Progress Report on the LTS/UMIACS Contract
RFP:MDA904-02-R-0151 – SOW: R4-02-0001.1
October 1,2004 – December 31,2004

1. Active Network Management (Mark Shayman, Samrat Bhattacharjee, Steve Marcus, Ray Chen, 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 previously developed a distributed DDoS detection system to detect TCP attacks originating from within the AS. The detection system consists of monitors which are deployed at different routers in the AS and which sample the network traffic. The traffic profiles collected at different monitors are correlated to detect DoS attacks and suppress false positives. We modified the algorithms at the monitors to collect enhanced statistics. We use these statistics to score the malicious quality of the suspected flows and use the score to detect attacks. The scoring of flows enables unified accounting of different characteristics of an attack. This simplifies the configuration space for the network administrators without compromising the flexibility of our detection system. The new algorithm also performed better by decreasing the detection times for the attacks.

We applied the new scoring algorithms (with a few changes) in autonomous systems with asymmetric traffic—i.e., where the gateway taken by the outgoing traffic on a connection is different from the gateway taken by the incoming traffic. With slightly increased processing overhead, we significantly reduced the time taken to detect the attacks and the number of legitimate flows that might be labeled as attacks by the detection system.

We used a packet level simulator to test our detection scheme and performed extensive simulations using two real Internet traffic traces and synthetic attack traffic. The two traces are collected at different locations and at different times, and characteristics of the two traces are significantly different. (For example, one trace has a packet rate of 3000 packets per second while the other has 100,000 packets per second.) We used the smaller of the two traces to configure the detection system and used the larger trace to test the detection system with the selected parameters. This approach helped us demonstrate that our detection system is practical and scalable.

We performed an extensive study of the system configuration parameters and, under the constraints of state and processing overheads, evaluated the optimal values for detection sensitivity. We have evaluated the detection system in an AS with a single border router under various attack scenarios. Our results show that with state requirements about 20% of per flow state and 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.

We also evaluated the detection system in an AS with two gateway routers and with different amounts of asymmetric traffic. Again, the detection system detects all attacks and signals less than one false positive on average for most of the asymmetric traffic scenarios. Only under extreme asymmetric traffic scenarios does the detection system report a few false positives.

1.2 Overlay-based Traffic Engineering

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

We are continuing our work on the traffic mapping (load balancing) problem using in-network overlays. We have generalized our algorithm to simultaneously load balance multiple traffic sources 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. We address the optimal multipath multicast routing problem in a more general framework than the traditional approach where there are multiple distribution trees. We have considered different network models with different functionalities. With this generalized framework, our goal was to examine the benefits observed by the addition of new capabilities to the network beyond basic operations such as storing and forwarding. Specifically, we have relaxed the usual assumption that from a given multicast tree each receiver gets multicast packets at the same rate. Note that each 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. We have efficiently overcome this potential bookkeeping problem by using a specific source coding called Digital Fountain codes, 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.

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 not possible 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 small neighborhood in the vicinity 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. We are investigating possibilities to optimize the overlay topology. Our goal will be to establish a 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.

Another issue we consider doing is to generalize our algorithm to wireless networks. In contrast to wireline networks that are characterized by high transmission rates and very low error probabilities, wireless networks are highly erratic radio propagation environments. The error probabilities are higher and the quality of the radio links is constantly changing due to movement of the nodes in addition to multipath propagation, interference and other transmission impairments. Consequently, the radio links must be carefully monitored and controlled to ensure reliability. As our algorithm depends on actual network measurements and reacts adaptively according to the changes observed, it can naturally be extended to the wireless network domains. We are planning to generalize our routing algorithm for wireless domains using a specific network cost structure that will not only consider utilization but also conditions related to the wireless nature of the network.

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 continued 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 networks 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.

In the proposed architecture, we consider two classes of traffic, real time (class 1) and best effort (class 2) traffic. Every source-destination class 1 and 2 traffic demands are known. The DTE structure distributes class 1 and 2 traffic between the network paths such that utilization of a link that carries 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. In this architecture, we only overprovision those links that carry class 1 traffic and there is no need for packet classification in the network, since those links that carry class 1 traffic are already overprovisioned and a simple FIFO queue would be able to accommodate QoS requirements.

We have identified two major optimization problems in this architecture. In the first problem, the Path to Class Assignment (PCA), we assume that we have a fixed set of paths and the problem is to distribute class 1 and 2 traffic between them. In the previous report, we briefly described the optimization algorithm based on Simulated Annealing that we have developed for PCA. The details of the DTE architecture, and PCA algorithm and the simulation results will be presented at Infocom 2005.

The second problem is Path Selection and Routing (PSR). In PSR, we select a set of paths to accommodate a specified class 1 and 2 traffic demand. We have developed an iterative gradient projection based algorithm for PSR. At each iteration, the algorithm uses link cost derivatives as the link lengths and applies the K-shortest path algorithm to select a good candidate paths set. At the end, the PSR algorithm provides a set of paths and an initial load distribution among them.

Theoretically, the PCA algorithm should be able to enhance the PSR performance. However, we do not observe any significant improvement in our simulations so far. We have to study the PSR performance and its interaction with PCA more carefully before making any conclusion. However, this is a very encouraging result if it holds, since PCA is much more complex and time consuming than PSR.

Our plan is to complete a thorough study on the PSR performance. Furthermore, we will study if it is feasible to have a distributed and asynchronous PSR implementation.

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.

In prior work, we have used a stochastic process model to generate the sequence of traffic matrices and to predict future traffic demands. That model consisted of a

deterministic part representing average daily variations in traffic as well as a random part (Brownian Bridge process) representing fluctuations about the deterministic part. During the current period, we have replaced this mathematical model with a model based on actual network data.

We have collected real traffic data from the Abilene Network. This data collection was a very long process since we were interested in source-destination traffic but the data available on the Abilene network was link traffic (number of packets going through each interface of each node). In order to generate source-destination traffic from link traffic we have to inspect each packet incoming into a node from external interfaces, find the destination address of the packet, and resolve the destination address using BGP tables available for the network to find the egress node of that packet on the Abilene network.

We intend to average the data for a few weeks and use that as a prediction of how the traffic changes over the course of a day. We expect to show that by having a daily trend for the traffic, our algorithm is able to better adjust to the changes in traffic and step-by-step transition to better topologies.

Previously, we had used an algorithm by Modiano with which to compare our algorithm. In recent work, we have applied rollout to this algorithm and showed that rollout improves this algorithm. We have also compared the rollout of Modiano's algorithm to our rollout algorithm. Each algorithm performed well based on the corresponding cost criteria when the network was not congested. But our algorithm performed better in congestion because Modiano's algorithm has a cost that depends only on the maximum link utilization. Therefore, it does not differentiate between having one link congested and having several links congested. Our algorithm fixes this problem by having dropped calls in the cost function.

We have run simulations with various rollout horizons. We started with horizon = 1 (heuristic) and went up to 5. We see great improvements when we increase the horizon to 2 and 3. However, there is little or no improvement when we increase the horizon to 4 and 5, so it does not appear to be worth the additional computation to extend the horizon to these values.

In order to determine how close to optimal our online algorithm is, we are trying to come up with offline algorithms that give near optimal topologies for given traffic patterns. We do this by generating random initial topologies and then running our rollout algorithm to find a local optimum. We use several of these random initial topologies, and then select the one that results in the best local optimum. One of the parameters that needs to be determined is the number of random initial topologies in order to obtain a nearly optimal final topology.

We are planning on comparing our algorithm to the performance of the offline algorithm explained above and to heuristic algorithms that exist in the literature.

Publications

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.

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

P. Fard, R. J. La, K. Lee, S. Marcus and M. A. Shayman, Reconfiguration of MPLS/WDM networks using simulation-based Markov decision processes, *Conference on Information Sciences and Systems*, The Johns Hopkins University, March 2005, submitted.

T. Guven, C. Kommareddy, R. J. La, M. A. Shayman, B. Bhattacharjee, Measurement based optimal routing on overlay architectures for unicast sessions, *Computer Networks Journal,* special issue on Network Modeling and Simulation, submitted.

T. Guven, R. J. La, M. A. Shayman, B. Bhattacharjee, Measurement based multipath multicast, *IEEE Global Internet Symposium*, Miami, Florida, March 2005, submitted.

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

Last quarter we proposed two areas of new work. The first was a systems analysis of retro-fitting mandatory access control (MAC) onto commodity operating systems. The results of that analysis, attached, were submitted to the IEEE Security and Privacy Conference. In summary, we found that Linux, Windows, and OSX all suffer from deficiencies in their application architecture that significantly limit the effectiveness of MAC. **See the attached paper for more details**.

The second area where we have focused is on developing an automatic method to expand the invariant space that Copilot can cover. Currently, Copilot only covers the text segments of a process. But, large sections of the data space are invariant as well, e.g. jump tables. By performing a static analysis of a program, we have shown that certain classes of static jump tables can be identified—effectively expanding the coverage space of Copilot, paper attached.

Our future work will focus on the continued automatic expansion of the coverage space of Copilot. We also hope to define what we are currently calling “semantic integrity”. That is we want to monitor long running processes ensuring that the data within that process remains within its semantic bounds.

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

3.1 Ubiquitous Wireless Interworking Test-bed (UWIN) (W. Arbaugh)

This quarter has focused on redesign and documentation of our past work as well as continued development of a novel authentication mechanism and expanding the capabilities of the iButton to include TPM functionality.

In the previous quarter, we identified power problems with the use of the Soekris boards. Rather than solve that problem, we've decided to change our strategy with the test-bed. Our original plans to deploy along I-95 require the solar/batter power, and the nature of the interstate limits the overall width of a mesh. As a result, we're refocusing on deploying a wide-area mesh on campus. This significantly reduces the bureaucratic aspects of the project, eliminates the power problems, and provides a wider mesh area. Our plans are to begin deployment of the test-bed during the summer of 2005.

We've continued development of the EAP-PAX authentication protocol. EAP-PAX is designed to avoid intellectual property rights associated with previous strong password mechanisms. In addition, we wanted to include key management and PIN bootstrapping. EAP-PAX is provably secure under the random oracle model and has been selected to move towards the standards track as an EAP Working group document. The current draft for EAP-PAX is attached.

With respect to task #3, we've begun implementation of an iButton TPM implementation. We believe that we can implement all (and more) of the TPM functionality on a higher assurance device than a smart card. We expect to have this implementation done by the end of this coming quarter.

We also completed a paper on modeling co-channel interference in Wi-Fi networks and using that model to assist in initial channel assignment for base stations. Our approach has shown experimentally and via simulations to reduce client co-channel interference by a significant amount—paper attached.

The remainder of the work this quarter has focused on the preparation of journal publications. We are currently preparing a journal paper summarizing our neighbor graph work, and we completed a paper for IEEE Proceedings (accepted for publication next year) that documents our work on roaming—**paper attached**.

3.2 Location Determination Using RSSI (A. Agrawala and U. Shankar)

The main objective of this effort has been to determine the location of an active unit having a NIC so that the RSSI from the APs can be measured by the unit. The measured RSSI is used to determine the location of the unit. Our earlier work demonstrated the technology in operation with measured accuracies of a few feet in an indoor environment with multiple APs.

Our recent work in this area has been to incorporate the heuristics in the location determination by reflecting additional knowledge including the past location estimates, direction of motion, location of barriers, etc. As a result, the accuracy and the robustness of this technology have improved significantly.

We have also integrated an audio channel between two nodes. Any node can establish an audio link by clicking on the icon of the other node as displayed on the map. When the called node accepts the connection the audio communication can start.

3.3 Nuzzer Technology (A. Agrawala and U. Shankar)

In this work we are measuring the effect of the presence of a passive entity (e.g. a person) in an RF field to determine the presence and the location of such passive entity. Recently we demonstrated the same effects outdoors for the presence of vehicles.

In order to take the next steps to make this technology robust we need to understand the electromagnetic properties of the fields. We have established a working relationship with another group on campus, in the Institute for Research in Electronics and Applied Physics and are collaborating with them to establish the experiments for the next phase of this effort. CISCO has also indicated their interest in participating in the effort.

In order to better understand the characteristics of the RF fields in indoor environments we have developed a simulation model and tested it through empirical measurements. We find that the observed variability is accurately predicted by a model which only takes the direct and once reflected waves into account. We are using this model to understand the changes in RF fields generated by the presence of people and other objects.

3.4 SALAM: A Scalable Anchor-free Localization AlgorithM For Ad Hoc Sensor Networks (A. Agrawala and U. Shankar)

There has been a growing interest in the applications of wireless sensor networks in unattended environments. In such applications, sensor nodes are usually deployed randomly in an area of interest. In order to correlate the reported data to the origin of the sensed phenomena an accurate knowledge of node location is essential. Further, awareness of the nodes' positions can enable employing efficient management

strategies, such as geographic routing, and conducting important analyses such as node coverage properties.

Most of the localization algorithms reported in the literature have focused on using a number of specialized nodes that know their positions. Such specialized nodes are usually referred to as *anchors nodes*. The rest of the nodes try to estimate their location by exchanging information to determine their distances to the anchors. Most of the anchor-based algorithms require a high percentage of anchor nodes in order to reach an acceptable accuracy. Most of these algorithms suffer from scalability problem either because they assume that anchors' positions are flooded into the network, or they require centralized computations.

SALAM is an anchor-free locally-centralized localization protocol that determines the position of sensor nodes consistently with low errors. We assume that there are no anchor nodes with known positions. The network is divided into clusters each with its own gateway node. Each gateway is responsible for building a local map of relative positions of the nodes. We formulate an optimization model to minimize the cumulative errors that may affect the accuracy of the established relative coordinate system. The gateways collaboratively combine their local maps to obtain the global relative topology of the network. This approach is highly scalable and is not affected by network partitioning while achieving accuracy similar to those obtained by a centralized approach.

3.5 Wireless Capacity Enhancement (A. Agrawala and U. Shankar)

Various properties of wireless networks, such as mobility, frequent disconnections, and varying channel conditions, make designing efficient protocols for these technologies a challenging task. We believe that enhancing the performance of wireless networks requires alleviating the effect of the physical layer characteristics (e.g., channel noise) and developing cross layer mechanisms to exploit those characteristics for enhancing network performance.

The wireless IEEE 802.11 MAC protocol, are based on Carrier Sense Multiple Access (CSMA) mechanism. In CSMA, a station may transmit if and only if the medium is sensed to be idle. The purpose is to prevent any station from causing interference to an ongoing transmission occupying the medium. We proposed an enhancement to the existing IEEE 802.11 Distributed Coordination Function (DCF) MAC to improve channel spatial reuse efficiency, and thus improve overall network data throughput. The modification, named the Location Enhanced DCF (LED) for IEEE 802.11, incorporates location information in DCF frame exchange sequences so that stations sharing the communication channel are able to make better interference predictions and blocking assessments. Utilizing an underlying physical layer design that supports frame capture, the LED enhanced interference estimation increases overall network data throughput by permitting more concurrent transmissions. Frame capture uses the well known "physical layer capture" phenomena in radio channels that allows the receiver to capture a frame if the frame's detected power sufficiently

exceeds the joint interfering power of interfering contenders by a minimum certain threshold factor. In [1], we studied analytically the potential performance enhancement of the LED over the original IEEE 802.11 DCF. The results are verified using the ns-2 simulator, which shows that up to 35% of DCF blocking decisions are unnecessary and our LED method can achieve up to 22% more throughput than the original DCF. The effects of node parallelism degree on the proposed protocol as well as its performance in the mesh networks are given in [2].

[1] T. Nadeem, L. Ji, A. Agrawala, and J. Agre. Location Enhancement to IEEE 802.11 DCF. To appear in *IEEE INFOCOM 2005*, Miami, Florida, USA, March 13-17, 2005.

[2] T. Nadeem, L. Ji, A. Agrawala, and J. Agre. IEEE 802.11 DCF Location Aware. Submitted for publication.

3.6 Wireless Networks in Noisy Environments (A. Agrawala and U. Shankar)

Wireless communication suffers from transmission errors due to channel noise. In order to increase transmission reliability, IEEE 802.11 standard implements a retransmission mechanism in which a packet is retransmitted over a link if no MAC layer acknowledgment is received. Since minimizing energy consumption during communication is an important goal in wireless networks, we developed techniques to compute energy-efficient reliable paths within the framework of on-demand routing protocols. In [3], we showed how the proposed schemes account for channel error rates in computing such paths. Simulation results showed that the proposed variants of on-demand routing protocols can achieve orders of magnitude improvement in terms of energy-efficiency. In addition, IEEE 802.11 adopts a fragmentation mechanism in which large packets are partitioned into smaller fragment to increase their transmission reliability. This fragmentation mechanism should be considered by the routing protocols in evaluating the reliable and energy efficient routes. In [4], we developed mechanisms exploiting the IEEE 802.11 fragmentation mechanism to account for channel characteristics and generate optimum energy-efficient paths. Our results show that our proposed variants of on-demand routing protocols can achieve orders of magnitude improvement in energy-efficiency of reliable data paths.

[3] T. Nadeem, S. Banerjee, A. Misra, and A. Agrawala. Energy-Efficient Reliable Paths for On-Demand Routing Protocols. *Sixth IFIP/IEEE International Conference on Mobile and Wireless Communication Networks (MWCN'04)*, Paris, France, October 25-27, 2004.

[4] T. Nadeem, and A. Agrawala. IEEE 802.11 Fragmentation-Aware Energy-Efficient Ad-Hoc Routing Protocols. *1st IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS'04)*, Fort Lauderdale, Florida, USA, October 24-27, 2004.

3.7 WLAN QoS (A. Agrawala and U. Shankar)

In previous studies, we showed experimentally that physical layer capture is the major reason for unfairness. Specifically, when a frame with a higher signal strength collides with a frame with a lower signal strength, it is often extracted by the receiver and the frame with weaker signal strength is lost. Therefore, there is unfairness in the physical layer, which is propagated up the protocol stack as imbalance in performance obtained. We found that the unfairness can be around 10% at the MAC layer, and upto 50% at the TCP layer. An important offshoot of our experimental work is that we learned how to synchronize the wireless transmission and reception logs of sniffers and nodes of a WLAN to within 5 microseconds. We examined several network simulators that handle 802.11 links (including *ns2* and Qualnet (the top-of-the-line commercial simulator), and found that they do not account for the physical capture effect. Fixes have been implemented for these simulators. This work appears in

Andre Kochut, Arunchandar Vasan, Udaya Shankar, Ashok Agrawala. **The Impact of Physical Layer Capture in 802.11b WLANs.** Published in IEEE ICNP 2004. Old version: http://www.cs.umd.edu/~shankar/Papers/mh2003.pdf

3.8 Timestepped Simulation for WLAN/WAN (U. Shankar and A. Agrawala)

***Introduction*:** Performance evaluation is crucial to advancing the state of computer networking. Yet existing evaluation techniques are not adequate for handling the large heterogeneous architectures and state-dependent control schemes that are intrinsic to modern computer networking. A state-dependent control scheme is one where control (e.g., source rate or window size, next hop, buffering threshold) is exercised based on the observed state of the network (typically averaged over some recent interval). Packet-level simulation, currently the most popular technique, is accurate but does not scale with link bandwidth and network size. Purely analytical techniques do not capture the effects of state-dependent control or realistic traffic mix with reasonable accuracy. In previous work, we developed a technique, called **Z-Iteration**, that yields the time-dependent evolutions of various instantaneous ensemble metrics of interest (e.g., loss rate and queue size of a link, throughput and delay of a connection) for TCP/IP networks of considerable size and speed. However, the Z-iteration handled time-dependent control mechanisms but not state-dependent control mechanisms. TSS overcomes this serious limitation.

Recent Work: We proposed a novel method, called *Timestep Stochastic Simulation (TSS)*, that combines discrete-event simulation and analytical approximations to yield *sample paths* of high accuracy and low computational cost (independent of packet rates and queue sizes). *TSS* computes the time evolution of queue sizes on a sample path at instants t_0, $t_0 + \Delta$, $t_0 + 2\Delta$, $t_0 + 3\Delta$… Within each interval $[t, t + \Delta]$ time evolution of the queue size is approximated using diffusion approximation. The state of the system at time $t + \Delta$ is chosen randomly based on the state at time t according to the probability distribution obtained using diffusion approximation. We formulated equations for merging and splitting traffic flows enabling us to extend *TSS* to networks of queues. The

comparison of simulation results with *TSS* reveals high accuracy of the method and low computational cost.

We have also implemented the *TSS*. The software package consists of the computational engine (written using C++), and a Python scripting interface. The scripting interface enables the user to describe the network's topology, traffic flows, as well as events taking place in the network (such as link failures or conditional starting of new traffic flows).

WLAN related efforts: High-speed 802.11 WLANs with data transmission rates of up to 54 Mbps are now commonplace. Such WLANs operate with a distributed MAC protocol (Distributed Coordination Function) that involves carrier sensing, ACKs, and ARQ with backoff. So far, discrete event simulations (DES) have been the only way of studying the transient performance of such WLANs (though analytical steady state performance models are available). Due to the distributed nature of the MAC protocol, the number of simulation events per unit amount of simulated time grows not only with the transmission bitrate, but with the number of stations also. Thus discrete event simulation of even reasonably sized WLANs becomes time intensive.

Timestepped simulation (TSS) offers the ability to generate sample paths of the simulated network's state at a fraction of the cost of discrete event simulation. In this context, we have developed a model which emulates the characteristics of the 802.11 MAC layer to a timestepped model of a higher layer protocol. Specifically, the method predicts the instantaneous MAC layer throughput obtained by each station in a WLAN in each timestep. The method is agnostic to the transmission bitrate; in fact, the accuracy is better with higher transmission bitrates. Further, the computation in each timestep is linear in the number of stations.

Our model predicts the mean and deviation of the instantaneous overall throughput and the instantaneous per-station throughput. The simulator works as follows. In each timestep, a station is first decided to be active or inactive depending on its throughput in the previous timestep. For each inactive station, the throughput is sampled from the distribution N(0,1). For the active stations, the throughputs are first sampled from N(0,1) independently. Then the throughput vector of active stations is adjusted to account for the mean, deviation, and cross-correlation of the active throughputs. All these operations take linear time in the number of active stations.

Current work: Our current efforts are devoted to the following major goals:

- Evaluating the accuracy of *TSS* for large network topologies
- Researching the possibilities of modeling aggregate traffic streams (e.g., multiple TCP connections modeled as one traffic stream)
- Researching the ways in which step parameter Δ may be adjusted dynamically (e.g., based on whether the queue of interest is in steady-state or still in transient period)
- Researching the ways to use the method in simulation of large networks with conditional events (e.g., starting new traffic flows or making route changes based on the dynamic network conditions)

- Optimizing the computational procedures (especially numerical methods used to obtain probability distribution of the system's state after time Δ)
- Extending the software package with graphical user interface

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

Sub-Task 1: Inter-working the Session Initiation Protocol (SIP) with Mobile IP in UMTS

The minute details of building the session initiation protocol (SIP) over the mobile IP stack continue to be resolved in the context of the IP multimedia subsystem interface specifications. The challenge lies in effectively working out the authentication procedures used by the proxy call session control function (P-CSCF), the interrogating call session control function (I-CSCF), and the serving call session control function (S-CSCF) to ensure that mobility management functions are not ceded to SIP on the one hand and, on the other, that seamless real-time connectivity is ensured between the parties involved in the multimedia call in question. Many steps have been taken by various organizations in attempting to inter-work SIP with mobile IP, since the time we began this work. However, to our knowledge, none of these attempt to preserve all the advantages offered by MIPv6. The first stage of this effort, namely the simulation of a basic SIP over mobile IPv6 stack implementation, is expected to see results by the end of this summer; further enhancements and optimizations that are possible in the case of a mobile communicating with a wired local area network (e.g., fast mobility and hierarchical mobile IPv6) will be subsequently incorporated.

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

As stated in the preceding report (dated September 2004), work on developing a energy-throughput-and-delay-efficient routing and scheduling algorithm for quasi-structured and ad hoc extension subnets is in progress. The subnet is modeled as a directed graph whose vertices correspond to the communicating nodes; an edge exists from a given node to another given node if the signal to interference plus noise ratio of this directed channel exceeds a specified threshold. Route discovery and scheduling are based decisions derived from a simple statistical formulation involving these signal to interference plus noise ratio subject to the constraints noted in the last report. This work will see completion by the summer of this year. The Glomosim package is being used and it is hoped that the present work will add to this package, in particular the incorporation of small-scale channel effects in addition to the existing path-loss/shadowing and additive interference and thermal noise. The schemes under investigation are expected to provide complexity-performance tradeoffs that are attractive compared with current standards such as ad hoc on demand vector (AoDV) routing and dynamic source routing (DSR). Extensions of this work that incorporate multi-access scheduling functions based on signal to

interference plus noise ratio criteria will further enhance its actual applicability to cases such as the ad hoc mode of the IEEE 802.11 suite.

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 of Section 3.1 – Business Case Development

One 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 addressing this aspect. Two of our other papers, "Budgeting Process for Information Security Expenditures: Empirical Evidence" and "Evaluating Information Security Investments using the Analytic Hierarchy Process," are also related to this task. During this quarter, *Communications of the ACM* accepted the former paper for publication and the latter paper, previously accepted by *Communications of the ACM* , was slightly revised. Also, during this quarter, a third paper, "The Economics of Investment in Information Security," formerly published in *Transactions on Information and System Security*, was reprinted as a chapter in a book.

Another aspect of 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. During the previous quarter, we developed a framework that 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 an earlier meeting, 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.

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, 2005.

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 on pages 129-142 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 of Section 3.3 – 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 completed and submitted a paper, "Corporate Disclosures of Cybersecurity Activities," 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 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.

We have also begun work developing a paper to submit to the Fourth Annual Workshop on Economics and Information Security that will be helping this year at Harvard. Our paper will explore issues related to evaluation of information security in federal government agencies.

Related Papers

Gordon, Lawrence A. and 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.

Gordon, Lawrence A., Martin P. Loeb, "Corporate Disclosures of Cybersecurity Activities." Manuscript submitted for publication, December 2004.

Task 4 of Section 3.3 – 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.

Concluding Remarks

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 continue to be 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 at Carnegie-Mellon University on November 5, 2004. This seminar was well attended and included faculty and Ph.D. students from Carnegie-Mellon's Tepper School of Business and its Heinz School of Public Policy and Management, as well as from the University of Pittsburgh's Business School. In addition, L. Gordon gave a modified version of this presentation at the Computer Security Institute's (CSI's) Conference on November 8. Based on this last presentation, we were contacted by Micosoft, the Gartner Group, and Booze-Allan to discuss our research and its implications for the activities of their organizations. We have also had numerous other contacts based on our research, including with individuals from various accounting/consulting firms, government agencies, universities (including foreign universities such as University of Tokyo and The London School of Economics and

Political Economy), journals (i.e., reviewing work related to our sponsored research), corporations and an industry-lobbying group (Cyber Security Industry Alliance). We have also been asked to serve as guest editors for a special issue of *Information Systems Frontiers* on the topic of "Economics of Information Security." Also, during this last quarter, we have been asked to be affiliated with the Center for Human Enhanced Security Systems (CHESS), a new center within UMIACS. Finally, we have corresponded with a research team in Spain that is applying for a grant with CICYT, The Comision Interministerial de Ciencia y Tecnologia (Interministerial Committee on Science and Technology) based, on our research. CICYT is Spain's version of the NSF in the U.S.

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

Task 2 of 3.3 – Research in the Impact of Pricing Strategies

Our empirical research on the affect of pricing in competition among ISPs has shown that price does not appear to be the main differentiator among ISPs. Through an analysis of our longitudinal data of ISP competition, we have found that although the mean price for ISPs has gone down somewhat over the last two years, the decrease is not statistically significant. Our insight is that although technology advances and increased competition in the telecommunications sector would drive prices lower, this effect is opposed by the impact of increased industry consolidation. Therefore, the net impact of the changing dynamics of the industry is that prices may not change over time. It is also interesting to note that price dispersion among ISPs is also fairly stable over time. This means that there is consistent differentiation among ISPs to sustain heterogeneous business models. Furthermore, the results from our survival model indicate that an ISP's price is not a good determinant for survival. This is somewhat surprising because economic theory would suggest that those firms with lower prices are likely to survive and those firms with higher prices are likely to exit. Our interpretation of these results is that although price is important, it is not the most critical element in determining a successful business model. This is excellent news for ISPs that would like to invest in security technology because it means that they may be able to develop a revenue stream through higher prices that would support investment in security.

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.

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.

Related Presentations
"Technology Adoption as Resource Acquisition: An Empirical Study of Internet Service Providers," J. Bailey and T. Porterfield, *Decision Sciences Annual Meeting*, Boston, November 2004.

Task 5 – The Business Case for Wireless Systems

Our empirical analysis of ISPs has indicated that wireless technology adoption may improve firm performance if the timing of its adoption is done appropriately. Interestingly, the empirical analysis of ISPs shows that wireless adoption in 2002 actually lead to ISPs exiting the market. Between 2002 and 2004, the surviving firms increased their adoption of wireless technology from 10.2% to 12.8%. And, perhaps most importantly, the entering firms in the market have a wireless technology adoption rate of 27.0%. The nature of this technology is often characterized as a "first mover disadvantage" because early adopters fail and the later adopters succeed in its successful implantations. There are two main reasons for this. First, learning from the failed implementations is transferred to the non-adopters who can therefore adopt at a later time with fewer costs. Second, the later adopters benefit from technology maturation. This is especially true with wireless technology because of standards uncertainty and the falling prices associated with economies of scale and learning effects. These findings indicate very clearly that the timing of wireless technology adoption is very important to the overall success of including it in a business case.

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 in two international workshops---October 2004 (Rutgers University) and January 2005 (Dagstuhl workshop)---to a noted and distinguished group of economists and computer scientists and received very well. It is currently being prepared for journal submission.

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. We are continuing to develop cutting edge optimization technologies to solve these massive optimization problems. 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*.

Related Presentations

Multi-Period Traffic Routing in Satellite Networks. INFORMS Denver, October 24-27 2004. I. Gamvros, S. Raghavan.

Multi-Period Traffic Routing in Satellite Networks. INFORMS Computing Society Conference, Annapolis, January 5-7 2005. I. Gamvros, S. Raghavan.

"Technology, Infrastructure, and Resources: An Empirical Analysis of Internet Service Providers," J. Bailey and T. Porterfield, INFORMS Computing Society Conference, Annapolis, January 5-7 2005.

5. Optical Networking (Gary Carter and Joel Morris)

5.1 High Speed Experiments (G. Carter)

During this period we have been reorienting our research to investigate the possibility of carrying out network related experiments with the DRAGON network in collaboration with MAX (Mid-Atlantic Cross Roads). This network is an installed experimental WDM network funded by NSF. In close collaboration with LTS personnel we have identified a major thrust of this effort to affects of stochastic impairments in the physical layer on the network routing and management system. The impairment we have identified is PMD (polarization mode dispersion) which can vary in time due to variations in the fiber caused primarily by stress. Compensating for PMD was long thought to be a viable solution until it became clear that the expense of the compensators would be prohibitive in a multi-channel (WDM) system. We realized that it would be possible to detect the level of impairment due to PMD with relatively simple sensors at the physical layer. As the PMD varied, for example due to temperature variations, the impairment would change. By communicating the value of the sensor to the routing system it would be possible to make routing decisions based on the level of impairment.

To carry out this work we have partnered with Jerry Sobieski from MAX and one of the P.I.'s on the DRAGON project. He has brought in collaboration with MOVAZ which is a commercial company that makes MEMS based optical routers. The UMBC, MAX, MOVAZ group identified ways to migrate the sensor data to the routing system which has the potential for routing a channel to a new path if the current path exceeds a pre-determined PMD level (determined by the sensor). Currently we are now in the next planning phase which will include introducing models for the sensor output.

5.2 Statistical Signal Characterization (J. Morris)

1. DAIS simulation results have been obtained for two more codes. The first code is the RCD code with $\eta = 17$ (code length $n = 289$, information length $k = 240$, code rate $r = 0.8304$) [1]. The minimum distance of this code is 6 and the code has 11560 codewords at the minimum distance [2]. Consequently, one may compute the Union bound for this code. Figure 1 shows the DAIS simulation results (WER and BER), the Union bounds on the WER and BER, as well as the standard Monte Carlo simulation results, wherever available, for the $\eta = 17$ RCD code. Excellent agreement of the DAIS results with the standard Monte Carlo results as well as the Union bound is observed. In fact, the DAIS results always fall within the 99% confidence intervals of the standard Monte Carlo simulations, wherever both simulation results are available (i.e., at 5 and 6 dB). The DAIS simulation for the $\eta = 17$ code at $E_b/N_0 = 10$ dB required the transmission of 3.46×10^9 codewords in both the constrained and the unconstrained simulation. The standard Monte Carlo simulation for the same code at

the same E_b/N_0 would require the transmission of about 4.1×10^{18} codewords for recording 10 codewords in error. This represents a gain of more than 9 orders of magnitude of the DAIS technique over the standard Monte Carlo technique. Using DAIS, we have been able to estimate word-error rates as low as 2.44×10^{-18} for this code. Such low WERs are practically impossible to estimate via standard Monte Carlo simulations.

The second code evaluated via DAIS is the Euclidean geometry (EG) code with n = 255 and k = 175 [3]. This code is a one-step majority logic decodable code [3] and has a minimum distance of at least 17. However, neither the exact minimum distance nor the number of codewords at the minimum distance are known for this code and, hence, the Union bound for this code cannot be computed. One of the chief reasons for evaluating this code was to validate the DAIS technique for codes with larger minimum distances than previously considered (all DAIS evaluations presented up to this point have been for codes with $d_{\min}$ = 6). Figure 2 shows the DAIS simulation results (WER and BER), and the standard Monte Carlo simulation results, wherever available, for this n = 255 EG code. Again, the DAIS results are seen to fall within the 99% confidence intervals of the standard Monte Carlo simulations, wherever both simulation results are available (i.e., at 4, 4.5, and 5 dB), thus indicating excellent agreement of the DAIS results with those of standard Monte Carlo. We observe that, in this case, the DAIS technique provides results for performance evaluations down to WERs as low as 10^{-26}.

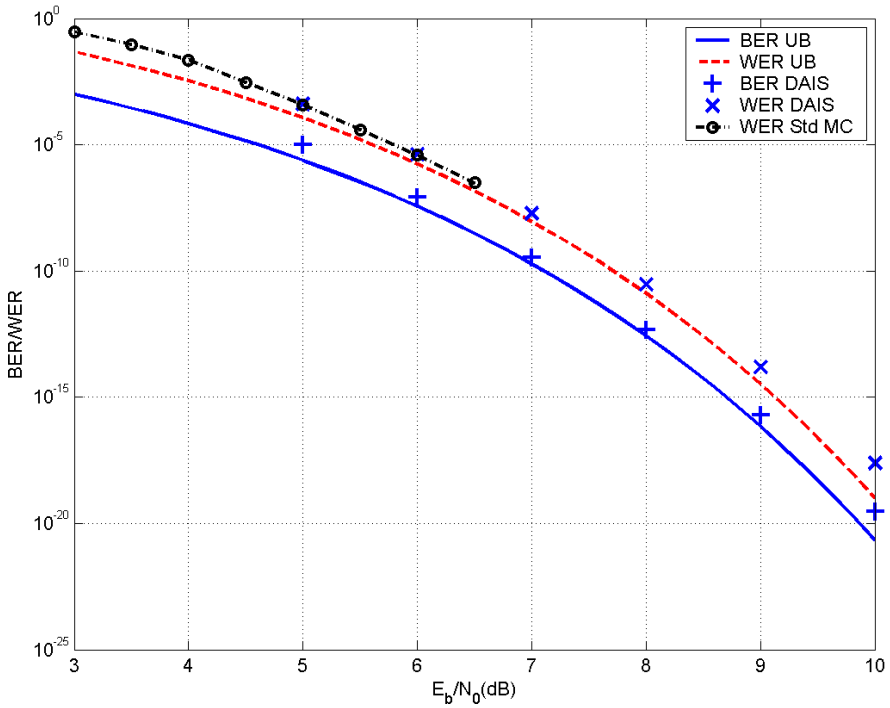

Figure 1: Union bound, standard Monte Carlo, and DAIS results for BER and WER for the η = 17 RCD code.

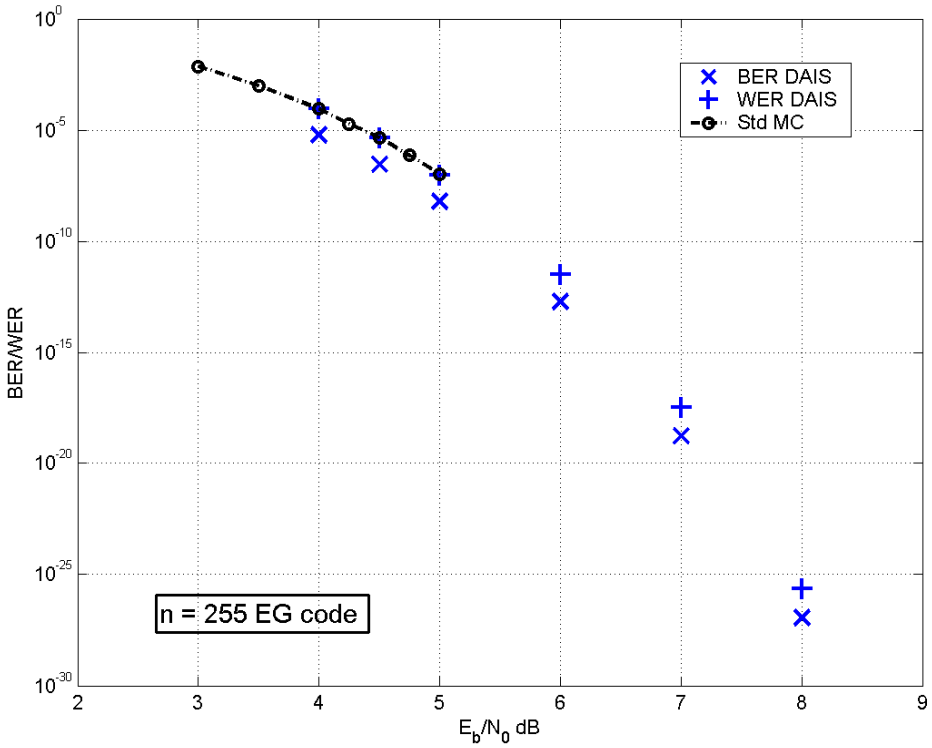

Figure 2: Standard Monte Carlo simulation and DAIS results for BER and WER for the n = 255 Euclidean Geometry code.

In addition to evaluating different codes using DAIS, we have also made some modifications to the program and added some new capability (this research partially supported by Maryland TEDCO). Importantly, a version of the DAIS software in C++ that can be used over different computing platforms such as Linux, Windows, and Macintosh has been developed. This program is also easy to read and incorporates a sound object-oriented design (the previous version, although it used C++, was haphazardly written and did not employ a true object-oriented structure). An easy to use graphical user-interface (GUI) has also been developed for the Windows platform. The algorithm that performs the scaling operation to combine the results of the constrained and unconstrained simulation has also been modified and made more robust. This new scaling algorithm considers only those unconstrained-simulation control-variable bins that have more than a user-specified fixed number (say 100) of word error events recorded, in order to perform the scaling. Finally, our paper that introduces the DAIS technique and presents performance evaluation results for a n = 96, k = 50 LDPC code is scheduled to appear in the Feb. 2005 issue of *IEEE Communications Letters* [4].

2. As mentioned in a previous report, 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. With an aim of improving FEC code performance on the optical

channel, soft-information may be introduced in the detection device by allowing for two decision thresholds and, hence, three decision levels. Arbitrary placement of the two decision thresholds will result in a binary-asymmetric channel with asymmetric erasures (BAC/AE) in general. Further, due to the asymmetry of the space and mark pdfs, it may not be possible to position the two decision thresholds to provide a binary symmetric channel with symmetric erasures (BSC/E) characterization. However, it is possible to adjust the two threshold locations to equalize either the error probabilities that define the transition from 0 to 1 and vice-versa, or the erasure probabilities that define the transition from 0 to *E* and 1 to *E*, where *E* represents an erasure. The former gives rise to a binary symmetric channel with asymmetric erasures (BSC/AE) characterization, while the latter results in a binary asymmetric channel with symmetric erasures (BAC/SE) characterization.

The task of manipulating the two thresholds to arrive at a BSC/AE characterization is observed to be much simpler than the manipulations needed to arrive at a BAC/SE. Also, although it is not clear whether equalizing the error probabilities is more beneficial as compared to equalizing the erasure probabilities, intuitively there seems to be some merit in allowing for equal error probabilities. Hence, we focus our attention on setting the two decision thresholds in order to obtain a BSC/AE characterization of the OFC channel with the effective noise modeled via chi-squared pdfs.

In the previous report, we described how one could compute the capacity and capacity-achieving prior probability distribution given knowledge of the transition probabilities of the BSC/AE. All that remains to be done to compute the capacity of the BSC/AE resulting from three-level decisioning of a realistic OFC is to develop a systematic technique to set the thresholds. Next we outline such a technique.

Let $f_0(x)$ and $f_1(x)$ represent the pdf of the space and mark, respectively. The pdf $f_0(x)$ is a central chi-squared pdf, while $f_1(x)$ has the form of a non-central chi-squared pdf. Let t_{ML} denote the ML decision threshold, and e_0 and e_1, the transition probabilities that define the resulting BAC model. Then

$$e_0 = p(1|0) = \int_{t_{\mathrm{ML}}}^{\infty} f_0(x)\,dx$$

and

$$e_1 = p(0|1) = \int_{0}^{t_{\mathrm{ML}}} f_1(x)\,dx\,.$$

Let t_l and t_r represent the left and right threshold, respectively, that we intend to set in order to obtain a BSC/AE characterization. In order to set the thresholds t_l and t_r, all we must do is specify a value for the equal transition error probabilities $\gamma = p(1|0) = p(0|1)$. Since it makes sense to have $t_l \le t_{\mathrm{ML}} \le t_r$, we constrain γ to satisfy $\gamma \le e_{\min} = \min(e_0, e_1)$. Figure 3 provides a diagrammatic view of the pdfs $f_0(x)$ and $f_1(x)$ and the different thresholds, while figure 4 shows the resulting BSC/AE channel model.

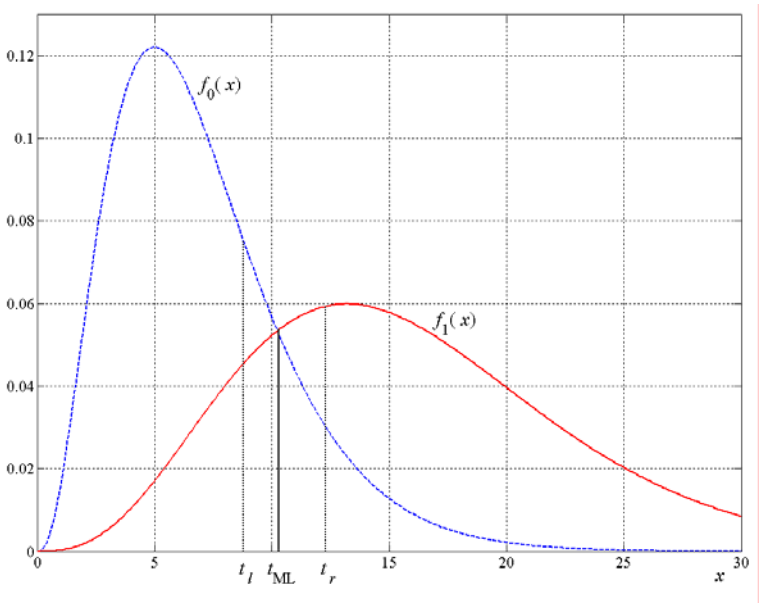

Figure 3: The pdfs $f_0(x)$ and $f_1(x)$ and the different thresholds t_{ML}, t_l, and t_r.

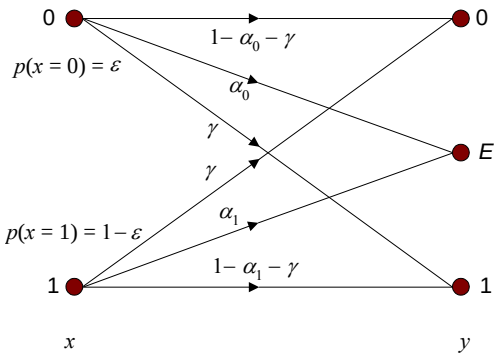

Figure 4: The binary symmetric channel with asymmetric erasures (BSC/AE) channel model

Once γ has been specified, we can numerically compute the thresholds t_l and t_r (say using the bisection method or secant method), since the following equations

$$\gamma = p(1|0) = \int_{t_r}^{\infty} f_0(x)\,dx$$

$$\gamma = p(0|1) = \int_{0}^{t_l} f_1(x)\,dx$$

must be satisfied. Here, we note that the computation of the two thresholds can be done independently. Once t_l and t_r have been computed, it is a simple matter to compute the erasure probabilities, since

$$\alpha_0 = p(E|0) = \int_{t_l}^{t_r} f_0(x)\,dx$$

and

$$\alpha_1 = p(E|1) = \int_{t_l}^{t_r} f_1(x)\,dx\,.$$

For fixed $f_0(x)$ and $f_1(x)$, the optimal choice of $\gamma \le e_{\min} = \min(e_0, e_1)$ (optimal in the sense of maximizing capacity) may be determined via numerical optimization. Figure 5 shows a plot of capacity versus the fraction $\gamma / e_{\min}$ for the BSC/AE for the OFC with the effective noise modeled via chi-squared pdfs and for parameters M = 3 and β = 14. For this M and β, we have e_0 = 0.001598 and e_1 = 0.001834 for the BAC resulting from single threshold ML decisioning. The capacity for this BAC is 0.9818 bits per transmission. From figure 5, we observe that the maximum capacity of 0.9891 bits per transmission is achieved when $\gamma / e_{\min}$ = 0.3. This corresponds to a BSC/AE with parameters γ = 0.000479, α_0 = 0.005815, and α_1 = 0.004296. Note that for the BSC/AE with $\gamma / e_{\min} = 0.3$, the erasure probabilities are roughly 10 times the transition error probability. Also, the combined probability of having an erasure or transition error for this BSC/AE channel (i.e., $\gamma + \alpha_0$ or $\gamma + \alpha_1$) is roughly three times the probability of having a transition error in the corresponding BAC channel (e_0 or e_1). Yet, the BSC/AE has a higher capacity than the corresponding BAC, although by only about < 1%. This is a strong indication of the benefits to be gained by introduction of minimal soft-information at the decoder.

Additionally, based on intuition, one would expect that a good choice for the BSC/AE model from an error control coding perspective would result when the erasure probabilities α_0 and α_1 are roughly twice the transition error probability γ. This follows since error control codes can correct twice as many erasures as errors. However, the validity of this claim needs to be investigated further via generation of code performance curves for the different BSC/AE channels.

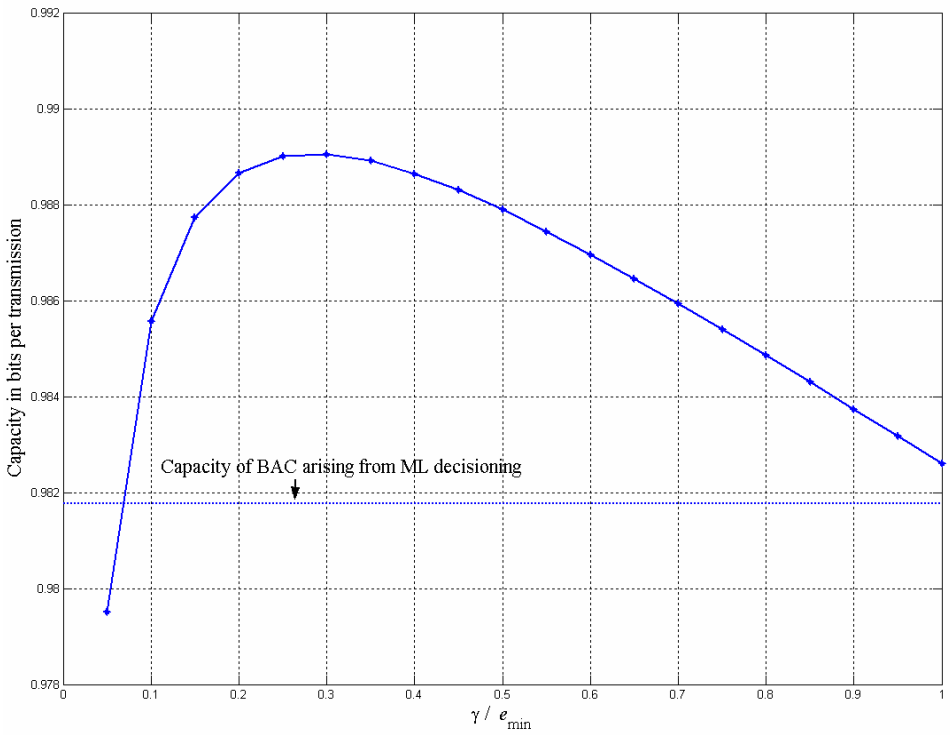

Figure 5: Capacity (bits) versus the fraction $\gamma / e_{\min}$ for the BSC/AE for the OFC with the effective noise modeled via chi-squared pdfs for parameters $M = 3$ and $\beta = 14$.

3. In the previous progress report, we discussed the fully-parallel implementation of the BFA decoder for the (255,175) Type-I EG LDPC code [1], which was augmented to include an input and output shift-register, the ability to allow for the BFA iterations to terminate as soon as the parity-check matrix is satisfied, and the capability to directly output the received word in the event of a decoder failure. This augmented version of the decoder written in structural Verilog has been debugged and further tested using waveform analysis. The test results indicate the decoder to be functioning as desired and all minor bugs seem to have been fixed. Once the testing is completed, we intend to convert the Verilog implementation to a corresponding layout for fabrication on a chip, which may then be interfaced with an optical fiber communications system testbed at UMBC.

   Along similar lines, effort has begun to also implement a fully-parallel BFA decoder for the $\eta = 23$ RCD code of length $n = 529$ and code rate $r = 0.873$ in Verilog. This BFA implementation will then be used as a basis to construct an extended-BFA decoder for the same code. Recall that the extended-BFA algorithm performs decoding on a binary symmetric channel with erasures (BSC/E). The ultimate goal of this exercise is to fabricate a chip that implements the extended-BFA decoder and use it for FEC code performance experiments on the UMBC testbed, and possibly in an appropriately designed optical fiber communications receiver.

References

[1] A. Mahadevan, "On RCD Codes as a Class of LDPC Codes: Properties, Decoding, and Performance Evaluation", PhD dissertation, CSEE Department, UMBC, Baltimore, MD 21250, May 2004 (anticipated).

[2] W. Martin, "The WEF for a class of regular LDPC Codes: The RCD Array Codes", PhD Dissertation, CSEE Dept., University of Maryland Baltimore County (UMBC), Catonsville, MD 21250, 2003.

[3] S. Lin and D. J. Costello, *Error Correction Coding: Fundamentals and Applications*, Prentice Hall, Engelwood Cliffs, NJ, 1983.

[4] R. Holzlöhner, A. Mahadevan, C. R. Menyuk, J.M. Morris, and J. Zweck, "Evaluation of the Very Low WER of FEC Codes Using Dual Adaptive Importance Sampling", to appear in *IEEE Communication Letters*, Vol. 9, No. 2, Feb. 2005.

6. Peer to Peer Networks and Network Behaviors *(M. Marsh)*

Lookup in Unstructured Networks

(*In collaboration with Bobby Bhattacharjee, Aravind Srinivasan, Jonathan Katz, and Sudarshan Chawathe .)*

Distributed applications are always constructed on some network topology. While it is often possible to optimize this topology for any particular aspect of an application, it is typically not the case that such optimization can be done for all aspects simultaneously. This leads to trade-offs in design. For instance, while efficient algorithms exist for performing decentralized data lookups, these algorithms require a tightly constrained overlay network topology. When other, more significant, optimizations preclude such a topology, it becomes necessary to employ a lookup algorithm that is reasonably efficient on an arbitrary topology. Previously, we have developed such an algorithm, called local minima search (LMS).

Based on LMS, we have designed a decentralized public key infrastructure (PKI) on the web-of-trust model (similar to PGP). Although LMS is a probabilistic algorithm, and hence cannot provide perfect assurance, through simulations we find that the probability of successfully locating any item (such as a public key) can be made arbitrarily close to 1 at the cost of a modest increase in network load. In addition, extensive simulations comparing an LMS-based PKI with other reasonable decentralized PKI designs (using a real-world sample web of trust) demonstrate that our design is more efficient for a given probability of success. This is true even when comparing with a distributed hash table (DHT), which provides best-case lookup efficiency for a highly structured network, and for which we have made very generous assumptions. The reason for this is that public keys must be accompanied by a certificate chain demonstrating their validity, and with LMS such a chain is constructed as a by-product of the lookup procedure (though this chain is often not

the shortest possible).

Starting from a complete existing implementation of the base LMS protocol, we have been developing an implementation of the PKI. We hope to be able to deploy this implementation in the near future. In addition, we are developing more sophisticated search techniques for the PKI, which at present is limited to matching email addresses to public keys.

Emergent Network Behaviors

Large networks pose significant challenges for simulations, yet simulations are typically the only way to assess a system's likely performance before an actual deployment. Most network simulations operate at either the packet level, which does not scale well to large networks, or the protocol level, which does not capture low-level interaction dynamics.

Other scientific disciplines motivate looking for a layered approach to both modeling and simulation of large networks. It is often the case that the complex interactions between elements at one layer lead to emergent behaviors that define a higher layer. For example, the quantum interactions of valence electrons in atomic physics lead to the basic interactions of chemistry, and these interactions are much simpler to calculate as chemical reactions than as solutions to the same problem formulated using the underlying quantum mechanics.

How to extend this to computer networks is not yet understood. One could imagine several possibilities. The most obvious is to aggregate traffic passing through border routers, effectively replacing each autonomous system (AS) with a single virtual host. While our initial investigations into such aggregation are encouraging, we do not have a model for how individual packet flows aggregate. It is also unclear whether chaotic effects will at some point dominate any attempts to perform aggregation. Other possible behaviors, not involving ASes as a fundamental unit, are also being considered.

This work is still somewhat preliminary, and it is difficult to predict how it will develop. While it is possible that some insightful new way to describe network dynamics will emerge, it is also possible that we will find that it is not possible to construct such a dynamics. This, however, would also be an interesting result, as it has implications for future simulations of large networks, as well as the ability to extrapolate behaviors from smaller simulations.