Progress Report on the LTS/UMIACS Contract
RFP:MDA904-02-R-0151 – SOW: R4-02-0001.1
September 30, 2004

1. Active Network Management (Mark Shayman, Samrat Bhattacharjee, Steve Marcus, and Richard La)

1.1 Overlay-based Security Services

This project represents work under Tasks 1 and 3. (S. Bhattacharjee, R. La, and M. Shayman)

We have developed an overlay based distributed DDoS detection system which is deployed within an AS and detects TCP based DoS attacks that originate in the AS by monitoring the network traffic. The detection system has several "monitors" which are co-located with the routers in the AS. These monitors poll the bidirectional traffic transiting through the co-located router with a chosen probability and maintain aggregate statistics. Using these statistics, each monitor tests if the flow aggregates conform to TCP specification (with respect to number of outgoing and incoming packets). The distributed architecture of our detection system gives the traffic profile from different vantage points and the detection system correlates these traffic profiles to single out DoS attacks originating within the AS.

We have used a packet level simulator to test our detection scheme and performed extensive simulations using real Internet traffic traces and synthetic attack traffic. We have undertaken an extensive study of the system configuration parameters, and under the constraints of state and processing overheads, evaluated the optimal values for sensitivity and detection time. We have evaluated the detection system in an AS with a single border router under various attack scenarios. Our results show that with very little state (40 counters at each router), with about 10% packet polling rate, our scheme has near perfect performance (it detects all attacks, and signals less than one false positive on average in each run). The attacks are detected within time inversely proportional to the attack source rate. The system has almost zero bandwidth overhead, and as expected, performs better with more severe attacks, and with higher sampling rates.

A much more sophisticated approach is required when the traffic in the AS is asymmetric with respect to the AS egresses (i.e., outgoing traffic goes through one router and incoming traffic for that same connection comes via a different router). This is likely in multi-homed ASs, and can lead to both false positives (asymmetric flows look like attacks), and false negatives (asymmetric flows mask real attacks, which are, by definition, asymmetric). We have extended our scheme to handle asymmetric flows, and our results show that asymmetric flows can be handled with the same amount of state and sampling rates, with slightly higher bandwidth overhead. The main tradeoff is that asymmetric flows take longer to detect. (We

believe this is a fundamental property of any stateless detection system.)

1.2 Overlay-based Traffic Engineering

This project represents work under Tasks 1 and 3. (S. Bhattacharjee, R. La, M. Shayman)

The focus of our work has been the traffic mapping (load balancing) problem. Given a certain source-destination traffic matrix, we are interested in finding a traffic assignment onto pre-established paths such that the overall system performance is optimized with respect to a selected performance measure (*e.g.*, sum of link costs throughout the network). Multiple paths between source and destination pairs are established using an overlay architecture. The paths from source to overlay nodes and from overlay to destinations are calculated according to min-hop principle. However, our basic approach can be directly applied to other types of networks, such as MPLS based networks.

We have established a distributed optimal routing algorithm based on stochastic approximation theory. We derive our *measurement based algorithm* from the idea of simultaneous perturbation stochastic approximation (SPSA). Optimal load balancing is achieved by running SPSA based stochastic approximation algorithm for each source-destination pair in an asynchronous and distributed manner and each source-destination pair relies only on local state information (*e.g.*, sum of the link costs only along the paths of the given source-destination pair). Use of SPSA allows us to greatly reduce the number of measurements required for estimating the gradient, while at the same time we achieve approximately the same level of accuracy as the classical finite differences method at each iteration. By reducing the *number of measurements*, we obtain a better overall convergence rate. (This is because each measurement requires a non-negligible amount of time in a networked environment.)

In the second phase of our work we have generalized our algorithm to load balance the traffic source with multicast nature. The algorithm is able to converge under different network models, where each model reflects a different set of assumptions about the multicasting capabilities of the network. Similar to IP multicast, most of the proposed solutions assume that multicast topology needs to be a tree from source to receivers. However, this assumption results in inefficient use of network resources as the load in the network cannot be balanced. We address the optimal multipath multicast routing problem in a more general framework than having multiple distribution trees. We have considered different network models with different functionalities. With this generalized framework, our goal was to examine the benefits obtained by the addition of new capabilities to the network beyond basic operations such as storing and forwarding. In particular, we first looked at the problem under the traditional network model without any IP multicasting functionality where multiple paths are established using a limited number of (application-layer) overlay nodes. Next, we evaluated the performance in a network model with multiple distribution trees. Finally, we have relaxed the usual assumption that from a given multicast tree each receiver gets

multicast packets at the same rate. Intuitively, relaxing this assumption may seem to create more problems than benefits. This is due to the fact that it potentially creates a complex bookkeeping problem since a source node has to make sure each receiver gets a distinct set of packets from different trees while satisfying the rate constraints along each tree. However, using a specific source coding called Digital Fountain codes, we overcome this problem in an efficient way, which gives an opportunity to observe the potential benefits of having different receiver rates on a multicast distribution tree. Our results show that the network actually benefits from such a setting by being able to balance the traffic load efficiently and achieving a better performance in terms of end-to-end throughput. Simulation results have shown that while IP multicasting functionality is highly essential for better performance, additional functionalities, such as probabilistic splitting of multicast traffic along a multicast tree, provide only limited benefits in relation to their required complexity.

An important issue is that the SPSA algorithm has the almost sure convergence property to the optimal solution provided that the step size parameter diminishes with the number of iterations. However, such a policy limits the practicality of applying SPSA under dynamic network conditions as the algorithm will not be able to react to the changes appropriately once the step size parameter becomes small. As a result, in practice, we have to reset the step size value after a certain time interval to ensure that the algorithm is able to react to network dynamics appropriately. An obvious alternative is the use of a constant step size. Even though it is difficult to obtain almost sure convergence in the constant step-size case, weak convergence (*i.e.*, convergence in distribution) which can be interpreted as convergence to a neighborhood of the optimal operating point(s) can be shown under appropriate conditions. Since the performance of the system near the optimal operating point(s) may be comparable to that of the optimal solution(s) in a network problem, the performance degradation, if there is any, due to a constant step size may not be significant. As a follow-up to the existing work we have shown the weak convergence of the optimal routing algorithms under constant step size policy for both unicast and multicast sources. Simulation studies have shown that the aforementioned performance degradation is negligible as expected.

The overlay architecture plays a key role in our study. By selecting the number, location and the connectivity of the overlay nodes we establish multiple paths between source- destination pairs over which we run our optimal routing algorithms. However, the performance of the routing algorithms is limited by the selection of overlay topology. Hence, it is of profound importance to optimize the overlay topology to improve the performance of the routing algorithms. As a future work, we propose to investigate possibilities to optimize the overlay topology. Our goal will be to establish topology control architecture with an offline component optimizing the overlay topology given the information at hand such as estimated traffic demands between SD pairs and an online component that will update the existing topology in accordance with major changes observed in network. Since the time scale of this topology control algorithm would be slower than that of the routing algorithms, methods such as simulated annealing or genetic algorithms can be used to attack this problem.

1.3 Multiclass Traffic Engineering Using Selective Overprovisioning

This project represents work under Tasks 1 and 3. (S. Bhattacharjee, R. La, and M. Shayman)

We are continuing our work on developing Differentiated Traffic Engineering algorithms for QoS provisioning. This platform is applicable in networks capable of source-based multi-path routing. Examples of such networks are MPLS network or IP networks with overlay nodes. The goal is to take advantage of multiple paths and simplify the packet level QoS enforcement mechanism such as class based queueing and scheduling in the core routers. More specifically, we intend to develop network architectures and traffic engineering algorithms that distribute packets at the edge routers between different paths according to their QoS requirement. Our objective is to overprovision links that carry QoS sensitive data such that a simple FIFO queueing is sufficient to provide required QoS.

In the proposed architecture, we consider two classes of traffic, real time (class1) and best effort (class 2) traffic. Every source-destination class 1 and 2 traffic demand is known. The DTE structure distributes class 1 and 2 traffic between the network paths such that utilization of links that carry any class 1 traffic is below a threshold value, say 0.5. However, other links that only carry best effort traffic do not need to be overprovisioned and their utilization could be as high as 1.

We have formulated this problem as a nonlinear optimization problem. The optimization problem is nonconvex and hence, common distributed traffic engineering schemes based on gradient-projection method are not directly applicable. We have so far, identified and worked on two complementary traffic engineering problems for the DTE architecture.

The first problem is Path to Class Assignment (PCA). Here, we assume that a fixed set of paths are selected and the problem is to assign paths to classes of service. For each path we have two options; either it is assigned to class 2 traffic or it is assigned to both classes. From the path assignments, we can obtain the link assignments, and consequently the maximum permissible utilization (virtual capacity) of links. The objective is to minimize total link cost, which is a function of link load and virtual capacity. We have designed an algorithm based on the well known simulated annealing optimization method. Simulated annealing is a randomized optimization method that can find the global minimum of a nonconvex function. We have simulated and studied the characteristics of the PCA method and they are very promising. The details of the DTE architecture, PCA algorithm and the simulation results are given in an Infocom submission.

The second problem is Path Selection and Routing (PSR). Here, we select a set of paths to accommodate a specified class 1 and 2 traffic demand. Currently, we are

actively working on developing an algorithm for this problem. Our general approach is to define an appropriate link cost (length) function for the class 1 and 2 traffic demands. Then, we use the K-shortest path algorithm to find an appropriate set of paths for optimization. The path set is then passed to the PCA module for optimal path to class assignment.

We plan to complete the PSR algorithm first, and then to work on integration of the PCA and PSR algorithms. Integration has its own challenges; we have to come up with an effective and stable set of interaction mechanisms and rules. For instance, we have to specify when to activate the PCA and PSR modules and ask for an update on the path assignment and path selection respectively; how to estimate and compute the traffic demand matrices that are used in PCA and PSR; how to avoid frequent updates and changes in the path set and path assignment.

1.4 Markov Decision Modeling for Integrated MPLS/WDM Traffic Engineering

This project represents work under Tasks 2, 3, and 4. (R. La, S. Marcus, and M. Shayman.

The results so far have been submitted to Infocom 2005. After this submittal, we have been investigating whether a simplified model can yield effective control of the reconfiguration process.

We have reduced the model to one time scale only in order to be able to compare the performance of our algorithm to that of other algorithms in the literature, including one published by E. Modiano and coworkers. If the network traffic changes smoothly and slowly, their algorithm is able to follow the changes in traffic by taking a sequence of branch exchanges.

We have compared our algorithm to Modiano's algorithm in both a static mode and dynamic mode. In static mode we start from a random topology in presence of static traffic. The goal is to find a sequence of branch exchanges that transitions the topology to near optimal topology. We have shown that our algorithm performs better in two ways. It has a better cost criterion and it looks ahead into possible branch exchanges and the performance of the network after those possible branch exchanges. This results in better performance than the algorithm by Modiano that looks only one step ahead.

In the dynamic case, Modiano considers smooth and slowly changing traffic. In addition to that, we consider real traffic collected from Abilene network. We have obtained real traffic matrices from Abilene that enable us to find a trend in traffic (average over a long time). This data and another set of data (not included in the long time average) are fed into the simulation program to find the response of our algorithm to this real world traffic. We are attempting to show that by using the historical daily trend to predict future traffic demands, our algorithm is able to better adjust to the changes in traffic and step-by-step transition to better topologies.

We have moved from a delay oriented reward function to a cost function that penalizes the dropped calls and the high utilization in the lightpaths.

We are also experimenting with the number of steps the algorithm looks into the future. Since looking more steps into the future will require more processing time, we want to find at what point we do not gain by looking further into the future. This will help us in optimizing the running time of the algorithm.

Publications

T. Guven, C. Kommareddy, R. J. La, M. A. Shayman and B. Bhattacharjee, Measurement based optimal multi-path routing, *IEEE Infocom,* Hong Kong, March 2004.

K. Lee, M. Kalantari and M. A. Shayman, Routing instability in the BGP protocol, *Conference on Information Sciences and Systems,* Princeton University, March 2004.

K. Lee and M. A. Shayman, Multicasting extensions to traffic engineering in MPLS networks, *Conference on Information Sciences and Systems,* Princeton University, March 2004.

K. Lee and M. A. Shayman, Optical network design with optical constraints in multi-hop WDM mesh networks, *International Conference on Computer Communications and Networks,* Chicago, Illinois, October 2004.

T. Guven, R. J. La, M. A. Shayman and B. Bhattacharjee, Measurement-based multicast on overlay architecture, *IEEE Infocom,* Miami, Florida, March 2005, submitted.

V. Tabatabaee, B. Bhattacharjee, R. J. La and M. A. Shayman, Differentiated traffic engineering for QoS provisioning, *IEEE Infocom,* Miami, Florida, March 2005, submitted.

P. Fard, R. J. La, K. Lee, S. Marcus and M. A. Shayman, Reconfiguration of MPLS/WDM networks using simulation-based Markov decision processes, *IEEE Infocom,* Miami, Florida, March 2005, submitted.

A. Rawat and M. A. Shayman, Preventing persistent oscillations and loops in IBGP configuration with route reflection, *IEEE Infocom,* Miami, Florida, March 2005, submitted.

2. Active Systems Security Management (William Arbaugh and Virgil Gligor)

We presented the Copilot paper at USENIX Security in San Diego in August. We had numerous discussions with other attendees about our approach and results—all positive.

Since that time, we have focused our attention on two major items. The first is to determine additional capabilities of Copilot beyond watching kernel text and data. The second is an examination of architectural guidelines for building operating systems that are mandatory access control (MAC) "friendly", i.e. what can operating system designers do to make life easier for those wishing to implement MAC.

There are two major research problems that we feel Copilot can provide a unique capability. The first is implementing two man control on administrative access, and the second is applying compiler technology to improve the ability of Copilot to monitor the data space of the kernel.

Copilot, by virtue of the fact that it can read host memory, can identify and parse the process table. By examining the process table, Copilot can determine the privilege level of each runnable process and determine if a process is authorized to run at administrator or root privilege. Two man control can be implemented simply by the central monitoring station informing Copilot that a particular process is approved to run before the process actually runs.

We believe that Copilot can not only detect malicious modifications to text and data, but when properly instrumented with additional type information Copilot can detect run-time logic and pointer errors. The general idea is to take type and map information from the compiler and linker to produce a template for Copilot to monitor. The extent to which we can do this given the limited typing of the C language remains to be seen. However, we are hopeful that we can convert a completely manual process at the current time to a fully automated process.

3. Wireless Networking (William Arbaugh, Ashok Agrawala, A. Udaya Shankar, and Joseph Thomas)

3.1 Ubiquitous Wireless Interworking Test-bed (UWIN) (William Arbaugh, Ashok Agrawala, and Udaya Shankar)

Over the past few months we have focused on the architecture for the test-bed. We deployed several test nodes on top of campus buildings with limited success. One of the nodes was damaged due to flooding, and the batteries of the other nodes could not sustain the nodes more than a few weeks. We next attempted to collect power information using an instrumented power regulator. Unfortunately, the power regulator itself drew more power and the batteries died sooner than previously. As a result of these experiments, we are currently investigating an even lower power

solution than the current Soekris hardware. This, however, has set us back in schedule since our goal was to have a hardware base to begin experiments. Until we have the hardware base in place, we can not begin experiments on the scale we need to investigate task #1 and #2 in the SOW.

We have been working on how the hardware base provided by the Trusted Computing Group (TCG) can assist in task #3.The general idea is that a node's identity will be based on a fingerprint of the software running on it combined with a random number selected from a sparse space. This will permit other nodes with the mesh to determine if a new participant is running a similar (or similar enough) software configuration to participate in the mesh. This will essentially form the first leg of trust. The second, and final leg, will be similar to a reputation system where a node is initially only granted limited privileges which are expanded as the node faithfully participates in the mesh.

3.2 Efficient IP-Based UMTS Networks (J. Thomas)

Based on the feedback obtained in program reviews, attention has been focused on the following two sub-tasks.

Sub-Task 1: Session Initiation Protocol in Mobile IP Environments

The session initiation protocol (SIP) is used to establish and manage stream-traffic calls, namely to determine source addresses, add new streams, add new participants, transfer calls, recognize QoS requirements etc. The SIP architecture includes a SIP-registrar and proxy and redirect servers. A SIP user is assigned a URI (universal resource identifier) which serves as an "address" for initiating calls; the user's IP address (which it registers with the SIP registrar, during initialization) is the address to which messages are routed. This allows a user to communicate with multiple interfaces. If SIP is used over MIPv4, packets for the mobile user are routed first to its home network and subsequently tunneled to its foreign network. Over an MIPv6 protocol, if the user's IP address that is registered with the SIP registrar is its home address, then tunneling is necessary, whereas if its care-of address is used, its destinations must re-invited at every change of this care-of address implying call disruptions and re-establishments. To circumvent these undesirable choices, the present effort consists of integrating SIP with suitable modifications over an MIPv6 environment (on the HUT-Dynamics platform). This is expected to see results by the summer of 2005.

Sub-Task 2: Cross-Layer Load Balancing and Routing in Ad Hoc Extensions to Structured Networks

In conjunction with the investigator's ongoing work on the tradeoffs between energy, throughput, and delay in structured and unstructured networks, specific attention is being devoted to the problem of load balancing and routing in ad hoc extensions to structured networks. The performance of a directed-graph-based load balancing and routing algorithm for this environment is being studied via simulations and analyses. In essence, the environment is modeled as a time-slotted system with all nodes being in one of three states, namely busy, available, or asleep. Nodes transition between these three states depending on

the history of recent traffic at the node using suitable thresholds. This is aimed at conserving battery power. In every time-slot a busy node broadcasts an indicator of its presence. Depending on the received power-levels of the acknowledgements from neighboring nodes weighted directed edges are dynamically created to these neighbors, i.e. if the acknowledgement received from a neighbor exceeds a stipulated signal-to-interference-plus-noise ratio threshold a directed edge is created with its weight specifying the SINR for that link. Routing and load-balancing decisions are implemented with the objective of meeting throughput, delay, and energy constraints. The scope of this work includes many special cases of interest ranging from quasi-structured subnets to fully ad hoc subnets. This effort is expected to see results beginning early in the spring of 2005.

4. The Economics of Communications/Networking Technology (Larry Gordon, Martin Loeb, Joseph Bailey, and S. Raghavan)

4.1 The Business Case Development and the Economic Impact (L. Gordon and M. Loeb)

Task 1 – Business Case Development

At the end of last quarter, we gave a presentation to Dr. William Semancik, Dan Foerter, and Justin McCann at LTS on our examination of how to organize and manage research activities. This task focuses on an economic methodology for determining under what conditions research should be done in-house and under what conditions in should be outsourced. The framework we developed focuses on three types of concerns: (1) organizational (2) neoclassical economic and (3) life-cycle process. Additionally, our analysis brought in a number of other concerns such as preserving confidentiality, balancing the portfolio of projects, and building upon the organization's core competency. Given the feedback from that meeting, during this quarter, we have refined the framework and identified characteristics of projects that fall in either the "make" or "buy" zones. We are now writing a draft of a paper based on this work.

During this quarter, we also revised and resubmitted our paper, entitled "Budgeting Process for Information Security Expenditures: Empirical Evidence," for publication to *Communications of the ACM*. We are delighted to report that **this paper has just been accepted for publication**.

Another aspect of this task is to analyze how to make the best case internally for funding of information security activities and projects. The goal is to develop a framework that information security officers can use to effectively compete for internal funds. We completed a literature review and are in the midst of developing an outline for a paper.

Related Publications and Papers

Gordon, Lawrence A., and Martin P. Loeb, "Budgeting Process for Information Security Expenditures: Empirical Evidence," *Communications of the ACM*, forthcoming.

Bodin, L., L. A., Gordon, and M. P. Loeb, "Evaluating Information Security Investments using the Analytic Hierarchy Process," *Communications of the ACM*, forthcoming.

Gordon, Lawrence A. and Martin P. Loeb, "The Economics of Investment in Information Security," *ACM Transactions on Information and System Security*, November 2002, pp. 438-457. (reprinted in Economics of Information Security, 2004, Springer, Camp and Lewis, eds.)

Gordon, Lawrence A. and Martin P. Loeb, " Return on Information Security Investments: Myths vs. Reality," *Strategic Finance*, November 2002, pp. 26-31. (awarded Certificate of Merit in June 2003 by the Institute of Management Accountants).

Gordon, Lawrence A., Martin P. Loeb, and William Lucyshyn, "Economic Aspects of Controlling Capital Investments in Cyberspace Security for Critical Infrastructure Assets," Proceeding of the 2nd Annual Workshop on Economics and Information Security, College Park, Maryland, May 2003.

Gordon, Lawrence A., Martin P. Loeb, and William Lucyshyn, "Information Security Expenditures and Real Options: A Wait-and-See Approach," *Computer Security Journal*, Vol 19, No. 2, 2003.

Task 3 – Economic Impact of Government, Industry, Academic Partnerships

This task addresses questions of economic welfare associated with the sharing of information by a government-corporate-academic partnership and the role of the government in facilitating private initiatives to enhance information security activities. To this end, we are completing a research note examining the welfare implications of government mandates to increase public disclosure of cybersecurity activities by publicly held firms. Provisions of the Sarbanes-Oxley Act of 2002, some of which have not yet gone into effect, will have an impact on cybersecurity disclosure. We have now examined the public disclosure of cybersecurity activities by telecommunications firms prior to the implementation of the Act. This will provide a baseline for further empirical study, once post-Act data becomes available.

Related Publication

Gordon, Lawrence A., Martin P. Loeb, and William Lucyshyn, "Sharing Information on Computer Systems Security: An Economic Analysis," *Journal of Accounting and Public Policy*, Vol 22, No. 6, 2003, pp. 561-485.

Task 4 – The Economic Effect of Information Security Breaches

This task calls for quantifying the economic effect of information security breaches on individual companies and the determining the spillover cost to other parts of society. As noted in earlier quarterly reports, several papers were either published or accepted related to this task (see below list of related publications). One of these papers looked at the stock market reaction to security breaches within U.S. Corporations (see the paper listed below published in the *Journal of Computer Security*). During this quarter, we have continued our work related to expanding this empirical study.

Related Publications and Papers

Campbell, K., L.A. Gordon, M. P. Loeb, and L. Zhou "The Economic Cost of Publicly Announced Information Security Breaches: Empirical Evidence from the Stock Market," *Journal of Computer Security*, Vol. 11, No. 3, 2003.

Gordon, Lawrence A. and Martin P. Loeb, "Expenditures on Competitor Analysis and Information Security: A Management Accounting Perspective," Chapter in Management Accounting in the Digital Economy (Oxford University Press), A. Bhimini (ed), 2003, pp. 95-111.

Gordon, Lawrence A., Martin P. Loeb and Tashfeen Sohail, "A Framework for Using Insurance for Cyber Risk Management," *Communications of the ACM*, March 2003, pp. 81-85.

Conclusion

Finally, some of our activities during the quarter relate to all of the above tasks and indicate the impact of our sponsored research. As founding members of the organizing committee of the Workshop on Economics and Information Security, we are again involved in planning for the annual Workshop. The next (fourth) Workshop is scheduled for June 2-4, 2005, at Harvard University's Kennedy School of Government (recall, the second Workshop was held at Maryland). In addition, L. Gordon presented a summary of the research we have done under the LTS sponsorship on September 29 at the 2004 State of Maryland Information Technology Security and Privacy Conference held at Johns Hopkins Applied Physics Laboratory. Finally, we have continued to broaden our contacts with individuals form various government agencies, universities, journals (i.e., reviewing work related to our sponsored research), and corporations concerning our research on economic aspects of information security research.

4.2 Internet Pricing and Network Management (J. Bailey and S. Raghavan)

Task 2 – Research in the Impact of Pricing Strategies

We continue to improve our publication submission in progress on "Ex-Post Internet Charging." The status of the paper is "revise and resubmit" to ACM Transactions on Internet Technology.

The impact of Internet pricing on the dynamics of Internet competition has yielded some interesting results. With our growing database of Internet Service Providers (ISPs), we are able to examine how technology and pricing strategies have changed over the past several years. Our starting hypothesis that security is an important determinant of ISP success—as defined by survivability in the market—is supported for some security technologies. We had hoped to find a linkage between adoption of security and price that the ISP charges in the market. In this way, we hoped to support the business case for security by explaining that the extra security cost is likely to be passed on to consumers. Unfortunately, our most recent analysis in this past quarter has led us to the conclusion that pricing is not tightly coupled with firm success and security adoption. Therefore, we are now under the impression that security and pricing are both driven by the competitive forces in the market and not a hedonic pricing model. We have just collected September 2004 data on these firms and we plan on confirming our finding soon.

Related Publications

"Ex-Post Internet Charging," J. Bailey, J. Nagel and S. Raghavan. To appear in *Internet Services: The Economics of Quality of Service for Networks, Grids, and Markets*, edited by L. McKnight and J. Wroclawski. MIT Press, 2004.

Task 5 – The Business Case for Wireless Systems

Our empirical analysis of wireless adoption among Internet Service Providers (ISP) indicates a much clearer business case for wireless. In our research, we have found that wireless adoption is increasing over time among surviving firms. From 2002 to 2004, wireless adoption increased from 10.2% to 12.8% among ISPs. However, the 10.2% figure is smaller than the level of adoption among exiting firms (14.2%). These potentially conflicting findings indicate that the business case for wireless systems is a more complex issue. Although the deployment of wireless increases firm performance, there are likely standards and security issues. Those firms that adopt wireless technologies while the standards are still immature may be more likely to exit the market. Furthermore, those firms that adopt wireless without security may be unable to survive in the market. It is this last point that we will be exploring further in the coming months.

Along with doctoral student Robert Day, we continued our study of spectrum auctions and the appropriateness of using CAMBO (our proposed combinatorial

auction framework) for spectrum auctions. In particular, we examined the proxy-auction setting proposed by noted economists Paul Milgrom (Stanford) and Larry Ausubel (Maryland) for combinatorial auctions. This particular auction has significant promise and is being considered for forthcoming FCC and possibly FAA auctions. However, it suffers from slow convergence. We develop a methodology to rapidly obtain the desired "bidder-pareto optimal core outcome" by solving the pricing problem (i.e., determining what amount winning bidders must pay in the auction) using constraint generation. This paper was presented recently to a noted and distinguished group of economists and computer scientists at a workshop in Rutgers University and received very well.

We continue our study of routing and design problems, for reliable/secure communication, in geostationary satellite communications. Satellites provide a secure communication technology that is relatively cheap, accessible everywhere, and difficult to breakdown. Several new geostationary satellite systems have been proposed by industry. Along with doctoral student Ioannis Gamvros we are investigating issues related to the design of satellite communication networks to meet demand over multiple periods of time. Issues such as satellite location, routing of traffic, and combining terrestrial links with satellite links are considered. Over the past quarter (and this year) we have developed a mixed integer programming model for the multiperiod routing problem. Briefly the problem can be stated as follows. We are given the satellite network topology that changes over time. There is also a cost to change the route that a customers' demand follows from period to period (this is because of the cost to repoint the satellite dish, stop communications, and then reestablish secure communications). Given a demand over multiple periods we would like to find the minimum cost routing plan. In particular, we have developed a column generation technique to solve this massive integer programming problem. Our results indicate significant benefits of performing multi-period routing. Additionally, our techniques solve problems an order of magnitude larger than those solved previously. The results of using our techniques should be a better understanding of the costs, and methods to minimize the costs of satellite communication networks. In the future, we also propose to develop models to deal with the inherent uncertainty in future demand in business cases/planning.

***Related Publications**:*

"CAMBO: Combinatorial Auctions using Matrix Bids with Order," R. Day and S. Raghavan, submitted for publication, *Operations Research*.

"Generation and Selection of Core Outcomes in Sealed Bid Combinatroial Auctions," R. Day and S. Raghavan. Presented at DIMACS Workshop on Computational Issues in Auction Design. October 7-8, Rutgers University, New Brunswick, New Jersey.

"The Multi-Level Capacitated Minimum Spanning Tree Problem," I. Gamvros, B. Golden, and S. Raghavan, To appear, *INFORMS Journal on Computing*.

5. Optical Networking (Gary Carter and Joel Morris)

5.1 High Speed Experiments (G. Carter)

We successfully carried out phase shift keyed (PSK) experiments using the fiber optics transmission set-up shown in Figure 1. The system was carefully characterized including the receiver. We were able to achieve receiver sensitivities that are comparable to the best published results. The experiment of interest is to find out what happens when a single PSK channel is transmitted among amplitude modulated channels. This experiment could emulate what might happen in a multi-optical format network. To realistically emulate an installed network we introduced a polarization scrambler into the fiber optic recirculating loop.

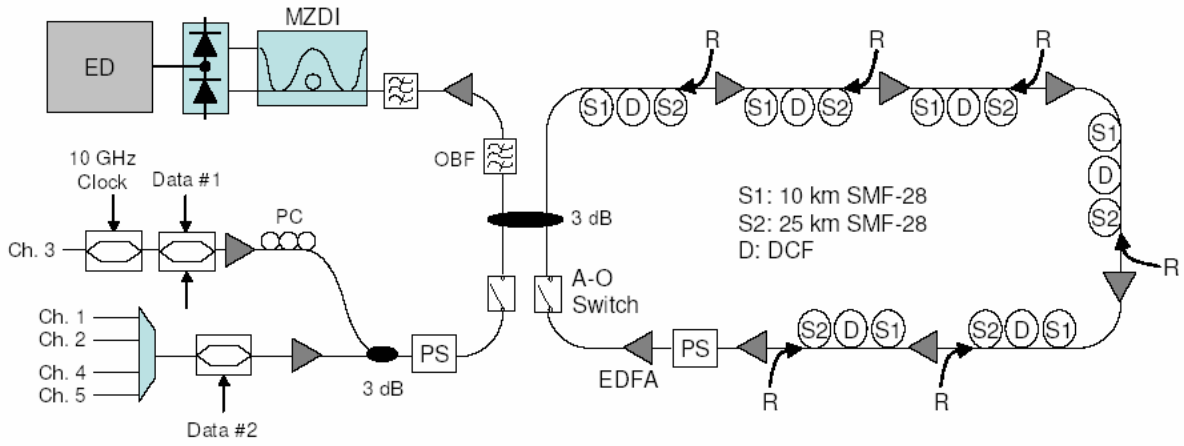

Fig. 1: Experimental schematic. R: Raman pump (1455 nm), PS: polarization scrambler, PC: polarization controller, OBF: optical bandpass filter, MZDI: Mach-Zehnder delay interferometer, ED: error detector.

The first set of experiments were to shut off the data modulation in the side channels to eliminate pattern dependent nonlinear interactions with the PSK channel. We have not verified this but we predict that this would emulate the situation where all of the channels are PSK modulated. We obtained low error rates even out to distances of 2500 km. By turning on the amplitude modulation of the side channels we observed severe degradation in the PSK channel limiting transmission to a few hundred kilometers. We are continuing investigate the performance of this configuration as a function of channel spacing and relative polarization of the channels. This work has enormous implications for multi-format optical networks.

5.2 Statistical Signal Characterization (J. Morris)

1. In previous work, we analyzed the binary asymmetric channel (BAC) and the binary symmetric channel with erasures (BSC/E) in order to compute their respective capacities and capacity-achieving prior probability mass functions

(pmfs). Along similar lines, we pursued the analysis of two other channel models with binary inputs and ternary outputs. These channel models are the binary asymmetric channel with symmetric erasures (BAC/SE) and the binary symmetric channel with asymmetric erasures (BSC/AE). The BAC/SE and the BSC/AE are both special cases of the more general binary asymmetric channel with asymmetric erasures (BAC/AE), which is shown in figure 1.

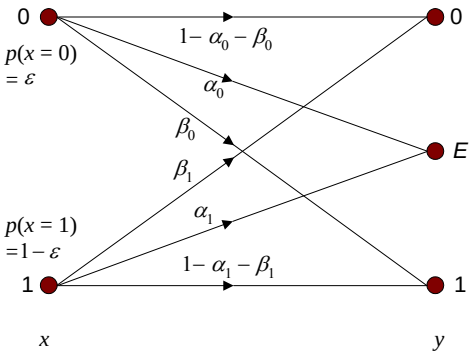

Figure 1: The binary asymmetric channel with asymmetric erasures (BAC/AE) channel model

The BAC/SE is derived from the BAC/AE if $\alpha_0 = \alpha_1 = \alpha$. Thus, the BAC/SE is defined by three parameters: β_0, β_1, and α. On the other hand, the BSC/AE results from the BAC/AE if $\beta_0 = \beta_1 = \beta$. Thus, the BSC/AE is defined by three parameters: α_0, α_1, and β. Here, we note that both the BAC/SE and the BSC/AE model can be arrived at from the optical fiber communications (OFC) channel with the effective noise in the electrical domain modeled via chi-squared probability density functions (pdfs) by appropriate positioning of the two decision thresholds that are required at the receiver to generate three decision levels (see progress report Feb-May 2004).

For the BAC/SE a closed-form solution for the capacity-achieving prior pmf (ε) exists and is given by

$$\varepsilon = \frac{\left[\theta\left(1-\alpha-\beta_1\right)-\beta_1\right]}{\left(1-\alpha-\beta_0-\beta_1\right)\left(1+\theta\right)}, \quad (1)$$

where

$$\theta = \exp\left(\frac{H\left(\beta_1,\alpha,1-\alpha-\beta_1\right)-H\left(\beta_0,\alpha,1-\beta_0-\alpha\right)}{1-\alpha-\beta_0-\beta_1}\right) \quad (2)$$

and $H\left(x,y,z\right) = -x\ln x - y\ln y - z\ln z$. This capacity-achieving prior can be used to compute the pmf of the output, and thus the capacity of the BAC/SE. Also, of practical importance, how close to this capacity is the mutual information when using the uniform prior pmf?

For the BSC/AE, a closed-form solution for the capacity-achieving prior does not exist. In this case, the capacity-achieving prior must be obtained by computing the roots of the following expression:

$$\frac{\left[\left(\alpha_0-\alpha_1\right)\varepsilon+\alpha_1\right]^{\alpha_1-\alpha_0}\left[\beta+\left(1-\varepsilon\right)\left(1-2\beta-\alpha_1\right)\right]^{1-2\beta-\alpha_1}}{\left[\beta+\varepsilon\left(1-2\beta-\alpha_0\right)\right]^{1-2\beta-\alpha_0}}=e^{H\left(\beta,\alpha_0,1-\beta-\alpha_0\right)-H\left(\beta,\alpha_1,1-\beta-\alpha_1\right)}$$

. (3)

Figure 2(a) shows a plot of the capacity-achieving ε for a family of BSC/AE channels defined by β = 0.01, and α_0 and α_1, each chosen from the set $\{0.01, 0.02, \cdots, 0.1\}$, while figure 2(b) plots the corresponding capacities.

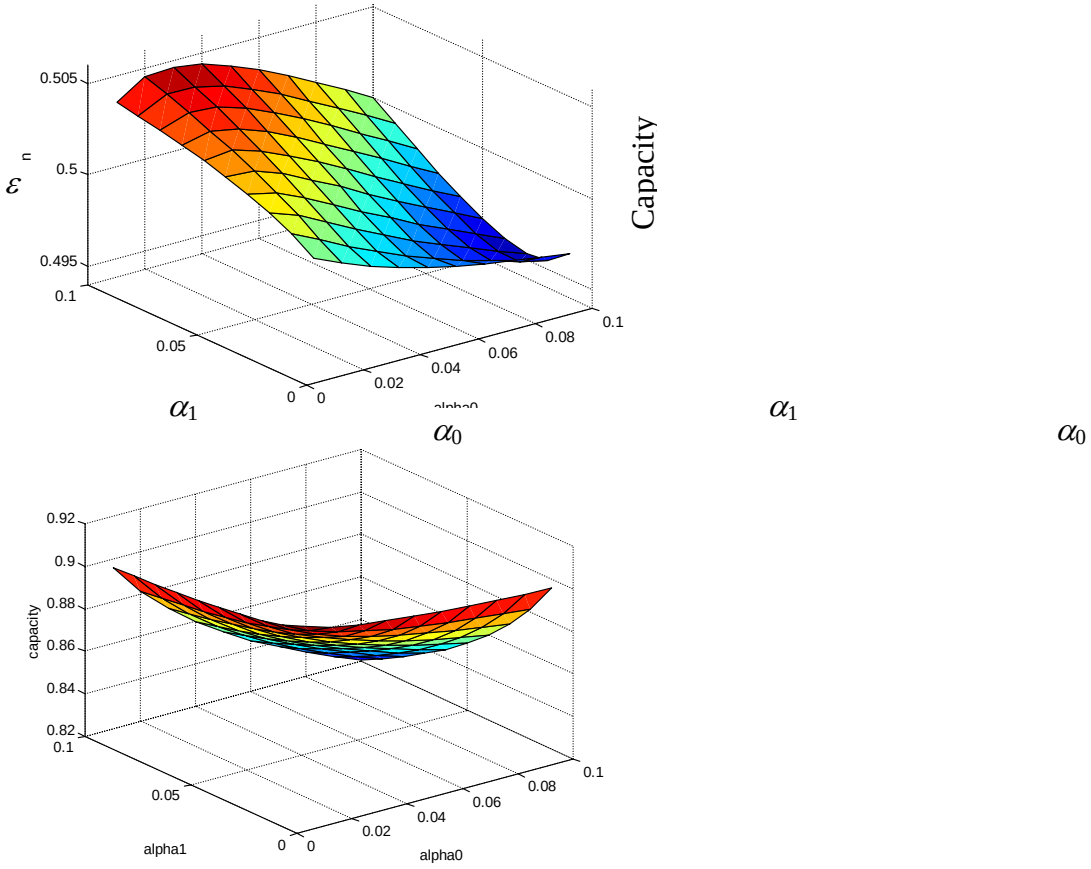

Figure 2: (a) Plot of the capacity-achieving ε and (b) corresponding capacities, for a family of BSC/AE channels defined by β = 0.01, and α_0 and α_1, each chosen from the set $\{0.01, 0.02, \cdots, 0.1\}$.

2. We have succeeded in proving some theorems that characterize the decoding of 3-error patterns for RCD codes defined by parameter η, where η is the size of the

codeword square array, via the bit-flipping algorithm (BFA) [1]. A classification scheme for 3-error patterns for RCD codes based on the maximum number of unsatisfied parity-check equations for any bit-position in the RCD array was presented previously. The new results specifically relate to classes 1, 2.1, and 2.2 and are as follows:

Theorem 1: Any arbitrary weight-6 RCD codeword can be partitioned into a unique pair of class-1 3-error patterns.

Theorem 2: Exactly $3\eta^2(\eta-1)$ class-2.2 3-error patterns each have a single unsatisfied triple-intersection point (*utip*), and each of these decodes to a codeword of Hamming weight 2η, thus generating a decoder error, in exactly 2 BFA iterations.

Theorem 3: For an RCD code with parameter η, at least $\frac{3}{2}\eta^2(\eta-1)(\eta-3)(\eta-4)$ class-2.1 3-error patterns decode successfully to the all-zeros codeword.

Here, a *utip* is defined as a location on the RCD code array that has all three of its associated parity-check equations unsatisfied. In addition to providing us with valuable insight into the structure of RCD codes and the BFA, these results will also enable us to compute bounds on the probability of decoding error for the family of RCD codes without resorting to computer simulations.

3. Maximum-likelihood (ML) decisioning at the receiver of the optical fiber communications (OFC) channel with the effective noise modeled via chi-squared pdfs results in a binary asymmetric channel (BAC) characterization. The ML decision threshold, t_{ML}, can be numerically computed from the space and mark pdfs, which have been presented previously. The ML-threshold can then be used to compute the transition probabilities ε_0 and ε_1 that describe the resulting BAC. It was previously observed that for all values of the chi-squared noise parameter M, the difference between ε_0 and ε_1 becomes negligible as the signal-to-noise ratio (characterized by the variable β) increases, indicating that the BAC was only mildly asymmetric and could be approximated by a BSC at high β.

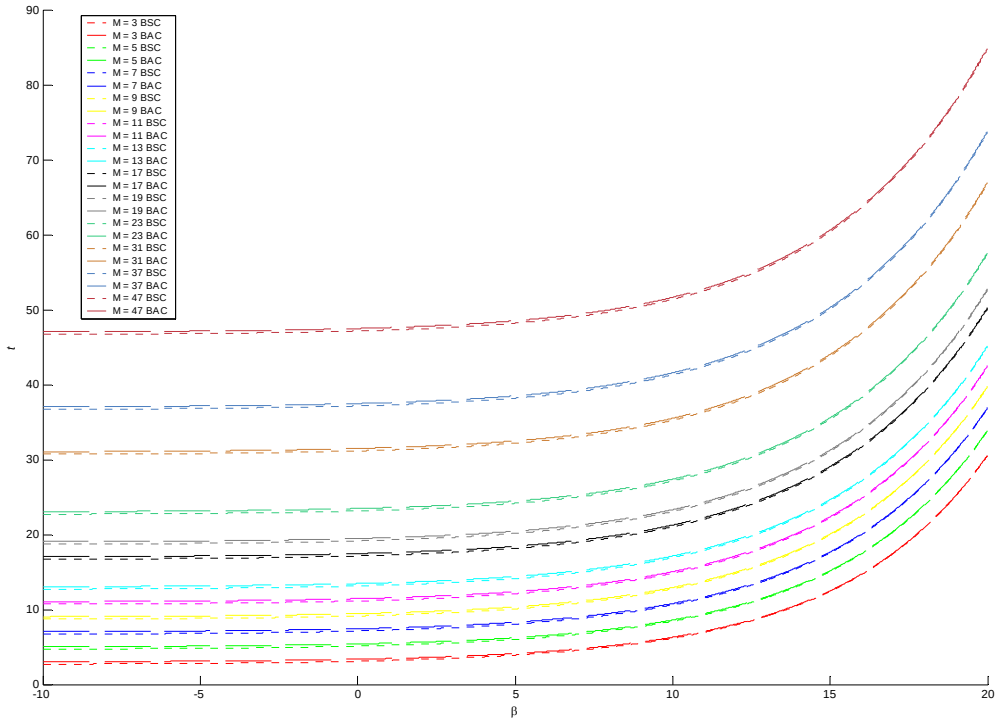

Figure 3: Plots of t_{ML} and t_{eq} as a function of β for various values of M.

Based on this observation, and using the chi-squared pdf to model the effective noise, we numerically computed the threshold t_{eq} that equalized the transition probabilities, and thus resulted in a BSC characterization as opposed to a BAC characterization.

Figure 3 shows plots of t_{ML} and t_{eq} as a function of β for various values of M. The suffix BAC in the legend represents the curves for t_{ML} while the suffix BSC represents the curves for t_{eq}. We observe that in all cases, the t_{eq} curves fall below the t_{ML} curves, and the difference between values of t_{eq} and t_{ML} for all M is small, especially for the larger values of β.

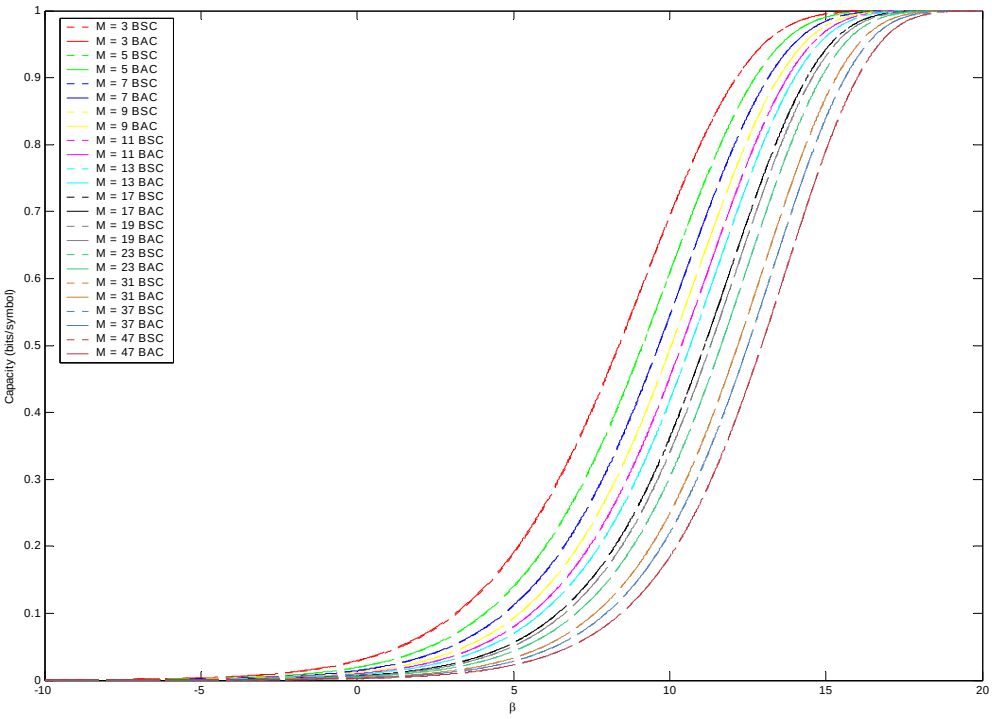

Figure 4: Plots of capacities of the BAC models based on t_{ML} and the BSC models based on t_{eq} as a function of β for various values of M.

In figure 4, the capacities of the BACs resulting from using thresholds t_{ML} are compared with the capacities of the BSCs resulting from using thresholds t_{eq}. It is clear from these plots that there is hardly any visible difference in the capacities of the two models for the same M and β. Thus, we suffer minimal loss in performance when we use the threshold t_{eq} and force the effective channel to be a BSC, instead of using the optimal threshold t_{ML} and, thereby, being constrained to an asymmetric channel model.

4. Figure 5 shows plots of the Shannon limit in dB versus code rate R_c for the binary symmetric channel with erasures (BSC/E) based on binary phase-shift-keying and additive white Gaussian noise with symmetric thresholds at $\pm t$ (this figure was presented in previous reports and is included here for reference). From figure 5, we observe that for every value of code rate, there exists a specific value of t that achieves the lowest possible Shannon limit among the values of t that are considered. Figure 6 shows a plot of the Shannon limit minimizing BSC/E threshold $t_{SL\,min}$ as a function of R_c. From figure 6, it is clear that $t_{SL\,min}$ decreases with increasing R_c and approaches 0 at $R_c \to 1$. This agrees with our intuitive observation that at $R_c = 1$ (absence of coding), a non-zero BSC/E decision threshold only decreases the probability of correct transmission and, thus, results in a worse channel as opposed to when the decision threshold equals 0.

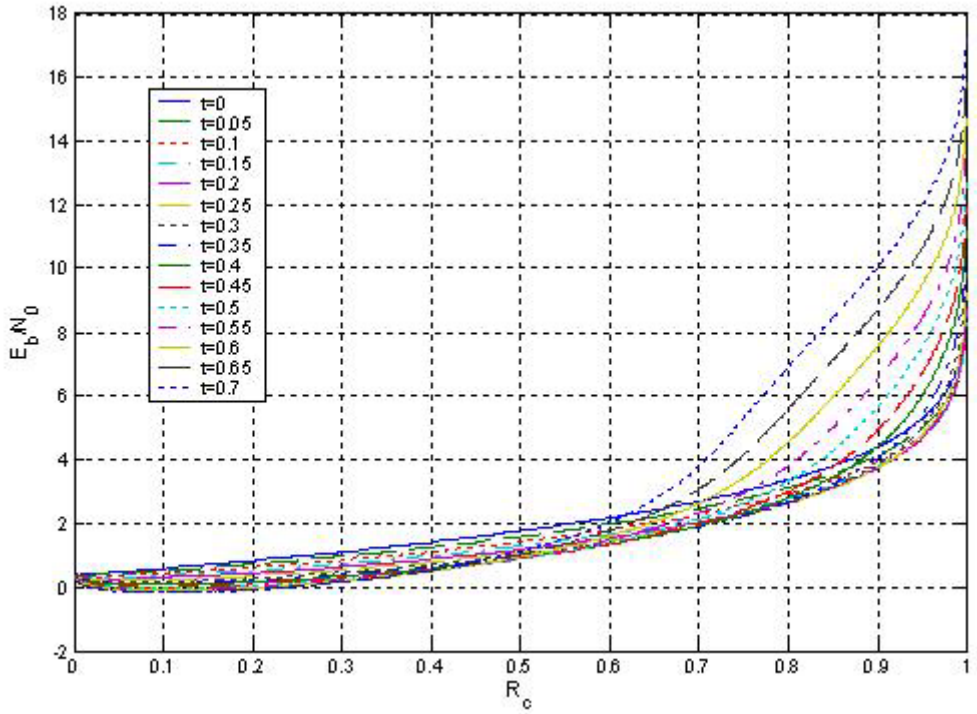

Figure 5: Plots of Shannon limit in dB versus code rate for the BSC/E for different symmetric thresholds at $\pm t$

Finally, figure 7 plots the minimum Shannon limit (achieved by setting the decision threshold to $t_{\mathrm{SL\,min}}$) as a function of R_c. These results clearly indicate that smaller Shannon limit values are achievable using the two decision thresholds $\pm t$ when compared with the single-decision-threshold-based BAC/BSC models.

5. The fully parallel implementation of the BFA decoder for the (255,175) Type-I EG LDPC code [Kou] was further augmented by allowing for the BFA iterations to terminate as soon as the parity-check matrix was satisfied (previously the iterations continued until the maximum limit was reached). An input and output shift-register have also been added to the design to enable serial-to-parallel and parallel-to-serial conversion, respectively, in order to enable interfacing of the decoder with a realistic communications system. Finally, the option of being able to directly output the received word in the event of a decoder failure has also been included in the current implementation.

References

A. Mahadevan, "On RCD Codes as a Class of LDPC Codes: Properties, Decoding, and Performance Evaluation", PhD proposal, CSEE Department, UMBC, Baltimore, MD 21250, Jan. 2004.

Y. Kou, S. Lin, and M. P. C. Fossorier, "Low-Density Parity-Check Codes Based on Finite Geometries: A Rediscovery and New Results", *IEEE Transaction on Information Theory*, Vol. 47, No. 7, pp. 2711-2736, November 2001.

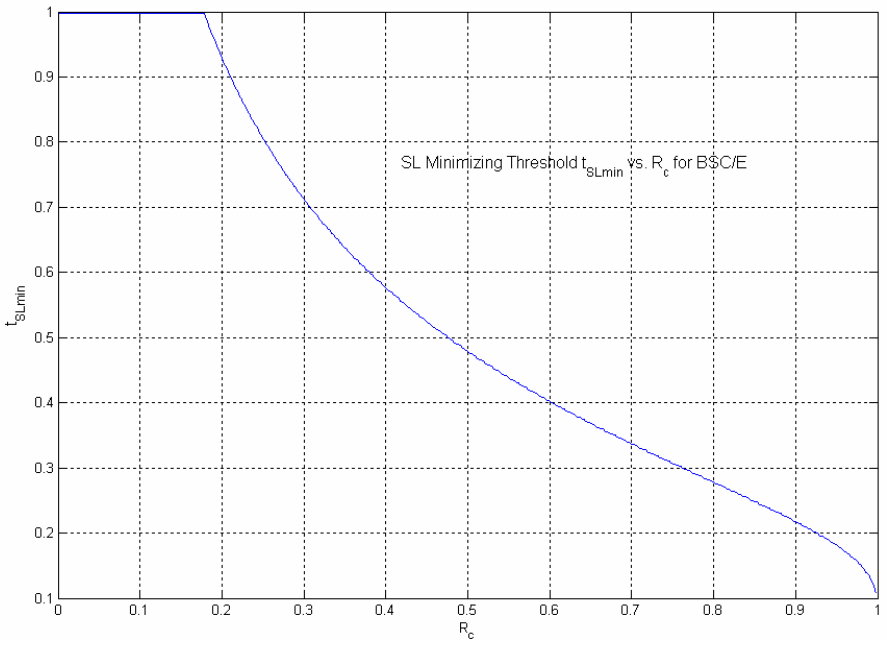

Figure 6: Plot of the Shannon limit minimizing BSC/E threshold $t_{\mathrm{SL\,min}}$ as a function of code rate R_c.

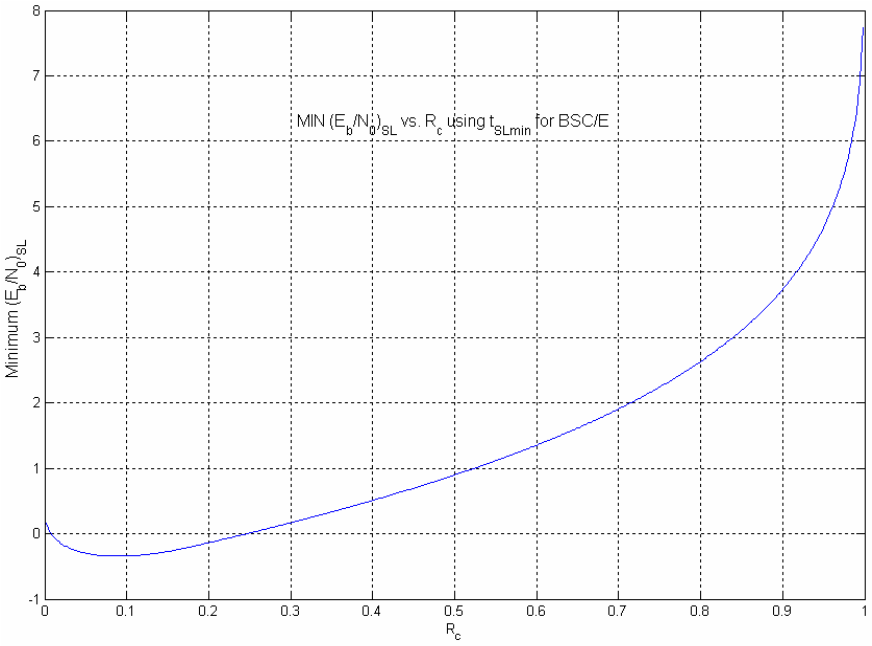

Figure 7: Plot of the minimum Shannon limit as a function of code rate R_c.

6. Peer to Peer Computing (Michael Marsh)

6.1 Lookup in Unstructured Networks (In collaboration with Bobby Bhattacharjee, Aravind Srinivasan, and Jonathan Katz.)

The ability to locate an object or resource in a network is crucial to many distributed applications. While efficient lookup algorithms exist, these require a network that possesses a particular structure. In contrast, many naturally occurring networks exhibit little structure, and constructing distributed applications on these networks requires a lookup algorithm that is relatively efficient without any structural requirements.

Previously, we had developed the Local Minima Search (LMS) algorithm, which provides probabilistic guarantees for locating items and scales better than other existing algorithms. To better understand various design parameters, we have begun the construction of a peer-to-peer public key infrastructure that extends the notion of a web of trust as employed by PGP. The nature of such a PKI is such that an efficient lookup algorithm, such as a distributed hash table (DHT), requires a recursive search for a signature chain linking the key owner and the principal requesting the key. By employing an unstructured algorithm, we are able to perform public-key lookup directly over a network constructed from the links implied by key signatures. Such a design scales better than the corresponding DHT-based design.

In conjunction with the public key infrastructure development, we have also considered a number of security enhancements to the initial LMS protocol. Message

tampering and injection are countered by adding digital signatures at various places in the messages, so that a peer has high assurance that a message has been handled correctly from its point of issuance to receipt. We are also investigating fault tolerance in the presence of malicious peers based on a probabilistic extension of the Byzantine quorum systems introduced by Dahlia Malkhi and Michael Reiter.

LMS is described in UMD Computer Science Department Technical Report CS-TR-4593. This report includes a rigorous analysis of the protocol as well as extensive simulation experiments as well as experiments on a working implementation.

6.2 Emergent Network Behaviors

Computer networks are very complex, and this complexity continues to increase as networks grow both in number of hosts and variety of applications competing for resources on a single host. Such complex systems often exhibit chaotic behavior, but this has as yet only been studied in a rather cursory manner. We hope to develop a more principled way of assessing the complexity of a network and characterizing its large-scale behavior.

Specifically, we are interested in investigating the phase changes that a network might experience, as well as other crises. An obvious taxonomy for network behavior comes from the field of fluid dynamics. In an under-utilized network, packets might interact like particles in a gas. As utilization increases, packet interactions might begin to more closely resemble a liquid. The transition point between these two regimes is likely to be marked by drastic changes in the overall behavior of applications running over the network. In addition, liquids tend to exhibit both laminar (smooth) flow and turbulent (chaotic) flow. We hope to identify what network configurations are likely to result in one type of behavior, and what configurations are likely to result in the other. Additionally, we expect many networks to exhibit additional crises in which laminar flow is exchanged for turbulent flow, and vice-versa. A number of publications in the past few years seem to motivate the treatment of a packet network as a fluid. These studies, however, have concentrated on a fluid-flow treatment of network traffic as a way to speed up simulations, rather than as a technique for understanding traffic dynamics.

Our work to date has focused primarily on developing a simulation infrastructure (based on a popular open source packet-level network simulator) for examining large networks. We have, concurrently, begun developing mathematical models for characterizing the behavior of these networks. We expect the bulk of our effort on this project to shift to modelling in the very near future.