic Simulation Model of a Router-Level Internet-like Network, K. Mills

3. MesoNetHS adds six congestion-control algorithms for networks with high bandwidth-delay products; algorithms include: binary-increase control (BIC), compound TCP (CTCP), fast AQM scalable (FAST) TCP, H-TCP, high-speed TCP (HS TCP) and Scalable TCP (STCP), K. Mills, E. Schwartz and J. Yuan.

4. EconoGrid: A detailed Simulation Model of a Standards-based Grid Compute Economy, C. Dabrowski and K. Mills.

5. MesoGrid: a mesoscopic scale simulation model of a computational grid economy, C. Dabrowski and K. Mills.

6. NGraph++: a simplified graph library for the analysis of complex networks based on ANSI/ISO C++. Includes basic graph operations (construction, intersections, subgraphs) and algorithms for computing degree distribution and clustering coefficients. R. Pozo

7. DiVisa: a multi-dimensional visualization tool. Can read any kind of data (simulation, statistics, text or numeric, etc.), and converters have been implemented to read several data formats without need for reformatting.

8. The NIST Data Flow System II, a generic cross-platform middleware for parallel process and distributed computing.

9. Distributed processing for multiple instances of Octave or Matlab computational engines for parallel matrix algebra needed in complex system simulations. Provided with the NDFS-II.

10. Distributed simulation architectures for the study of emergent behaviors in large populations of simple agents that form complex systems. The distributed processing was shown to support very large agent population capable of solving large combinatorial optimization problems. Provided with the NDFS-II.

11. The ant colony Engine that simulates the behavior that biological ant colonies show in nature when foraging for food and finding efficient paths from the nest to the food sources. Provided with the NDFS-II.

12. B. Cloteaux , A C++ implementation of a sequential importance sampling algorithm, created by us, for estimating the number of spanning trees in a graph, Sept. 2007.

13. B. Cloteaux , A C++ implementation of various network metrics including minimum vertex cover (MC). Uses several of the kernelization techniques of Langston et al. The kernelization method involves taking a graph G and producing a smaller graph G', the kernel, with a value k such that MC(G) = MC(G')+k. For the smaller G' we solve for minimum vertex cover exactly, Feb. 2008.

14. B. Cloteaux , A program for generating random graphs with a given degree sequence. Took code written in R by Joseph Blitzstein at Harvard, converting it to C++, and then optimizing it with the use of novel datastructures. These datastructures allowed the C++ program to run much larger cases than the original program. In particular, it could then handle the Autonomous Systems data from the UCLA database. Software has been released to J. Blitzstein (Harvard), C. Priebe (Johns Hopkins), J. Devinney (IDA Center for Computing Sciences), March 2008.

15. B. Cloteaux , Software to generate graphs with specified s values. These used the Monte Carlo technique of threshold acceptance, May 2008.

16. B. Cloteaux , Software to calculate the maximum s value for a set of graphs. This includes a new approximation technique we developed, namely a deterministic version of Blitzstein's algorithm to produce the upper bound, May 2008.

17. B. Cloteaux , Software to enumerate all graphs with a given degree sequence.Done for illustration purposes for small graphs, May 2008.

18. B. Cloteaux , C++ implementations of Wilson's algorithm to uniformly sample spanning trees in graphs and Welsh's Monte Carlo algorithm for counting bases of a frequent matroid, Aug. 2008.

19. B. Cloteaux , Software implementing our algorithm, based on sequential importance sampling, to approximate the coefficients of the reliability polynomial of a graph, Sept 2008.

20. B. Cloteaux , A C++ implementation of Colbourn's algorithm to estimate the coefficients of the reliability polynomial, Sept 2008.

21. B. Cloteaux , Software to estimate specific values of the reliability polynomial. Uses Chernoff bounds to determine convergence. It produces a graph of connected probability vs edge probability, Sept 2008.

Visualizations
1. Animation of simulation of Abilene network. C. Houard, J. Hagedorn

2. Animations of firefly and pacemaker experiments on NDFS-II testbed L. Diduch

3. Distributed Ant Colony Optimization Simulation is a new method of parallelization based on the NDFS-II middleware. I. Hamchi