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

Joseph JaJa
University of Maryland Institute for Advanced Computer Studies (UMIACS) and
Department of Electrical and Computer Engineering

This report represents work completed during the time period Oct 1, 2003 through January 31, 2004 under the LTS/UMIACS Contract. The table below gives a list of the projects and the pages where related progress is described.

1. Active Network Management (Participating Faculty: Mark Shayman, Samrat Bhattacharjee, Steve Marcus, Ray Chen, and Richard La)

1.1 Overlay-based Networks Services (S. Bhattacharjee, R. La, and M. Shayman)

This project represents work under Tasks (Active Network Management Techniques) and 3 (Simulation and Experimental Testbed)

We are continuing our work on using in-network overlays for network management, monitoring, and security. Specifically, our current effort is in monitoring DoS attacks that are initiated within the domain. Detecting DoS attacks at the source has several benefits, the obvious being that network resources in the ISP/backbone network are not abused. A somewhat obtuse benefit comes from avoiding legal issues that can ensue when particularly destructive attacks are initiated from a given domain.

In our work, we are constructing a distributed attack detection system which uses in-network overlay nodes to probabilistically sample two-way traffic and detect attacks using predefined signatures. There are several advantages to our approach:

- It is more scalable than a single IDS system at the network edge since smaller detectors (which do not have to scale with domain size) can be used.
- Attacks can be detected quicker than in monolithic detection systems since packets can be examined closer to the source (before they are multiplexed with other traffic).

A particularly novel component of our work is the development of an advanced network interface card (NIC) that logs all incoming and outgoing packets at individual hosts. The NIC is constructed such that the packet logging cannot be stopped or the log erased by host software. Once the in-network overlay nodes trigger an alert, the host packet logs can be examined to verify the alert. We believe this technique will substantially reduce false positive probabilities while improving the overall efficiency of the entire detection system. In our current work, we are:

- Developing algorithms to compress header information such that several minutes worth of data can be stored in a few megabytes of storage.
- Developing primitives which would be broadly useful in the NIC in detecting a broad range of attacks.
- Acquiring and programming NIC hardware for use in a testbed to verify our entire system.

1.2 Overlay-based Multicast Traffic Engineering (S. Bhattacharjee, R. La, and M. Shayman)

This project represents work under Tasks 1 (Active Network Management Techniques) and Task 3 (Simulation and Experimental Testbed).

Implementing multicast services in the application layer has drawn considerable attention since IP multicast service is not widely available at the network layer. With the use of a logical overlay network, data is transferred using unicast transport services. 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. Recent proposals to overcome this limitation using network coding appear to be impractical for realistic networks. However, our work suggests that proper integration of *source* coding with multipath load balancing can generate a practical solution to optimally distribute multicast load.

The incoming flows are separated into segments (segments may or may not overlap.) and are encoded using a rateless code such as LT or Raptor coding. The source will continue to send coded packets of a particular segment until all receivers acknowledge successful reception of the segment being sent. Receivers can successfully decode the data when they receive any distinct $K+\varepsilon$ packets, if the original segment was consisting of K packets. Therefore, we do not need to do any bookkeeping in the sense that we do not need to keep track of which packets are sent to each receiver. We only need to guarantee that each receiver should receive a certain amount of distinct packets. This gives us the flexibility to send packets to each receiver along multiple paths in an efficient way. Besides, due to the rateless nature of the coding, a packet loss does not create a major problem in the decoding procedure unlike network coding.

Multiple paths between source and destination pairs are established using overlay nodes. The paths from source to overlay nodes and from overlay to destinations are calculated according to min-hop principle. It is assumed that the depth of the overlay network is one in the sense that there exists only a single node between a source node and destinations along any given path. However, the definitions may be modified to cover a setting where the depth of overlay network is greater than one.

The algorithm extends the one we developed for unicast load balancing using Simultaneous Perturbation Stochastic Approximation (SPSA). Optimal load balancing is achieved by running SPSA based stochastic approximation algorithm per each source-receiver pair in an asynchronous and distributed manner. In other words, we treat each source-receiver pair as unicast and obtain its splitting ratios using SPSA. However, while sending the data to the overlays, we make use of the multicasting nature of the traffic. Particularly, when a source node sends packets to overlays, instead of sending at a rate that is sum of all path rates passing through the given overlay, we only send with maximum path rate passing through the overlay. This allows us to minimize the link stress. Since the algorithm is measurement based, with a proper convex cost function we are able to reach the optimal solution in an iterative way.

1.3 Multiclass Traffic Engineering Using Selective Overprovisioning (S. Bhattacharjee, R. La, and M. Shayman)

This project represents work under Tasks 1 (Active Network Management Techniques) and Task 3 (Simulation and Experimental Testbed).

We have started working and developing a traffic engineering and multi-path routing algorithm for QoS provisioning. In particular, the algorithm can be used in a network with overlay nodes or in an MPLS network. The general research trend for QoS provisioning in a network relies on per-flow or per-class bandwidth provisioning and scheduling in the routers. However, these algorithms assume that the routers are capable of classification and scheduling of the packets in wire-speed, and furthermore, the network manager can set up the corresponding router parameters appropriately. In practice, these assumptions do not hold and service providers resort to over-provisioning to provide QoS. In this approach, routers use simple FIFO queues with no sophisticated classification and scheduling mechanism. In order to maintain the QoS requirements, the utilization of the links is kept at a relatively low level. The over-provisioning approach simplifies the routers and the network management requirements by sacrificing network resources such as bandwidth. For the real-time applications such as VoIP over-provisioning factors of 2 to 4 are recommended, which means that the utilization of the links should be limited to between 25 and 50 percent.

In our proposed routing algorithm, based on the customer demand, a subset of the paths in the network are used to carry the real-time traffic (possibly along with best effort traffic), and the required over-provisioning factor is only maintained on these paths. The rest of the network links do not need to be over-provisioned since they are only carrying best effort traffic. In this way, the routers inside the network do not need to perform complicated scheduling and classification tasks and also most of the links can be fully utilized since they are only carrying best effort traffic. The problem is to determine how much real-time traffic (if any) and best effort traffic should be placed on each path.

We have formulated this problem as a non-linear optimization problem. The over-provisioning constraints are introduced indirectly as additional terms in the cost function. We have used the gradient projection method to solve the problem and have gotten some promising initial results. However, it is shown that the problem is inherently a non-convex optimization problem, and hence we plan to extend the algorithm and use global optimization techniques such as simulated annealing to be able to reach the global optimal solution. We also plan to study and consider the practical issues related to distributed realization of this algorithm in a network such as feedback delay, measurement noise and asynchronous update of the

1.4 Fast Timescale Multiclass Traffic Engineering (S. Bhattacharjee, R. La, and M. Shayman)

This project represents work under Tasks 2 (Stochastic Control) and Task 3 (Simulation of Experimental Testbed).

The project considers an MPLS network with two classes of service--a high priority class consisting of voice and video traffic and a low priority class consisting of TCP traffic. In the previous Progress Report, we described our work on the *flow migration problem* and the *duplication problem*. During the current period, we have implemented controllers for each problem and evaluated them using the policy evaluation algorithm TD(0). Then we strengthened the evaluation with packet-level simulation.

In the flow migration problem, control strategies were developed to migrate high priority flows on congested paths onto alternative paths while taking into account potentially adverse consequences for best effort (TCP) traffic on the alternative paths. We have designed our controller in a LQG (Linear Quadratic Gaussian) context. Then we improve this controller using Rollout Algorithm. With TD(0) and packet-level simulation, we compared our controller with other controllers. The empirical study demonstrates the effectiveness of our migration controller under the bursty network environment

In the duplication problem, we exploit an unused backup path in order to increase the quality of service of high-priority traffic. Note that many network providers already setup backup paths for each active data path. These backup paths will protect active paths upon their failure and they are unused by high priority traffic in normal conditions. We propose a scheme in which duplicates of high-priority packets are transmitted in the backup path when congestion is detected in the main path. We have designed a controller that is based on Certainty Equivalence Control (CEC). Then we improve this controller using Rollout Algorithm. With TD(0) and packet-level simulation, we compared our controller with other controllers. Our empirical study demonstrates the effectiveness of our intelligent duplication controller under the bursty network environment. This empirical study shows a clear advantage of the duplication scheme, especially when we duplicate intelligently. It can enable network providers to increase their high-priority traffic share without compromising the quality of service. Moreover, our study points out the advantage of using traffic prediction based on traffic models for video and voice.

*1.5 **Markov Decision Modeling for Integrated MPLS/WDM Traffic Engineering (R. La, S. Marcus, and M. Shayman)***

This project represents work under Task 2 (Stochastic Control), Task 3 (Simulation and Experimental Testbed), and Task 4 (Integration with Optical Layer).

We have fine-tuned the topology reconfiguration policy and we have used the simulation to compare the results from this policy with that of a heuristic policy and a fixed policy that does not use state information.

The fine-tuning was aimed at reducing the computational load of the decision maker and at the same time maintaining the performance of the algorithm. At each slow time scale step, when we have to pick a branch exchange (BE) from a set of possible BEs, if we were to compare all BEs in the set of possible BEs it would be very time consuming. We can rule out a large portion of the BEs without running an internal simulation to

compare them with the others. We have had an algorithm to limit the action space but the old algorithm eliminated some good BEs from the list as well, and as a result it would reduce the performance of the algorithm. This was improved to enhance the performance while reducing the time of computation.

The number of dropped calls as a result of topology reconfiguration has always been a concern. We have developed call migration to reduce the number of dropped calls.

Our simulations show that the rollout policy performs significantly better than the heuristic policy (as expected). Also when comparing the fixed policy (which is computed based on a deterministic traffic matrix) with our policy, as the random part of the traffic demand grows larger our policy performs better. In particular, when the congestion occurs locally in the network our policy performs much better than the fixed policy. In case of global congestion (all S-D pairs have high volumes of calls) reconfiguration rarely results in improved blocking rate because when all S-D pairs are overloaded no matter how we select our topology all lightpaths are going to be overloaded.

In addition to comparing the reward and the number of serviced calls during a day for each of the policies above, we are now comparing the call blocking ratio resulting from each of the policies.

We are working on a paper to be submitted to Globecom by the end of February.

We are looking into trying to change the reward function to not include delay and instead put a limit on the end-to-end delay for each S-D pair. This will bring linearity to the reward function and may help us to use linear programming to solve the moderate time scale problem of bandwidth assignment. At the moment we are running a heuristic algorithm for the moderate time scale.

1.6 Integrated Logical Topology Design and Traffic Grooming for Reconfigurable MPLS/WDM Networks (M. Shayman)

This project represents work under Task 3 (Simulation and Experimental Testbed) and Task 4 (Integration with Optical Layer).

The goal is to develop algorithms that optimize the configuration of the optical layer and the set of MPLS label switched paths that are routed over the logical topology that results from configuring the optical layer. In this work, it is assumed that the traffic matrix is known. The traffic matrix specifies the expected traffic rate between each source-destination pair. When a significant change in the traffic matrix is observed, the algorithm(s) are triggered and re-optimize the MPLS/WDM network. Our objective function is to maximize network throughput, the portion of traffic that can be accommodated in the network.

Most of the existing work in this area has described the problem using integer linear programming (ILP). However, finding optimal solutions using ILP is computationally

prohibitive even in small networks. So, most algorithms proposed use simple heuristics making use of traffic demands, logical hop count, physical hop count, etc.

We propose two algorithms for the logical topology design and LSP path selection algorithm. Each algorithm makes iterative use of the *linear programming* algorithm for solving multi-commodity flow problems. We refer to the proposed algorithms as the Addition Algorithm and the Deletion Algorithm. The Addition Algorithm finds a feasible topology by adding a link in each step. It works as follows.

1. Configure full mesh graph setting each link to be temporary with its capacity infinite.
2. Solve the multi-commodity flow problem with current network topology.
3. Find a link which is temporary and has maximum link utilization among temporary links.
4. Mark the selected link as permanent and assign it the capacity of one wavelength, and delete all temporary links violating degree constraints.
5. Go back to step 2 until all remaining links are permanent.

The Deletion Approach is similar to the addition approach except it deletes a link with minimum utilization rather than adding a link with maximum utilization. The Deletion Algorithm works as follows.

1. Configure full mesh graph, each link capacity is that of a wavelength.
2. Solve the multi-commodity flow problem with current network topology.
3. Find and delete a link which has minimum utilization among the links that violate the degree constraints.
4. Go back to step 2 until all nodes satisfy degree constraints

Our algorithm tries to find a near-optimal solution for logical topology design and LSP allocation problem using multi-commodity flow optimization approach. Since the problem can be defined as LP rather than ILP, we can find the solution within polynomial time. Currently, we analyze the performance of the proposed algorithms using simulation. In this simulation, we configure the 16-node NSFnet topology and use MATLAB. We compare the proposed algorithm with other simple heuristic algorithms such as HLDA, MMHA, and simple lightpath deletion algorithm.

Publications

K-I. Lee and M. A. Shayman, Single and multipath logical topology design and traffic grooming algorithms in IP over WDM networks, *International Conference on Computer Communications and Networks,* Dallas, Texas, October 2003.

K-T. Kuo, S. Phuvoravan, S. Bhattacharjee, R. La, M. Shayman and H-S. Chang, On the use of flow migration for handling short-term overloads, *IEEE Globecom,* San Francisco, California, December 2003.

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

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

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

1.7 GMPLS Actively Managed WDM Testbed (R. Chen)

This project represents work under Task 4 (Integration with Optical Layer) and Task 5 (Interaction Between Active Network Management and MPLS Network Stability).

Introduction

Generalized-MPLS (GMPLS) is being standardized as a unified control plane to facilitate intelligent interoperations among various types of data layers and optical layer. Typical modeling methods for GMPLS networks consist two extreme cases, which are the overlay model and the peer model. Under the peer model, all layers are supervised under the same control plane, sharing information among each other. Since in the peer model each node has complete knowledge of the network status, resource allocation can thus be more efficiently optimized. It is generally recognized that the peer model will be more widely deployed in future networks. Our research of traffic engineering is based on the peer model. In this report we document the traffic engineering considerations and procedures in our actively managed WDM testbed.

The main objective of TE is by means of routing and signaling mechanisms to optimize the use of network resources while, at the same time, satisfies traffic requests with QoS constraints. In a dynamic network environment, computing routes with multiple QoS constraints is usually a NP-Complete problem. [1] Simple and efficient heuristics are always preferred. A common approach is to utilize constrained shortest-path first (CSPF) algorithms that are extended from standard SPF algorithms like the Dijkstra algorithm. The CSPF algorithms perform shortest path computation based on a simple graph extracted from the physical network topology and calculate a feasible route (though may not be optimal) that satisfies all QoS constraints of the request. The simple graph represents the current state of the corresponding network and QoS performance limitations of different network components. After a route is successfully calculated, the signaling protocol, such as CR-LDP [2] and RSVP-TE [3], is invoked to establish the LSP and reserve network resources along the LSP. However, in time of network congestions, high priority QoS requests may preempt existing low priority ones. In this case, before the establishment of new LSPs, selected low priority LSPs need to be preempted so as to give room to high priority ones. The preempted low priority LSPs may be rerouted after the establishment of high priority LSPs. Figure 1 illustrates the flow chart of the TE heuristic.

Network structure

Using the peer model, a GMPLS node can be viewed to consist of two layers: the optical layer and the electronic layer. The major components at the optical layer are wavelength selective optical Crossconnect switches (OXCs) that execute switching at wavelength or high bandwidth granularity [4]. The structure of the generalized GMPLS node is illustrated in Fig. 2. For such GMPLS node, there may exist three types of forwarding schemes: O-O-O bypass, O-E-O switching, O-E-O label-switching/store-and-forwarding, each of which is associated with a specific level of quality of service (QoS), with O-O-O bypass having the highest QoS and O-E-O store-and-forwarding having the lowest QoS. In our current research, we have examined typical QoS factors including bandwidth, delay and traffic priority, all of which have been incorporated into the proposed TE routing heuristic.

Graph representation model

One of the key ingredients for traffic engineering is routing. To carry out efficient routing, we proposed a simplified graph representation model to construct an auxiliary graph based on the physical network topology and the node structure proposed above, as shown in Fig. 3. In this abstracted graph, for each wavelength the optical layer is represented by two logical nodes which stand for the aggregation of input ports of the OXC and the aggregation of output ports of the OXC, respectively. Each wavelength on a physical link is represented by a logical link, and if a lightpath is established between nodes, that lightpath is also represented using a logical link and its constituting physical links are removed from the original graph. In the auxiliary graph, each link has an associated tuple metric (*b, d, c*), where *b* represents the reserve-able bandwidth of the link, *d* represents the delay of the link, and *c* represents the assigned cost to the link which can be either calculated or administratively assigned according to carrier's policy. By adjusting link cost *c* in the auxiliary graph, different TE objectives can be achieved.

Path computation

When a traffic request arrives at the network with specific QoS requirements, a shortest path with respect to the cost metric is calculated. Based on the graph model, we proposed an efficient Dijkstra SPF-based algorithm called two-pass Dijkstra, which calculates a delay-constrained least-cost route for individual request. The two-pass Dijkstra algorithm first calculates projected delay QoS from the destination node to all other nodes, then it starts its second Dijkstra pass from the source node by pruning links whose projected QoS does not satisfy the QoS requirement. The priority of traffic requests is incorporated into the cost metric and, during network congestions, high priority traffic requests may preempt a set of low priority traffic requests, which are determined during the path computation. We also compared the impact of centralized preemption policy versus distributed preemption policy to the traffic engineering heuristic.

References

1. Z. Wang and J. Crowcroft, “Quality-of-Service Routing for Supporting Multimedia Applications”, IEEE JSAC, Vol. 14, No. 7, Sept 1996.

2. L. Andersson, et al., “LDP Specification”, RFC 3036.

3. D. Awduche, et al., “RSVP-TE: Extensions to RSVP for LSP Tunnels”, RFC 3209, Dec 2001.

4. “GMPLS Testbed Based on Optical Ethernet”, Zhonghua Zhu, Aihua Guo, Wenlu Chen, Wei Chen and Yung. J. Chen, LEOS Annual meeting, Tucson, (2003) (Invited).

Figures

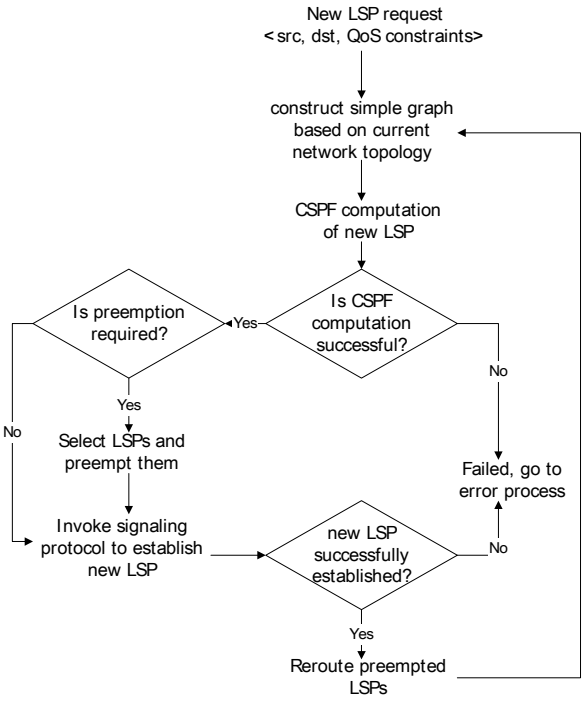

Fig. 1 A TE Heuristic

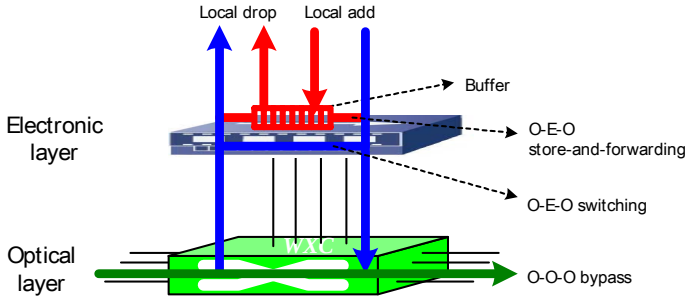

Fig. 2 *Generalized GMPLS node*

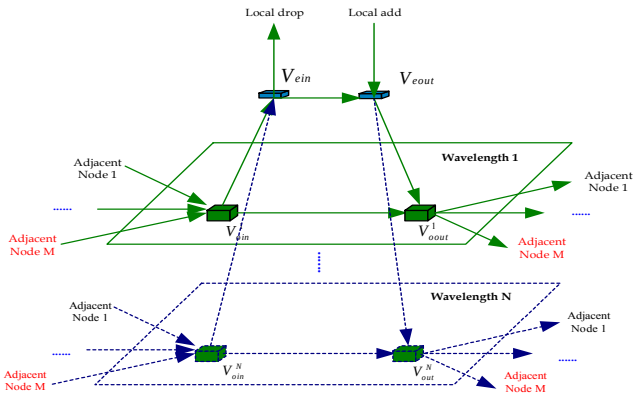

Fig. 3 *Graph representation model*

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

This project represents work under Task 1 (Costs for Systems Security Management).

We abandoned the FPGA effort due to a lack of experienced students working in the synthesis environment and focused our attention on using our previous prototype board- the EBSA-285.

Using the EBSA-285, we now have the capability to detect the installation of malice into the kernel within thirty seconds of its installation. The cost of this detection capability is approximately 1% overhead on the memory bandwidth of the protected host. To date, we have tested twelve different "root kits" found on the Internet and we have detected all twelve. We have not yet found a root kit that we can not detect. The results of this effort have been submitted to the USENIX Security Conference 2004. A copy of the paper is attached.

Our plans for the future of this project include:

0. Identifying a new board that includes a network interface for out of band communications between boards in both a distributed and centralized fashion.
1. Reconstitution of the system after the detection of malice: this may either mitigating the effects of the malice or completely removing it.
2. The detection of unauthorized privilege elevation, i.e. a process gains administrator or root privileges without authority.
3. Configuration management testing, enforcement, and remote patch installation.

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

3.1 WLAN Location Determination (A. Agrawala and U. Shankar)

This project represents work under Task 2 (Wireless Local Area Network Research)

We have been working on location determination based on WLAN signal strength. Briefly, the approach involves (1) offline construction of a radio map of AP signal strengths, and (2) online determination of client position based on correlating currently received signal strengths to the radio map.

Previous Work: Our previous work addressed the following issues:

- Developing clustering techniques to reduce the energy consumption of WLAN location determination systems
- Handling correlation
- Gaussian approximation
- Handling small-scale variations

Recent Work: We addressed the following issues:

- Analytical analysis methods for performance
- Analyzing the effect of the user profile on performance
- Allowing continuous-space radio map
- Studying the effect of the calibration parameters

In recent work, we developed an analysis method for studying the performance of WLAN location determination systems. The method can be applied to any of the WLAN location determination techniques. It does not make any assumptions about the signal strength distributions at locations, independence of access points, or the user profile.

We also studied the effect of the user profile on the performance of the WLAN location determination systems. We used the analytical method to obtain the optimal strategy for selecting the user location, which is not implemented by any of the current WLAN location determination systems. We also used analytical methods to study the effect on performance of averaging multiple signal strength vectors.

The results show that averaging multiple signal strength vectors reduces the variance of the resulting distribution and hence reduces the overlap between distributions. The less the overlap, the less the error.

We used simulation experiments to validate the analytical results and to study the effect of user profile on the performance of the location determination systems. The results show that incorporating the user profile in the location determination system can enhance the accuracy significantly when the available hardware is limited. However, when a reasonable number of access points can be heard at each location, the

performance of the location determination system is consistent under different user profiles.

We also developed two techniques to allow the continuous tracking of a node, instead of returning a discrete radio map location. The two techniques are: the *Center of Mass* technique and the *Time-Averaging* technique. Using the *Center of Mass* technique, the performance of our system, called *Horus,* is enhanced by more than 13% for the first testbed and more than 6% for the second testbed compared to the basic technique. The *Time-Averaging* technique enhances the performance of the *Horus* system by more than 24% for the first testbed and more than 15% for the second testbed. The two techniques are independent and can be applied together.

We also compared the performance of the *Horus* system to the performance of the *Radar* system from Microsoft. Our results show that the *Horus* system is more accurate than the *Radar* system by more than 11 feet, on the average, for the first testbed and more than 9 feet, on the average, for the second testbed. Moreover, the tail for the distribution of the distance error for the *Horus* system is significantly better than that of the *Radar* system. We also showed that the *Horus* system leads to more than an order of magnitude savings in number of operations required per location estimate compared to the *Radar* system.

We also studied how sensitive the Horus system is to the calibration parameters. Our experiments show that a training time of 15 seconds per location is enough to construct the radio-map information for the *Horus* system. The *continuous-space estimator* component of the system allows us to choose the radiomap locations as far as 14 feet while maintaining the high accuracy of the system.

Recent Implementations: We have implemented a Java based 'Calibrator' which is used for collecting data for the radio map, and a 'Tracker' which keeps track of the location of a user as the user moves around. We used both these tools for the 4th floor of AVWilliams building. The current tracker implementation has a 2.5 second averaging delay. In spite of this delay the tracker keeps track of a person walking around to a very high degree of accuracy. We are in the process of making formal measurements on the precision of this implementation. This implementation verifies the analytical developments done to date.

References

Moustafa Youssef, Ashok Agrawala, A. Udaya Shankar, "**WLAN Location Determination via Clustering and Probability Distributions, "** IEEE International Conference on Pervasive Computing and Communications (PerCom) 2003, Fort Worth, Texas, March 23-26, 2003. http://www.cs.umd.edu/~moustafa/papers/percom03.pdf

Moustafa Youssef, Ashok Agrawala, A. Udaya Shankar, and Sam H. Noh, "**A Probabilistic Clustering-Based Indoor Location Determination System, "** CS-TR 4350, Department of Computer Science, University of Maryland, College Park, March 2002. http://www.cs.umd.edu/~moustafa/papers/locdet_tr.pdf

Moustafa Youssef, Ashok Agrawala, " **Small-Scale Compensation for WLAN Location Determination Systems, "** IEEE Wireless Communications and Networking Conference (WCNC) 2003 New Orleans, Louisiana, March 16-20, 2003. http://www.cs.umd.edu/~moustafa/papers/wcnc03.pdf

Moustafa Youssef, Ashok Agrawala, **"On the Optimality of WLAN Location Determination Systems, "**CS-TR 4459, Department of Computer Science, University of Maryland, College Park, March 2003. http://www.cs.umd.edu/~moustafa/papers/opt_tr.pdf

Recent Publications:

Moustafa Youssef and Ashok Agrawala, "**Handling Samples Correlation in the Horus System**", IEEE Infocom, March 2004. http://www.cs.umd.edu/~moustafa/papers/corr.pdf

Moustafa Youssef and Ashok Agrawala, **"On the Optimality of WLAN Location Determination Systems",** Communication Networks and Distributed Systems Modeling and Simulation Conference, January 18-24 2004, San Diego, California. http://www.cs.umd.edu/~moustafa/papers/cnds04.pdf

Moustafa Youssef and Ashok Agrawala, " **Handling Samples Correlation in the Horus System",** CS-TR 4506, Department of Computer Science, University of Maryland, College Park, June 2003. http://www.cs.umd.edu/~moustafa/papers/cs_tr_4506.pdf

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

This project represents work under Task 2 (Wireless Local Area Network Research) and Task 7 (Traffic Engineering Across the Wireless Core Network).

***Introduction*:** QoS in WLANs is an active area of research. It has been known that even the simplest form of QoS - throughput fairness- is not provided by the random-access version (DCF) of the 802.11 MAC. There is a time-scheduled MAC protocol (PCF), but it is not implemented so far and also suffers from higher service-times at low loads due to the polling involved. Existing studies on QoS focus on the MAC, specifically the Binary Exponential Backoff (BEB) algorithm used. But almost all 802.11 wireless cards limit BEB to 4 retransmissions, which precludes long-term unfairness.

Previous Work: In previous studies, we showed experimentally that physical layer capture is the major reason for long-termunfairness (i.e., over intervals longer than 1 second). 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, even if the stronger signal arrives after the weaker signal. 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 developed a compensatory QoS mechanism that augmented the DCF mode with a cyclic link layer control. In each cycle, clients provide current load estimates to the AP and the AP provides the clients with their "fair" share in the next cycle. The outgoing traffic of the clients is shaped at the queue between the OS device driver and the wireless card. We implemented this mechanism in Linux clients and a hostAP based Linux AP. The control information exchanged between clients and AP was piggybacked transparently on the regular 802.11 packets.

We extended the span of our experiments to include multiple UDP and TCP sources working in infrastructure mode. The results confirm the importance of the physical layer capture effect on fairness in 802.11b WLANs. Specifically, for four TCP sources the difference in throughputs may be as high as 100%. More details may be found in "*The Impact of Physical-Layer Capture on Higher-Layer Throughput in 802.11b WLANs*".

Recent Work: We have examined several network simulators that handle 802.11 links, including *ns2*, the most popular open source simulator, and *QualNet*, a state-of-the-art and probably the best commercial simulator. We found that they do not account for the physical capture effect. Basically, a collision is resolved in favor of the stronger signal iff it arrives first, whereas our work shows that this happens even if the stronger signal arrives after the weaker signal. Hence the results they provide can be very erroneous.

We fixed the ns-2 simulator, and in the process discovered other flaws in its 802.11 component; specifically, ns-2 never allows colliding RTS/CTS frames to be resolved in favor of the stronger frame, even if the stronger frame is the first to arrive! We also communicated this flaw to the Qualnet support staff, and they have since scheduled a fix to Qualnet in their next release. A paper has been submitted for publication.

References

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

Andre Kochut, Arunchandar Vasan, Udaya Shankar, Ashok Agrawala. **An Empirical Characterization of Instantaneous Throughputs in 802.11b WLANs.** Available at http://www.cs.umd.edu/Library/TRs

3.3 Characterizing User Behavior and Network Performance in 802.11: The Wireless Side (A. Agrawala and U. Shankar)

This project represents work under Task 7 (Traffic Engineering Across the Wireless Core Network).

Introduction: We have studied the behavior of 802.11b environments and have found that there is a significant variability in the performance of different wireless cards in use today. In particular we found that the receiver sensitivity affects the throughput and delays. The signal strength reported by the firmware depends on the implementation

approach taken by the manufacturer. Some cards report the signal strength as a fine-grain value while the others may only report it as a coarse-grain value. Cards use the signal strength as a basis for deciding the transmission rate. When the coarse-grain values are used the transmission rate may change frequently.

Previous Work: In order to better understand the workload and the wireless traffic we captured the packet and frame traffic over a 24 hour period which is being used as the base for creating suitable analytical models for the workload. We are starting with the models that have been reported in the literature and evaluating their applicability to the data we have collected.

We conducted 24-hour measurement in A.V. Williams Building with three sniffers. We have those sniffers capturing the IEEE 802.11 traffic, which are transmitted to/from only one AP. To merge those three traces, we consider one sniffer's timestamp as a reference timestamp for time synchronization. Then we apply the least square fitting method to convert the other two sniffers' timestamps to be synchronized with the reference sniffer's timestamp. We applied the technique during every interval between two beacons, which are commonly received by the three sniffers.
We successfully merged three traces into one big trace, obtaining the synchronization errors of less than 2 microseconds on average and 30 microseconds at maximum. This result well satisfies our error requirement, i.e. less than 50 microseconds, which is the minimum inter-frame time between any two 802.11 frames.

Using the merged trace, we conducted some preliminary traffic workload characterization of wireless LAN traffic. For the purpose of comparison, we also analyzed some well-known traffic traces, e.g. Sigcomm 2001 wireless LAN traces, Lawrence Berkeley Laboratory WAN traces and Bellcore Ethernet LAN traces.
We obtained the traffic counting process C_n and inter-arrival time process A_n from each trace and applied self-similar and multi-fractal analysis techniques to those traces. The followings are the results we obtained:

- We obtained the first-order statistics for MAC/PHY level, such as type and retransmission distribution, number of fragments per type, number of frames in Power-Save mode, Disassociate/De-authentication reasons and per-AP statistics.
- Wireless LAN traffic shows very high degree of non-stationarity, which results in non-Gaussian bimodal distribution of counting process C_n.
- Self-similar analysis on wireless traffic produces incorrect estimation of Hurst parameter due to such high non-stationarity.
- By applying multi-fractal analysis, we examined *spikeness* (multi-fractality) of traffic. If the adjacent traffic signals are characterized to be far different from each other, then such traffic is said to be spiky (or multi-fractal). Our result shows that the wireless LAN traffic in A.V. Williams Building has similar (low) spiky characteristics to well-known Bellcore Ethernet LAN trace. However Sigcomm 2001 wireless LAN trace shows high multi-fractal characteristics very similar to that of WAN trace. Sigcomm 2001 traffic is more like WAN because in such

conference setting, very few traffic is internal, i.e. most of the traffics flow from/to the network outside.

These results are being documented at present.

Recent Work: We studied the performance of wireless LAN measurement techniques. We compared three techniques, i.e. SNMP measurement, measurement at wired vantage point and our wireless monitoring technique. The followings are the results of the performance comparison:

- SNMP provides precise traffic statistics only for inbound (To-AP) traffic, which is the traffic from the wireless stations to the Access Point (AP). The outbound (From-AP) statistics are incorrect because it cannot know whether the packets departing from the AP are successfully transmitted to wireless stations or not.
- Measurement at wired point also shows incorrect statistics for From-AP traffic due to the same reason as SNMP.
- Our wireless monitoring technique provides the statistics (packet numbers, bytes and errors) in both To-AP and From-AP traffics very close to (nearly 100%) what we can obtain in end-to-end measurements. Moreover, it can provide more detail MAC traffic statistics, such as number of packets retransmitted once, twice and three times.

Another study we have done recently is to determine the best locations for wireless sniffer in order to improve the performance of wireless monitoring technique. By placing the wireless sniffer in proper location, we can have the sniffer observing more traffic than end-to-end measurement can collect. For this purpose, at different locations we measured signal strength and SNR (Signal to Noise Ratio) of the signal from the AP, the traffic through which we want to measure. Then we obtained a contour map where the contour line indicates the locations with the same signal condition. Using the contour map, we can find the coverage area of the target AP, so that we can place the sniffer at the center of the coverage area. With the contour map, we can also detect some bad locations for sniffers, where signals are fading more rapidly than other locations.

References

Jihwang Yeo, Suman Banerjee, Ashok Agrawala**. Measuring Traffic on the Wireless Medium: Experience and Pitfalls**. University of Maryland, CS-TR-4421, Dec 2002. http://www.cs.umd.edu/users/jyeo/TR.pdf

3.4 Z-iteration for WLAN/WAN (A. Agrawala and U. Shankar)

This project represents work under Task 7 (Traffic Engineering Across the Wireless Core Network).

***Introduction*:** Performance evaluation of TCP/IP networks is problematic. Current analytical models do not adequately capture many properties important to researchers and network designers. Packet-level simulation is reasonably accurate but prohibitively expensive for realistic topologies and workloads (e.g., link speeds of 10,000 packets/s).

We have 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. Internally, the method models the IP layer by a time-dependent queuing network and the transport sources by time-varying stochastic processes. Different TCP versions and router queuing disciplines (e.g., FIFO, RED, CBQ) can be handled. Comparisons against ns2 simulations show that the method scales well and provides reasonable accuracy.

Previous Work: The previous version of the Z-iteration software, called NetSolver, was implemented in Visual C++/Windows and had a restrictive GUI. We completed a platform-independent implementation of the Z-iteration software, called NetSolver. The implementation is in ANSI C and has a Python scripting interface which allows one to conveniently define networks and TCP/UDP/IP workload, including starting and stopping flows at pre-defined times as well as in response to logical events (e.g., start flow 2 when flow 1 has transferred 2MB or encounters delay greater than 200 ms. This latter capability is not present in the network simulators we are aware of (including the de facto *ns* simulator).

Our experimental studies of 802.11b WLANs (outlined in the **WLAN QoS** section) have brought us much closer to obtaining the needed throughput profiles for 802.11b clients, not only in determining the quantitative values but also the functional dependencies (e.g., should the throughput at time t depend on the number of competing stations at time t *and their relative signal strengths*).

A 3-hour tutorial on the Z-iteration was given in June 2003 at the ACM SIGMETRICS conference, the premier forum for computer systems performance evaluation. Slides are available at http://www.cs.umd.edu/~shankar

We have worked on improving the accuracy of the solver, refining the equations modeling the time evolution of metrics and incorporating arrival processes more complex than the existing nonstationary Poisson, including phase type distributions and their solutions using matrix geometric methods.

Recent Work:

WLAN QOS and WLAN model for Z-iteration: We have identified the correct model for physical layer capture in current 802.11b cards from our experimental data and our novel measurement technique. Specifically, we show that in a collision, the stronger frame is almost always retrieved even if it starts second (provided it starts within the physical layer preamble of the previous frame). This has significant impact on the way simulators model physical layer capture. Both NS-2 and Qualnet do not account for this behavior. We have communicated our findings to the maintainers of NS-2 and Qualnet, and Qualnet folks have already scheduled a fix.

We are currently developing a WLAN model for Z-iteration. Existing analytical models provide handles on throughputs computed over long durations (order of seconds) under the unrealistic assumption of saturation (i.e., each wireless node always has a packet to send). On the other hand, we are interested in obtaining a handle on the short-term (e.g. 50 msecs) throughputs obtained by different nodes as our ultimate objective is to use this in a heterogeneous (wired cum wireless)network with TCP workloads.

The 802.11 MAC causes throughputs to be unfair on the short term due the random choices of backoff values. For instance, over 10 packets sent in the WLAN, node A could have sent just one, while B could have sent everything else. Owing to the high overhead of the MAC, a 10 packet horizon would correspond to more than 50 msecs at 2Mbps. Therefore, it is important to model this unfairness appropriately. We are trying to obtain a model for predicting short-term throughputs and associated unfairness.

Z-iteration for M(t)/M/1/K networks: We have validated z-iteration solver for M(t)/M/1/K networks. In simple cases of a single M/M/1/1 and M(t)/M/inf queues we proved formally correctness of the method. For M(t)/M/1/K queues as well as networks of such queues we used simulation and statistical methods to explain validity of functional approximation used by z-iteration. We are currently working on a more formal proof for a case of M/M/1/K queue.

Z-iteration for TCP/IP networks and hybrid siimulator: We have implemented flexible python-based interface to the solver (for both unix and windows environments). We are currently working on developing a better approximate equation for IP queue. We are interested in a forlmula that could be applied to compute the distribution of the state of the queue after larger time interval (on the order of 0.5 sec.). We plan to use it for fast solving of IP networks with TCP traffic. We also plan to use the equation as a key part of hybrid simulator. The hybrid solver will combine discreete event simulator with system of equations modeling queue and TCP dynamics.

3.5 Assertional Security Analysis of 802.11 (W. Arbaugh)

This project represents work under Task1(Trust Based Routing in Support of Ad-Hoc Networking Protection) and Task 2 (Wireless Local Area Network Research).

Previous Work:

We started to develop an assertional reasoning framework for security properties, with the expectation of obtaining the same precision and power that assertional reasoning brings to standard correctness (i.e., safety, progress, compositionality) properties of concurrent systems.

We developed a simple but powerful extension to the assertional formalism to reason about security properties of distributed systems in general, including properties of authentication, confidentiality, and key management. We will use this extension to develop *provably secure* protocols for wireless and wired networks.

The fundamental extension is to introduce for each system of interest (i.e., each principal) a state variable, called **key set** and denoted Kset, that represents the set of keys and potential secrets maintained by the system. Kset consists of (1) all keys generated by this system, (2) all keys received from other systems, (3) all strings that may potentially be keys (e.g., the key space that an intruder system may search through). For a regular user, Kset is updated whenever the system generates a key, receives a key (usually encrypted by another key that is already in the system's Kset), or deletes a key. However, the Kset of an *attacker* system can also be updated by dictionaries of potential passwords and keys, pairs of encrypted and plain text, and other such information relevant to security attacks. In any case, Kset denotes the space in which the attacker would search for potential matches.

The importance of the Kset variables is that it allows one to reason about security properties within the well-understood framework of assertional reasoning, without the need for new logics (e.g., BAN) and interpretations. First, security properties can be expressed by standard predicates and safety and progress assertions in system program variables (including Kset). Second, they can be verified using the standard assertional proof rules. Third, they can be tested using our testing harness.

Current Work: We are applying the theory to 802.11 protocols and to peer-to-peer systems. In the latter, we are developing a common service specification for various peer-to-peer systems, including Gnutella, Chord, and NICE.

3.6 ***Wireless Hand-off and Security (W. Arbaugh and U. Shankar)***

This project represents work under Task 1 (Trust Based Routing in Support of Ad-Hoc Networking Protection), Task 2 (Wireless Local Area Network Research), and Task 3 (Protocol and Functionality Requirements).

During the period October 2003 and January 2004, we completed the following work:

a. (Task 3) Previously, we abstracted a wireless hand-off into three phases: The probe or discovery phase, the reassociation phase, and the authentication phase. Our prior work with IAPP reduced the latency of the reassociation phase from approximately 12 ms to 1.2 ms- an order of magnitude improvement. During this reporting period, we completed implementations and made experimental measurements for the remaining two phases.

   We first completed work on “Proactive Key Distribution” which pre-positions key material ahead of a mobile station via modifications to the AAA server and the RADIUS protocol. Measuring the results experimentally, we found that we were able to reduce the cost of authentication from approximately 800 ms for a full EAP/TLS to 20 ms. The results of this effort appeared in the IEEE Wireless Communications Magazine, February 2004.

In the probe phase, we once again applied Neighbor graphs. However, this time we augmented them with non-overlapping graphs. The combination of these two algorithms permitted us to reduce the probe time from several hundred milliseconds to adjusted (taking into account are not in the firmware) value of 30 ms. The results of this effort will appear in MobiSys 2004.

This is gives a combined layer 2 latency of approximately 50 ms. Some what higher than our goal of 35 ms, but significantly better than the over 1.2 second previously. Our focus in the next reporting period will be on combining our previous implementations into a single cohesive experiment and report the results to a transactional journal for publication.

b. (Task 2) Work is continuing with the IETF EAP working group on formalizing the state machine for EAP. A second draft is currently being prepared and will be submitted to the working group in a few days. The chairman of the working group, Bernard Aboba of Microsoft, is extremely pleased with our contributions to date.

c. (Task 1) We have completed a paper on probabilistic routing and have submitted it to Mobihoc 2004. The results of the program committee will be released shortly, and we are hopeful the paper will be accepted.

*3.7 **Efficient IP-Based UMTS Networks (J. Thomas)***

This project represents work under Task 5 (Efficient IP-Based UMTS Networks Devising Schemes for Reliable Mobile Packet Transmission).

Sub-Task1: Mobile IP – Compatibility of MIPv4 and MIPv6 Based on Dual Stack Model

Experiments with MIP v4 to test for functionality and potential vulnerabilities have been completed in cooperation with BBN Technologies. The incorporation of the dual stack model into the Dynamics (Helsinki) package has been under progress, since 22 September 2003 when IP v6 connections became available at our laboratory. With reference to the configuration shown in Figure 1 below, the “lab router” is being configured for functionality and most of the obstacles to the implementation of this architecture now appear to have been overcome.

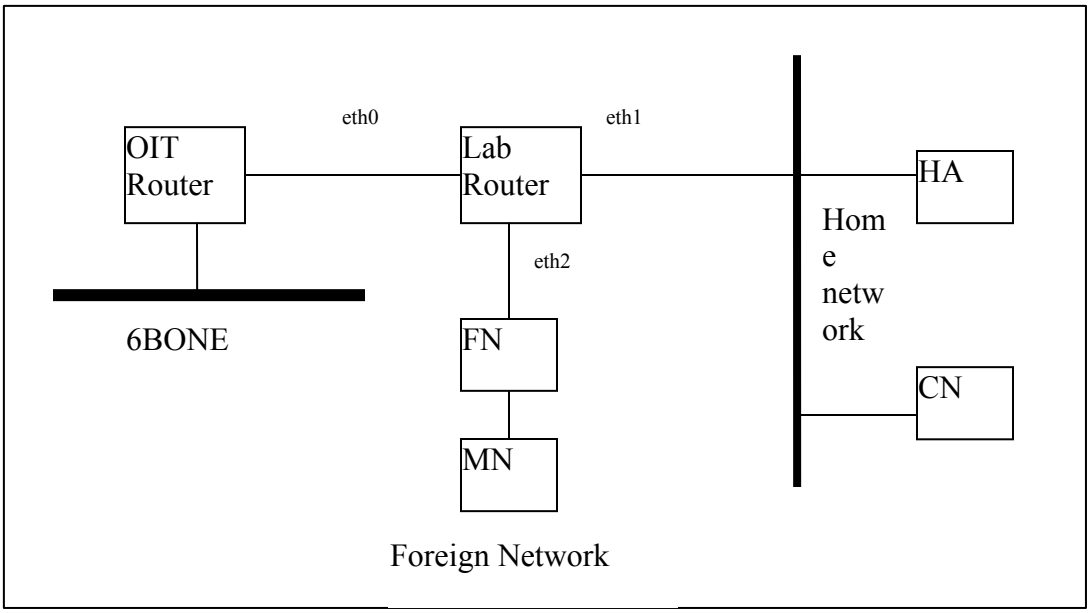

Figure 1: Mobile IPv6 testbed configuration at UMBC laboratory

Sub-Task 2: Integration of MIP in ad hoc networking environments

Our interest in this area revolves around two issues, namely ad hoc address autoconfiguration and load balancing. Thus far, we have focused primarily on the first of these issues. Most address allocation protocols employ the notions of network merge and partition, and IP release and reuse. However, this implies that every node is involved in some information exchange when an entering node seeks an address. The consequent overhead and increased latency are undesirable. We seek to emulate the scalability of wireline networks which owe this attractive feature to their inherent hierarchical structure. Could we employ ideas similar to clustering as in routing in ad hoc networks? If clusters can be used here, only clusterheads need to keep track of address allocation. Each clusterhead keeps a copy of the address information for the whole network. When a new node joins a cluster, the clusterhead chooses an unused address and polls other clusterheads. This address is assigned to the entering node only when all the other clusterheads agree. With clustering, most of the packets are sent out by unicast; this reduces traffic overhead in ad hoc networks. Since the clusters effectively constitute a partition, this approach allows network merge and partition relatively easily. We continue to study means of effectively tracking changes in cluster configuration.

Sub-Task 3: Stream Traffic over Mobile IP

There are essentially two standards for supporting voice over IP, namely the ITU-H.323 and the IETF SIP (session initiation protocol). In H.323, stream traffic (i.e. real-time traffic, namely voice and video) are encapsulated into RTP (a protocol over UDP for fast handoff and low latency). SIP is an application layer protocol designed to be independent of lower layered transport protocols. H.323 does not support host mobility. The IETF SIP, on the other hand, supports personal mobility. Each user has a unique ID

which she registers with the current IP address to the SIP server. When user A wants to setup a session with user B, SIP servers find the current location of B using B's registration information. Mobility support allows a change in user location (IP address) during an active session in a manner that is transparent to the user. Mobile IP is a potential solution here. Mobile IP and SIP have some similar features: mobile host registration, message relays (Home Agent, SIP server), message redirect (Mobile IPv6). We can integrate SIP server with Home agent to make SIP support mobility. Two ways have been used for this integration, namely by modifying SIP (with inspiration from MIP) or by running SIP over MIP. The latter approach causes some duplication since the mobile host has to register with both the SIP server and the home agent (or have the SIP server query the home agent about the mobile host's location). The main issue in mobile stream IP sessions is that of smooth handoff. The foreign agent hierarchy reduces the frequency of home agent registration and the handoff delay by changing the remote registration to local registration (with foreign agent). To reduce packet loss due to handoff, lost packets can be forwarded from previous foreign agent to current foreign agent. However, this may be unsuitable in real-time environments.

Our preliminary experiments used the testbed configurations shown in Figures 2 with one home agent (HA: 130.85.95.52), one correspondent node (CN: 130.85.95.50), one mobile node (MN: 130.85.95.49), two foreign agents (FA1: 130.85.93.96 and FA2: 130.85.168.175). Dynamics (MIPv4) runs on the HA, the two FAs and the MN, while Linphone (a voice over IP application using SIP) runs on the CN and MN. The following sequence was executed.

1. HA, FA1 and MN are set up (MN binds to FA1 and registers with HA);
2. An SIP call connection between CN and MN is set up (voice transmission begins);
3. FA2 is set up (MN can see both FA1 and FA2, MN stays with FA1, voice transmission continues);
4. FA1 is turned off (voice transmission stops immediately; MN detects this change after ~60 seconds, and then switch to FA2; after MN receives the confirmation of the new registration, voice transmission resumes.)

The handoff delay when the MN moves from FA1 to FA2 is over 90 seconds, which is unacceptably high. The reason for the MN's inability to detect loss of contact with FA1 immediately relates to the lifetime of the router advertisement. If the MN fails to receive another advertisement from the same agent within the specified lifetime, it assumes that it has lost contact with the agent. So the MN has to wait till the lifetime of FA1 expires to decide whether or not to register with FA2. The lifetime parameter can be tuned by the FA configuration file. Reducing the lifetime reduces the handoff delay but destabilizes both the MN and the FA. If registration is moved to step 3 then the handoff delay can be reduced significantly. This requires that HA keep both the bindings of MN-FA1 and MN-FA2. Currently, when HA receives the registration message of MN with FA2, HA will the remove the binding of MN-FA1and replace it with MN-FA2. HA forwards packets to both FA1 and FA2. So even if FA1 is down, MN can still receive packets from FA2. When MN receives duplicated packets, it can just discard the duplicates. When MN loses contact with FA1, it will inform the HA to delete the MN-FA1 binding. Another method

is for the MN to send message to FA1 through FA2 to let FA1 forward packets to FA2 during the time of registration with FA2.

The testbed configuration with foreign agent hierarchy is shown in Figure 3. Here, both FA1 and FA2 are up at the same time, so that the MN can see both FAs. Linphone was used with this two-layer foreign agent hierarchy (with ping programs) to test handoff between two local foreign agents. The result was that the MN chose FA1and FA2 alternately; the handoff delay when switching FAs is less than one second. Running Linphone during the handoff does not affect audio transmission. The handoff latency is the same as that between two foreign networks.

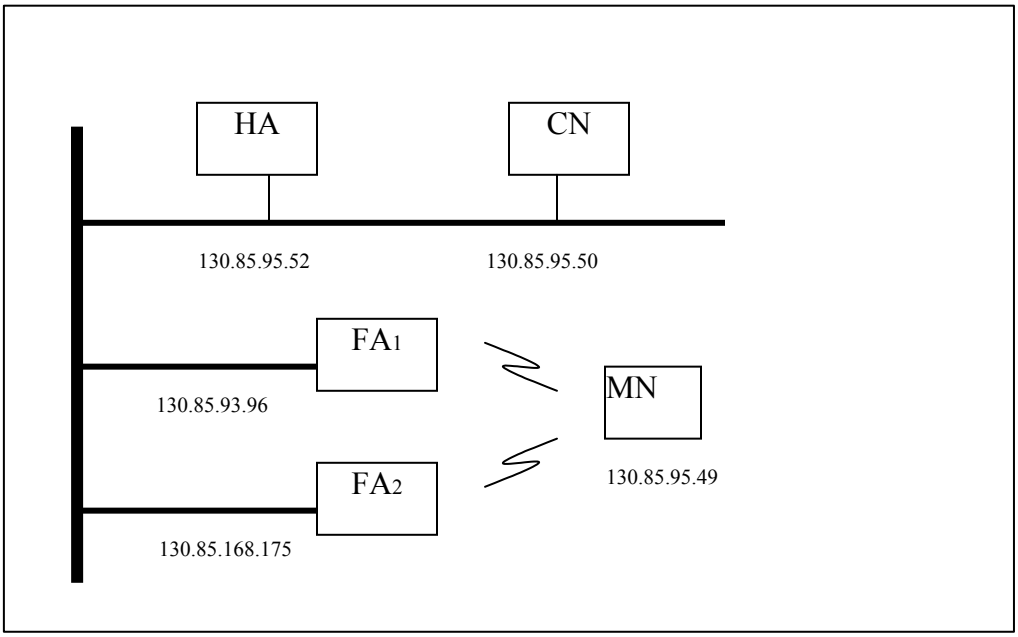

Figure 2: Handoff between two foreign networks

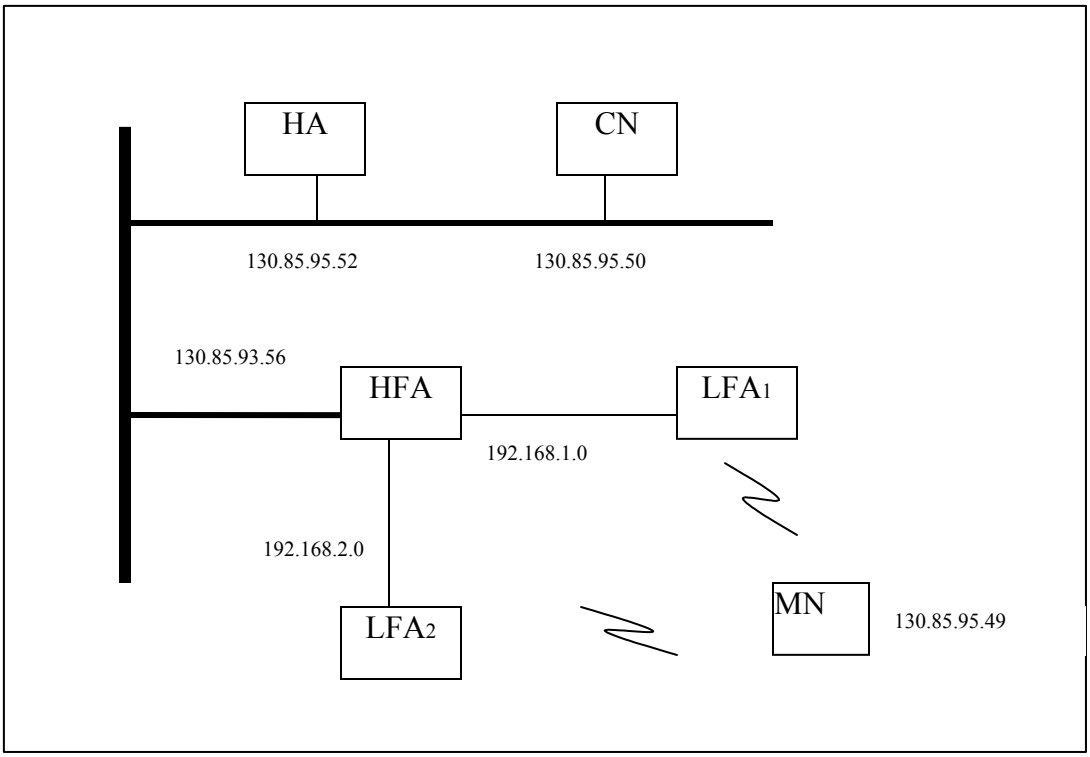

Figure 3: Handoff in a two-layer foreign hierarchy

Thus, the key question now is to discover how the MN can switch to FA2 when FA1 is abruptly turned off. If the HA can keep both bindings of MN-FA1 and MN-FA2, then even if the MN-FA1 tunnel is down, the connection should remain alive. We continue to study this and related issues.

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

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

This project represents work under 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 the last quarterly report, several papers were either published or accepted related to this task (see below list of related publications).

During this quarter, we completed another paper, entitled "Budgeting Process For Information Security Expenditures: Empirical Evidence," and submitted it for publication. In addition, during this quarter, our paper, entitled ""Expenditures on Competitor Analysis and Information Security: A Management Accounting Perspective," was published as a chapter in the book entitled Management Accounting in the Digital Economy out of Oxford University Press. Finally, during this quarter, our paper, entitled "The Economics of Investment in Information Security" and originally published in *ACM Transactions on Information and System Security*, was selected for reprinting in a book called Economics of Information Security co-edited by faculty from Harvard University and Cambridge University.

Publications

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

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

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

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.

4.2 Economic Impact of Government, Industry, Academic Partnerships (L. Gordon and M. Loeb)

This project represents work under 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. The task not only seeks to examine potential benefits from joint ventures by such combinations, but also to examine incentive issues related to realizing these potential benefits. We (along with William Lucyshyn) have addressed these issues in the context of modeling the information sharing of computer security breaches and attempted breaches. The paper, entitled "Sharing Information on Computer Systems: An Economic Analysis," demonstrates that information sharing by firms result in each firm spending less on information security. Although firms spend less on security, we provide conditions by which sharing leads to overall higher levels of information security. Furthermore, we examine the problem of free riding by partner firms, and examine ways of enhancing the incentive system. Our work in this area is directly related to Information Sharing Analysis Centers (ISACS) and the Department of Homeland Security's strategy to encourage information sharing to secure cyberspace. The paper was published this quarter in the November-December 2003 issue of the *Journal of Accounting and Public Policy* (see below list of related publications).

Publications and Papers

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, "Sharing Information on Computer Systems Security: An Economic Analysis," *Journal of Accounting and Public Policy*, Vol 22, No. 6, 2003, pp. 561-485.

4.3 Business Case Development (L. Gordon and M. Loeb)

On August 27, 2003, we met with Dr. William Semancik to receive further

clarification of this task. Based on a 2.5-hour meeting, the specific nature of task was clarified and given a much sharper focus. Specifically, the task is now defined to focus on examining the differing paradigms for organizing and managing research activities. Since the above noted meeting, we have spent a significant amount of time on this task. In this regard, we have decided to use a "transactions cost economics framework" as the basis for our analysis. We are now ready to present our preliminary findings related to this task to Dr. William Semancik. We have contacted him in this regard, and are waiting for him to indicate a time that would conveniently fit his schedule.

Finally, some of our activities during the quarter relate to all of the above tasks and indicate that impact of our sponsored research. Last quarter our work was highlighted in an interview broadcast on Maryland Public Television (MPT). In this regard, Lawrence A. Gordon was the guest expert interviewee on "Business Connections," which is hosted by Mr. Jeff Salkin, and this quarter Dr. Gordon was asked to schedule a follow-up interview. Second, during this quarter, Larry Gordon presented a summary of our work related to the above tasks at the I-4 meeting. Third, Drs. Gordon and Loeb have continued contacts with officials of the Homeland Security Department. Fourth, Dr. Gordon reviewed an NSF grant related to the economics of information security and reviewed a submission to *ACM Transactions on Information and System Security*. Fifth, both Drs. Gordon and Loeb were reviewed submissions to *Communications of the ACM* that build on the above referenced research. Sixth, Dr. Loeb reviewed a submission to *Information Systems Research* that builds on our work in the information security area and has been asked to review another related paper for *Management Science*. Seventh, Dr. Loeb is serving as member of the doctoral dissertation committee for a computer science student at Harvard working on issues dealing with the economics of information security.

4.4 Research in the Impact of Pricing Strategies (J. Bailey and S. Raghavan)

This project represents work under 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.

Our empirical research in the area of Internet pricing is progressing nicely. Along with our doctoral student, Toby Porterfield, we have had a chance to analyze data from October 2003 and compare it with biannual data starting with March 1998. As expected, we are seeing more consolidation in this industry over time. Our hypothesis is that this is occurring because of the pressures of price competition, technology development, and security. We have had a chance to analyze the different dimensions of competition among ISPs and have found that less than 10% advertise their offerings in terms of their security investments. This figure is far smaller than the percentage of ISPs that post prices and offer different advanced technologies that are not security-related. Furthermore, we have linked security investment to an overall increase in the connectivity speed of the ISP to the Internet. We have decided to separate out the task of analyzing such a rich data set somewhat along the lines of the tasks supported through

LTS. Our first paper is to examine the role of security from a resource based view of the firm. Building upon our work on ex-post charging, we will attempt to find linkages between security and pricing in our data.

4.5 The Business Case for Wireless Systems (J. Bailey and S. Raghavan)

This project represents work under Task 1 (Business Case Development)

Our database of competition in the ISP industry will also serve as a foundation for research into wireless technologies. Many ISPs are currently deploying wireless systems and we are interested in the type of diffusion seen among these firms. We hope to test hypotheses such as our proposition that firms with wireless technologies likely to also be investing in security technology. Although we have not yet started analyzing the data yet, we will likely start this longitudinal study soon. A likely approach to this research will be a nested logit model so we can understand the decision to deploy wireless systems conditional on the firm's use of security technology.

Along with doctoral student Robert Day, we continue our study of spectrum auctions and the appropriateness of using CAMBO (our proposed combinatorial auction framework) for spectrum auctions.

We have started a study of satellite technologies, specifically geostationary satellites, for reliable/secure communication. 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. The results of the study should be a better understanding of the costs, and methods to minimize the costs of satellite communication networks. We also propose to develop models to deal with the inherent uncertainty in future demand in business cases/planning.

Publications

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

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

4.6 The Business Case for Optically Transparent High-speed Networks (J. Bailey and S. Raghavan)

This project represents work under Task 5 (The Business Case for Optically Transparent High-Speed Networks).

Our database on ISP competition will also help us understand the role of deploying high-speed networks in the market. Although there appears to be little diffusion of these technologies in the ISP market, the firms that do appear to be leading the efforts are the same ones that invest in security. In our paper on pricing we hope to separate out the effect of firms deploying high speed networks with their investment in security. However, it is very likely that these two variables moderate the overall success of an ISP.

Along with doctoral student Daliborka Stanojevic, we are continuing our study of Integrated Logical Topology Design (ILTD) for MPLS/WDM Networks. We have made significant progress on a comprehensive literature survey, identifying areas of investigation where mathematical programming methods can be effectively used of for the design of MPLS/WDM Networks. We have implemented a linear programming based column generation procedure to solve the ILTD problem. We find that we are able to solve problems with up to 20 nodes within an hour of CPU time on a fast Pentium 4 PC. This is a tenfold increase in running time over the traditional arc-based formulation for the ILTD problem. The solutions obtained are all within 1% of optimality. We are now working on trying to solve the ILTD problem to optimality using column generation for integer programming.

Publications

"Heuristic Search for the Generalized Minimum Spanning Tree Problem," B. Golden, S. Raghavan, and D. Stanojevic, Accepted, INFORMS Journal on Computing.

"Long-Distance Access Network Design," R. Berger and S. Raghavan, to appear Management Science, March 2004.

5. Optical Networking (Gary Carter and Joel Morris)

5.1 High Speed Experiments (G. Carter)

This project represents work under Task 1 (Long Haul Experimentation Testbed) and Task 2 (Long Haul Transmission Experiments).

We have completed and analyzed the first round of experiments on ATDNET to BossNet and back. The transmitter and receiver were at LTS. The experiments were looped back at BossNet's Wilmington, Delaware and New York nodes. The timing jitter was accurately measured as a function of transmitter pre-compensation and receiver post-compensation dispersion. These two dispersion elements compensate for the dispersion primarily in the link on ATDNET. Figure 1 below shows the experimental data for the Wilmington loop back along with a calculation of the timing jitter using a sophisticated theoretical model which includes the power and dispersion distribution for the link.

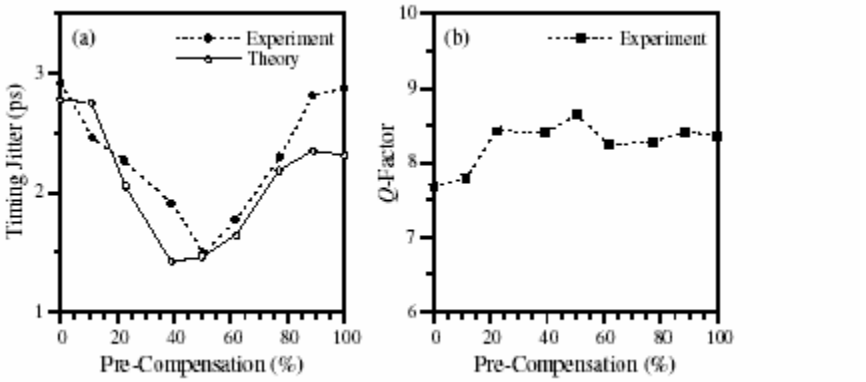

Figure 1. The measured and theoretical timing jitter for the LTS-Wilmington link as a function of pre-compensation (a); and the measured Q factor as a function of pre-compenstation.

The data clearly shows the minimum in timing jitter at approximately 50% pre-compensation. This demonstrates that uncompensated links present a challenge in all-optical networks. It will be insufficient to completely compensate the dispersion all at one point. As can be seen from Figure 1, it will be necessary to apportion the compensation between the receiver and the transmitter. This distribution will vary depending on which links the data traverses. Thus the network management system will ideally have sufficient knowledge to "dial-in" the appropriate pre- and post-compensation. Note that these results were obtained over the relative short distance of approximately 400 km. This work has been accepted for publication at the 2004 Conference on Lasers and Electro-optics.

5.2 *Statistical Signal Characterization (J. Morris)*

This project represents work under Task 3 (Statistical Signal Characterization).

Recent progress on Task 3 of the High-Speed Optical Networking project comprise several areas:

1. An exhaustive analysis of the decoding ability of the bit-flipping algorithm (BFA) for 3-error patterns for the RCD codes has been completed for η = 3, 5, 7, …, 23, where η is the size of the square array. Each and every 3-error pattern was generated, classified, and decoded via the BFA. The results were then tabulated. The knowledge of the number of 3-error patterns for the RCD codes that decoded successfully to the all-zeros codeword was used to strengthen the upper bound on the word error rate (WER) of the BFA. The number of 3-error patterns that resulted in decoder failure or decoding to the wrong codeword were also used to lower bound the WER of the BFA. The adjoining figure shows the tighter bounds based on decoding of 3-error patterns for the RCD code with $\eta = 23$ (code length 529).

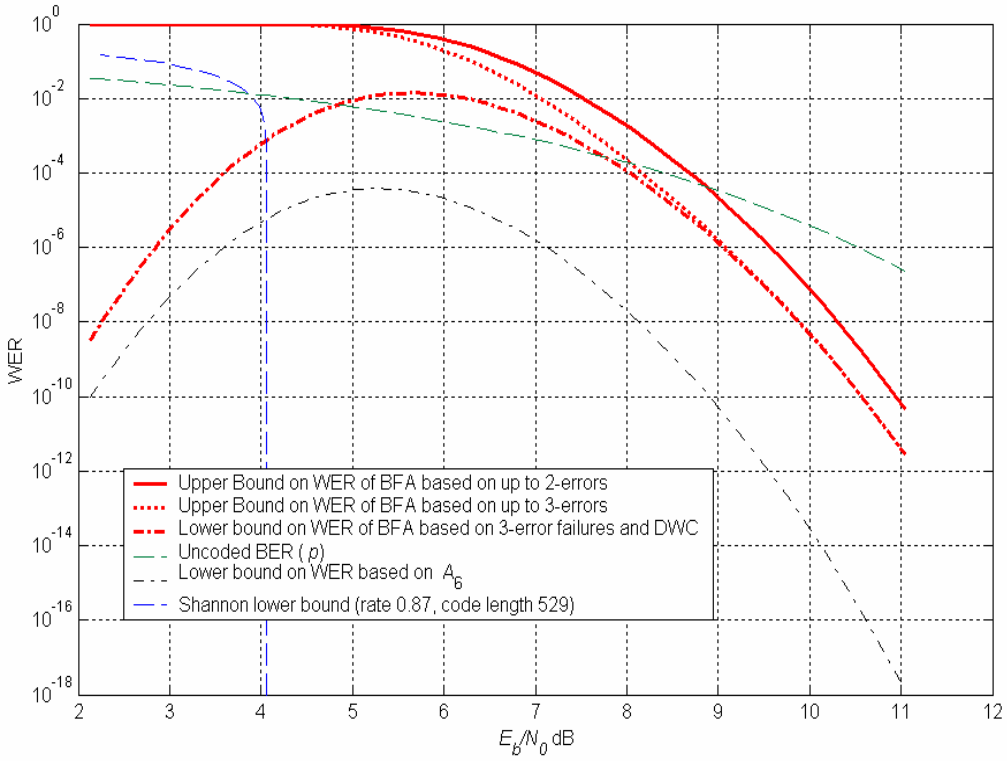

Figure 1: Bounds on the WER performance of the BFA for RCD code with $\eta = 23$.

To further tighten the bounds on the WER of the BFA in the region of $E_b/N_0 < 9$ dB, investigation of the behavior of 4-error patterns and higher weight error patterns is necessary. Note that for $E_b/N_0 > 9$ dB the upper bound and lower bound on the WER based on the 3-error patterns merge suggesting a very tight bound.

Certain combinatorial expressions have been conjectured to lower bound the number of 3-error patterns that decode successfully to the all-zeros codeword, and to lower bound the number of 3-error patterns that result in decoder failure or decoding to the wrong codeword. It is our aim to prove the accuracy of these combinatorial expressions. These combinatorial expressions will enable us to obtain tighter bounds on the BFA performance for 3-error patterns for RCD codes with $\eta > 23$ without needing to exhaustively generate, classify, and decode every 3-error pattern, thus saving valuable computational time and resources.

Preliminary investigation of 4-error patterns indicates that there are 16 unique classes possible. An exhaustive analysis of the decoding ability of the BFA for 4-error patterns for the RCD codes has been completed for $\eta = 3, 5, 7, \ldots, 15$.

2. Further progress has been made with Prof. Curtis Menyuk's group on the dual-adaptive importance sampling (DAIS) technique to evaluate FEC code probability of error down to extremely low probabilities. DAIS simulations were performed on LDPC codes with length 20 and 96 under sum-product decoding. The DAIS simulation results were found to be in excellent agreement with the Union Bound for these codes, and in the latter case, with results based on standard Monte Carlo simulations down to a BER of 10^{-8}. Further developmental work is continuing and DAIS simulations are now being carried out for a much longer code – the RCD code with $\eta = 37$ (code length = 1369).

3. An exact combinatorial expression for the number of codewords at the minimum distance of 6 (A_6) has been derived for the class of RCD codes with parameter η. The expression is $A_6(\eta) = \frac{\eta^2(\eta-1)(\eta-2)}{6} = \eta\binom{\eta}{3}$. A combinatorial expression has also been conjectured for the number of codewords of weight 8 (A_8). The expression is $A_8(\eta) = \frac{\eta^2(\eta-1)(\eta-3)(4\eta-11)}{8}$. A proof for A_8 for all η has not yet been established. Knowledge of A_6 and A_8 enables us to compute the Union Bound for an RCD code with any η. The Union Bound is an upper bound on the maximum likelihood (ML) decoder performance of a code under soft-decision decoding on a binary phase-shift-keying (BPSK) additive white Gaussian noise (AWGN) channel. The Union Bound is known to be tight at high values of E_b/N_0. The Union Bound, thus, indirectly serves as a tight lower bound for code performance for practical decoders, such as the sum-product decoder, that are known to closely approximate the optimal ML decoder. Figure 3 shows the Union Bound on BER based on A_6 for RCD codes with $\eta \in \{7, 23, 37, 71\}$.

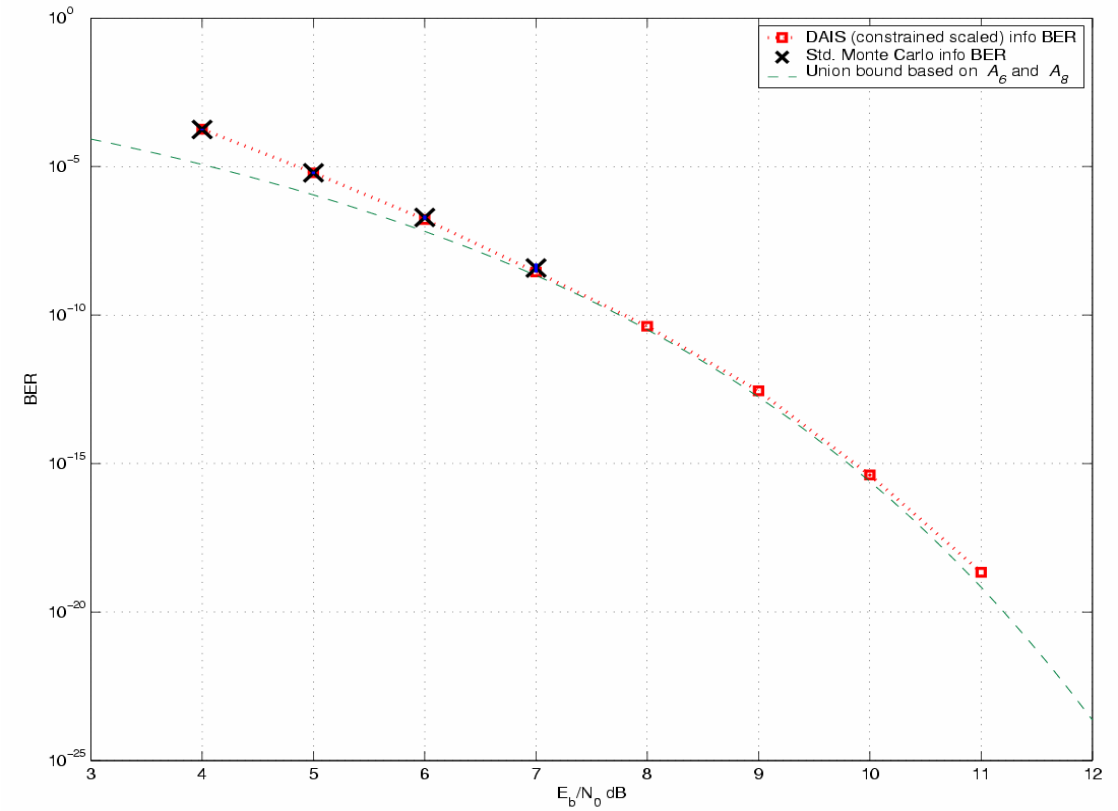

Figure 2: Performance curves for the BER of an $n = 96$, $k = 50$, LDPC code under sum-product decoding.

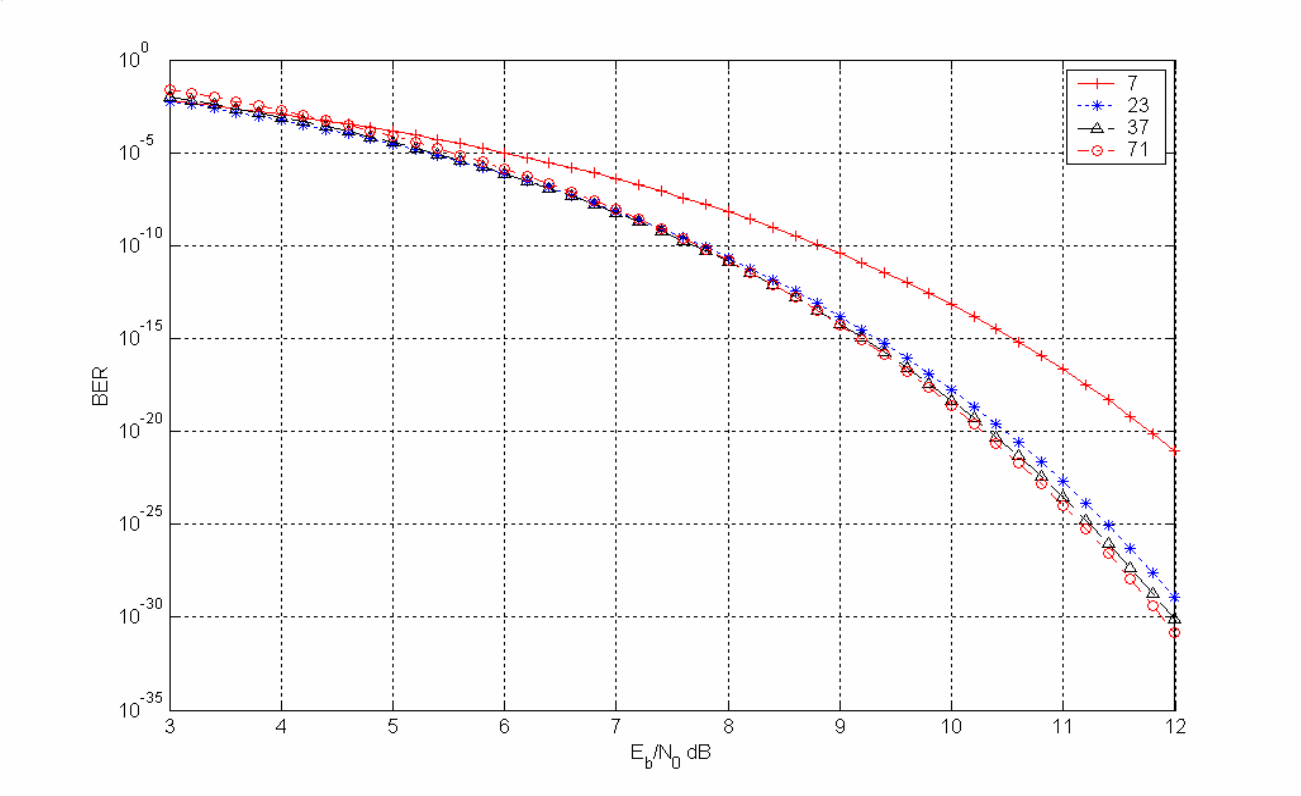

Figure 3: Union Bound on BER based on A_6 for RCD codes with $\eta \in \{7,23,37,71\}$.

6. Technical Exchanges and Peer to Peer Networking (J. JaJa and M. Marsh)

6.1 Seminar Series (J. JaJa)

This project represents work under Task 1 (Technical Exchanges for Telecommunications Sciences)

During the Fall of 2003, we ran a seminar series that focused on peer to peer networking, which included some of the top researchers in this area. Below is a list of seminars, which includes for each seminar the name of the speaker and the abstract of his presentation.

October 15, 2003 – Resilient Multicast Using Overlays
***Speakers*: S. Bhattacharjee and A. Srinivasan**
University of Maryland

Abstract
We first present a protocol for anonymous communication over the Internet. Our protocol, called P5 (Peer-to-Peer Personal Privacy Protocol) provides sender-, receiver-, and sender-receiver anonymity. P5 is designed to be implemented over the current Internet protocols, and does not require any special infrastructure support. We present an outline of the P5 protocol, and a set of performance results based on detailed simulations.

We next present a low-overhead media streaming system, called SRMS (Scalable Resilient Media Streaming) that scales to very large groups. SRMS leverages a probabilistic loss recovery technique to provide high data delivery guarantees even under large network losses and overlay node failures. We describe the randomized protcols that are used to provide resilience in SRMS, and present simulation and implementation results showing its efficacy in practice.

November 5, 2003 - Randomized Algorithms for Network Security and Peer-to-Peer Systems
Speaker: Micah Adler
University of Massachusetts at Amherst

Abstract
This presentation will consist of two independent shorter talks, with a question and answer period in between. The first talk will consider probabilistic packet marking (PPM) for IP Traceback. PPM is a technique for tracing a sequence of network packets back to an anonymous source. An important consideration for such schemes is b, the number of packet header bits that need to be allocated to the marking protocol. In this talk, we introduce a new PPM scheme applicable when a sequence of packets is sent along the same path. This new technique allows the packets to be traced back to their source using only a single bit in the packet header. With this scheme, the number of packets required to reconstruct the path is $O(2^{2n})$, but we also show that it is not

possible to do better when $b=1$. We also study the tradeoff between b and the number of packets required. We provide a protocol and a lower bound that together demonstrate that for the optimal protocol, the number of packets increases exponentially with n, but decreases doubly exponentially with b.

In the second talk (joint work with Eran Halperin, Richard Karp and Vijay Vazirani), we analyze a load balancing question for distributed hash tables (DHTs), a fundamental tool in peer-to-peer networks. In our DHT, the number of queries required to find a key in the table is $O(\log n)$, the number of pointers each processor maintains is $O(\log n)$, and during a sequence of n processor arrivals, the ratio between the maximum load of a processor and the minimum load of a processor is always $O(1)$ (with high probability). To the best of our knowledge, this is the first analysis of a DHT that achieves this performace. This analysis reduces to a simple stochastic process executed on a graph; this process can be viewed as a structured version of the well known coupon collector's process.

November 10, 2003 – Structured Peer-To-Peer Overlay Networks: A New Foundation For Distributed Applications?
Speaker: Dr. Peter Druschel
Rice University

Abstract
Peer-to-peer (p2p), initially conceived for the purpose of sharing music in the Internet, is emerging as a much more general paradigm for the construction of resilient, large-scale, distributed services and applications. We define p2p systems broadly as self-organizing, decentralized distributed systems that consist of potentially untrusted, unreliable nodes with symmetric roles. The scalability and resilience of p2p systems lends itself to a growing domain of applications beyond file sharing. At the same time, the scale, decentralization, diversity and potentially open membership in these systems pose difficult problems, particularly in resource management and security.

Recent work on structured p2p overlay networks like CAN, Chord, Pastry and Tapestry has made significant strides towards providing a general substrate that simplifies the construction of robust, large-scale distributed applications. These overlays effectively shield application designers from the complexities of organizing and maintaining a secure overlay network, tolerating node failures, balancing load, and locating application objects.

In this talk, I'll present an overview of the state-of-the-art in structured p2p overlays that provide self-organization, fault-tolerance, efficient object location, and proximity-aware overlay construction. I will also sketch the design of several applications, including cooperative network storage, scalable endsystem multicast, and content distribution. I'll conclude with an outlook on key research problems and future directions.

December 10 - New Tools in Applied Cryptography
Dr. Dan Boneh

Stanford University

Abstract

Over the past three years we have seen a number of exciting new cryptographic constructions based on bilinear maps. In this talk, we will survey some of these new constructions and their applications. For example, bilinear maps give rise to a new digital signature scheme with remarkable properties: it can be used to reduce communication in secure routing protocols and shrink certificate chains. Bilinear maps have also been used to construct the first practical public key encryption scheme where public keys can be arbitrary strings (i.e., an identity-based encryption scheme). In this talk we will survey some practical applications of bilinear maps and describe several open problems in this area. The talk will be self contained

6.2 Peer to Peer Networking Research (M. Marsh)

This project involves work that was approved under an expansion of Task1.

Trust Inference

Determining whom to trust is a challenging problem. For people, trust is built through social interactions: an individual who has acted in good faith previously is more deserving of trust than both someone who has acted in bad faith and a stranger. While trust is not in general transitive, we can infer trustworthiness in that someone considered trustworthy can vouch for the trustworthiness of another.

We explore the applicability of these notions of trust to networks of peers. Good faith actions might include correct computations or fair exchanges of resources, while bad faith actions would be the converses. By quantifying the effects of actions, we can establish a trust metric that allows a peer to make judgments about other peers.
These judgments can then be exchanged with others to propagate at least a minimal value for a peer's trustworthiness. We also consider accusations of bad actions, and the trustworthiness of the accusers.
Given such an infrastructure, we can construct distributed applications that require trust.
This is joint work with Samrat Bhattacharjee and Jonathan Katz.

Lookup in Unstructured Networks

In a content-storing peer-to-peer system there is a general problem of finding items. We limit ourselves to the case where the identity of an item is known, so that lookup and retrieval are the only relevant issues. While much work has been done on efficiently searching on peer-to-peer networks where the topology is constrained (ie, distributed

hash tables), much less has been done to efficiently search on unstructured topologies. We have developed a technique for searching on an unstructured topology which is efficient and scales to very large peer-to-peer networks. As with distributed hash tables, we virtualize both peers and items with 160-bit identifiers (such as the output of a collision-resistant hash function) and define a distance between identifiers. Where distributed hash tables place items (or pointers to items) at the closest nodes, which requires long-range knowledge of the network, we instead place items at the locally closest nodes, which we refer to as local minima. Since this requires only short-range knowledge, it scales to large unstructured networks.
Both placement and searching employ the same basic algorithm, which we call local minima searching (LMS). Beginning placement or searching with a sufficiently long random walk serves to randomize at which local minimum a placement or search will terminate. This is joint work with Samrat Bhattacharjee and Aravind Srinivasan and has been submitted to ACM's SIGCOMM '04 conference.

As a network grows, additional copies (replicas) of items must be placed at increasingly large distances from their owners. We introduce an adaptive protocol that determines the number of replicas needed based on search results that occur during normal system operation. We achieve long-distance propagation by resuming the random-walk phase of placement (with a longer random walk) whenever placement terminates at a peer that already holds a replica.

Undiscoverable Secure Communications

In joint work with LTS, we consider a system involving mobile nodes in a hostile environment. Communications must be maintained among nodes and between the nodes and a trusted center. Security is paramount in this scenario, so strong authentication and secrecy are needed. To the environment, which might be controlled by a powerful adversary, it must appear at best that no communications are occurring, and at worst that communications originate from indeterminate sources. Traditional notions of anonymous communications are insufficient, since hiding in a crowd is infeasible when all members of the crowd must remain hidden.

Trust Exchange

Establishing trust is difficult, so once established a trustworthy principal has great incentive to assure other principals that it will remain trustworthy. If any individual processor might be compromised, we can maintain trustworthiness by distributing trust among a set of processors in such a way that (under some failure assumption) enough correct processors participate to guarantee that an action is performed correctly. Moreover, by employing cryptographic techniques we can ensure that no compromised processors will ever learn secrets for which they are not authorized.

We consider specifically a service, which we call CODEX, that stores secrets for clients and allows them to specify authorized recipients of the secrets. CODEX is short for the Cornell Data Exchange, and was developed at Cornell University as joint work with Fred Schneider.

While the bulk of the work was performed at Cornell, substantial efficiency gains were subsequently made at UMCP based on discussions with Jonathan Katz. In addition, the code base for CODEX was used (and further enhanced) to construct an implementation of the LMS protocol. A journal paper describing CODEX is currently in preparation.