ITR: Detection and Localization of Anomalous Network Behavior

Alfred Hero(PI), Rich Baraniuk(co-PI), Mark Coates(collaborator),
Stéphane Lafortune(co-PI), Mingyan Liu(co-PI), George Michailidis(coPI),
Brian Noble(co-PI), Robert Nowak(co-PI), Sandeep Pradhan(co-PI), Atul Prakash(co-PI),
Rolf Reidi(co-PI), Demosthenis Teneketzis(co-PI), Dan Wallach(co-PI)

March 14, 2002

1 Summary

See next page

Today's private and public communications networks are critical systems of data terminals, routers, and sensors which provide the backbone of our information society. The focus of this grant proposal is to develop an infrastructure for distributed detection and localization of anomalous behavior in large heterogeneous data networks. The proposed effort has four components: 1) development of a comprehensive program in network security education; 2) dissemination of a collaborative open-source software suite for data collection, visualization, and analysis; 3) development and implementation of distributed algorithms for pattern recognition, anomaly detection, and measured response; 3) demonstration of these algorithms for distributed intrusion detection and denial of service (DDoS) attacks.

The sheer size and complexity of the Internet makes the anomaly detection problem extremely challenging. We propose to develop and implement efficient network probing strategies and a powerful new class of distributed algorithms which monitor multi-dimensional information flows of packets sizes, packet rates, source-destination addresses, and other packet attributes. This will lead to improved off-line and on-line algorithms for identifying anamolous/malicious behavior, for implementing dynamic routing and flow control for congestion, for network provisioning, and for remote monitoring and service verification. Our research approach is a potent combination of emerging techniques in signal processing, networking, discrete events, and stochastic optimization including: novel probing methods to perform network measurement and tomography from multiple sites; distributed data compression of collected statistics; and decentralized pattern recognition and change detection in traffic flows, packet types, and network transport characteristics.

Crucial to detecting anomalous changes in aggregate behavior of networks is our ability to determine traffic flow statistics throughout the network and to characterize what constitutes a significant change in flow patterns. This would normally require measurements of internal link traffic and a reliable baseline against which to test for anomalous traffic variations. However, the largely unregulated structure of networks makes it impractical for every router and terminal to cooperate in collecting and forwarding local traffic statistics. Furthermore, even if massive amounts of link data could be collected, data-analysis presents difficult problems of large scale computation and statistical validation. Our team proposes an innovative and comprehensive approach to this problem in which we will use network tomography to identify and predict internal traffic behavior patterns based on measurements from a few edge nodes of the network. While we will implement a variety of pattern recognition and detection algorithms, particular focus will be on a flexible multivariate time-series packet-level model that will be coupled to a higher-level dynamical discrete event system (DES) model of network dynamics. DES's have been recently applied to network fault detection and isolation, but this is the first time that these models are being proposed together with time-series models as part of an integrated and comprehensive framework for detection of anomalous traffic behavior in large-scale networks.

This project will involve precollege, college, and continuing education. A new cross-disciplinary undergraduate and graduate curriculum in global network security will be introduced. Students in these courses will participate in data collection, software development, and data analysis as part of instructional lab projects. We will also create a summer internship program for qualified high school and middle school students. These students will participate in various educational and recreational signal processing and networking activities which we will organize. An educational innovation of our project is the development of an isolated networked environment at Rice and Michigan for emulation of attacks. A yearly summer competition will pit Rice students against Michigan students in a contest for most effective attacks on and defenses of their respective subnetworks. This laboratory will serve dual purposes: 1) an exciting environment for learning about network security; and 2) a testbed for student generated attack scenarios to be used for our research. Our industry collaborators will provide guidance and inputs to this project including: help in emulating realistic attack scenarios, response strategies, and deployment of distributed data collection software.

1. Proposal Overview

We are a multi-disciplinary team of researchers in the fields of adaptive signal processing, network traffic measurement and modeling, network topology characterization, discrete event systems (DES), multivariate statistics, tomography, distributed computing, network performance analysis and content delivery, network security, and stochastic control. We propose to develop and disseminate a methodology for distributed data collection and rapid detection and localization of the onset of spatio-temporal changes in global network traffic. If succesfull, the outcome of this project could allow prediction and mitigation of distributed attacks and other disruptions before they fully evolve and cause damage.

Remote detection and localization of anomalous internal traffic patterns in heterogeneous networks require the synthesis of accurate algorithms for detection and recognition of deviant patterns of packet flows in the network, inference of characteristics of internal links that transport the traffic, and distributed measurements that can be acquired and integrated with minimal overhead on computation and communication. Quickly detecting and localizing such changes from measurements of a few cooperating nodes is an extremely challenging problem that requires a large scale effort and new approaches. Our approach is a broad, systematic and integrated strategy of distributed algorithm development, collaborative data collection, and education that has the following features:

1. Integration of a discrete event system model and a spatio-temporal traffic model for classification of detectable sequences of changes in network transport characteristics.
2. A flexible combination of passive measurements and active probing network tomography methods for two stage detection of changes in spatio-temporal packet distribution patterns and link statistics.
3. Application of centralized and decentralized detection algorithms to track changes in network connectivity, changes in loss and delay distributions, changes in traffic correlation patterns, and other potential anomalies.
4. Experimental validation with real multivariate data traces from the Internet and from large private networks. Data will be collected and analyzed in collaboration with Internet2, Los Alamos National Labs, NASA JSC, Sprint, and Texas Instruments, (Letters of support are attached).
5. Creation of an open collaborative infrastructure for data collection and distributed pattern recognition for monitoring the global network. This will involve free and wide dissemination of open-source software tools for data collection, computation, and transmission to Rice and UM. SImilarly to the SETIhome program the software would take advantage of a participant's free cpu cycles to perform data collection and computations.
6. Development of an instructional laboratory on network security using an isolated IP network to teach students about security through a multiterminal computer game. In this game teams of students, sponsored by our industrial partners, will match wits against each other on developing and mitigating attack strategies. Data traces will be collected and used as test scenarios for our research.

This project differs from previous "Internet Mapping" projects that focus on mapping a large portion of the Internet core or on mapping performance metrics throughout the network. Here we depart from this line of thinking. Instead of globally mapping the Internet, we focus on rapidly detecting and localizing anomalous performance or traffic behavior. This is a move from the conventional estimation/mapping approach to a detection-theoretic approach utilising distributed pattern recognition algorithms.

The research project will likely result in major advances: 1) a fuller understanding of the limitations of network inference methods for detecting and localizing potentially debilitating attacks and link

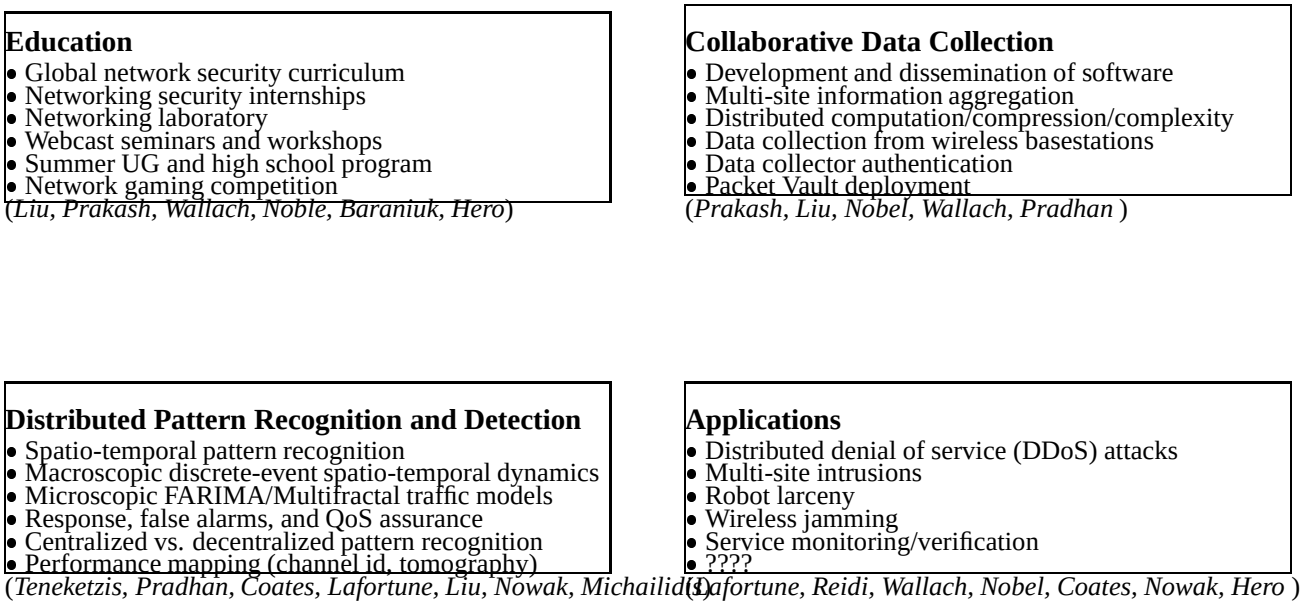

Figure 1: Commutative diagram of research and education activities.

failures; 2) an integrated and flexible multiple time series analysis relating traffic measurements at a few monitoring sites to internal traffic and link behavior; 3) scalable decentralized algorithms for detecting emerging attack patterns from edge-node measurements; 4) a software tool for topology generation, visualization, and simulation which incorporates an accurate spatio-temporal model of the network; 5) instruction of undergraduates on computer security through a fun computer game which will serve to both generate data for validating our models and generating new patterns of attack and mitigation. An added benefit will be multi-disciplinary training of graduate and undergraduate students in Signal Processing, Networking, Statistics, Discrete Event Systems, Optimization, and Software.

A commutative diagram showing how the proposed effort is compartimentalized is given in Fig. 1. Principal associated co-PI's are listed with each of the activities,

2. Prior NSF Support

Below we include one paragraph each from each PI who has been supported on pertinent NSF project. We need everybody's projects here!

1. "Information Visualization through Graph Drawing: Modeling, Analysis and Optimization Issues," NSF/IIS-9988095; 01/01/01-12/31/03, PI, George Michailidis,

   Summary: This ongoing project focuses on (1) developing a flexible modeling framework based on graph theoretical concepts that allows the efficient representation of complex data structures, (2) formulating information visualization as an optimization problem and (3) developing efficient, robust, and simple algorithms for solving the problem.

2. "Information theoretic analysis of tomographic systems," NSF/BCS-9024370(1993-1995), PI A.O. Hero

Summary: Tomography allows spatio-temporal characteristics of a source or medium to be reconstructed from a few edge measurements or projections. In this NSF grant, which ended in 1995, the fundamental limitations on performance of tomographic systems were characterized in terms of edge sensor placement, properties of the medium, and statistical variability of the measurements. This resulted in more pertinent criteria for design of tomographic data collection systems and in new high performance algorithms for reconstruction [21, 23, 19, 15, 38, 36, 35, 37, 20, 17]. It also resulted in establishment of rates of convergence of iterative reconstruction algorithms in terms of information and entropy measures [18, 16, 12], development of improved iterative reconstruction algorithms based on information decoupling [13, 14, 11], and development of a class of image reconstruction algorithms based on recovering a confidence region in object space [39, 22, 40]. The paper [21] describing the Uniform Cramer-Rao bound design criterion won a Best Paper Award from the IEEE Signal Processing Society in 1998.

3. Distributed Pattern Recognition and Detection (Leader: Teneketzis)

- Spatio-temporal pattern recognition (Michailidis, Teneketzis)
- Macroscopic discrete-event spatio-temporal dynamics (Lafortune, Teneketzis)
- Microscopic FARIMA/Multifractal traffic models (Reidi, Michailidis)
- Responses, false alarms, and QoS assurance (Liu)
- Centralized vs. decentralized pattern recognition (Teneketzis, Hero, Pradhan)
- Performance mapping (channel id, tomography) (Nowak, Coates)

(*Teneketzis, Pradhan, Coates, Lafortune, Liu, Nowak, Michailidis*)

The research approach in this proposal represents a dramatic departure from existing activities in traffic analysis and network monitoring. We will combine novel and flexible multiple stream traffic data collection, adaptable information aggregation strategies, decentralized diagnostic algorithms to detect and localize abnormal network behavior, and control actions to mitigate undesired behavior.

3.a Pattern Identification/Recognition

We will investigate a dynamical system framework for capturing dependencies between multiple spatio-temporal samples of traffic and packet level information (http/ftp requests, ping, netflow) taken at different points in time and space (localized on topology). The systems models are sufficiently general to be adapted to different traffic types, e.g. those occurring during an attack at a victim's host site and at an upstream router site. This general framework extends extant pattern recognition and event correlation techniques in two ways: residuals from our proposed hybrid dynamical system model permit detection of deviations from a baseline of traffic flow and associated sequences of events; 2) control actions can be evaluated in the context of maintaining stability and minimal levels of overall QoS.

Figure 2 encapsulates the dynamical system model for the case of a single data collection site. One way to capture traffic flow dynamics is to identify a multivariate spatio-temoral time series traffic model combined with a discrete event system model (module labeled DES at left of diagram). The time series model captures microscopic (fast) behavior of the network while the DES model captures its macroscopic (slower) behavior. This powerful combination of models allows us to decouple packet level behavior from network transport level behavior to quickly zero-in on a wide range of spatio-temporal anomalies in the network. Traffic measurements are collected and integrated from a few monitoring sites (nodes of the network) using a combination of ambient traffic measurement and network tomography with active probes (module in middle of diagram). Based on these measurements detectable changes in macroscopic behavior of the network will be performed by estimating the states of the DES model (module at right of diagram). This will be accomplished in two stages: 1) estimation of microscopic traffic parameters using the multivariate spatio-temporal traffic models; and 2) use of these estimates to

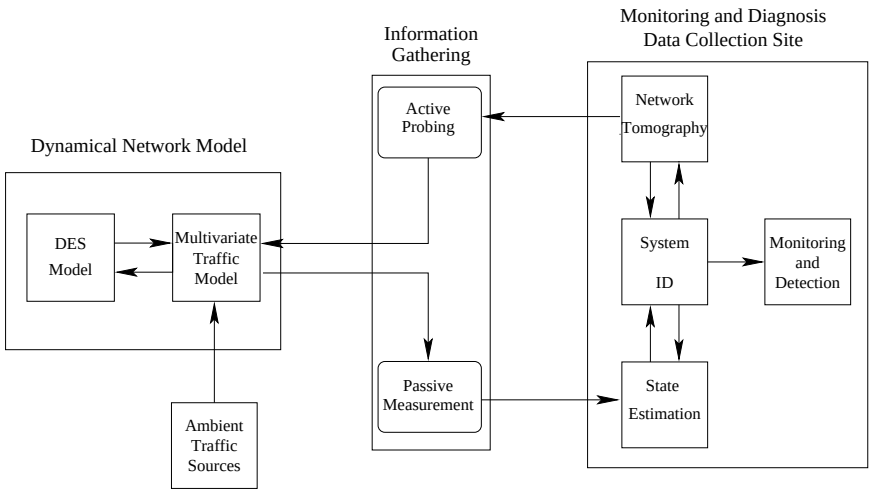

Figure 2: *Block diagram of our approach for a single collection site.*

identify states of the DES model. The information gathering configuration may be adapted as evidence of an anomaly emerges by more intensive probing or by switching to alternative measurement sites in order to focus on suspected sections of the network.

The above describes a model-based approach to detecting anomalous behavior using a single collection site. In some cases the model may be overly complex to implement in real time and we will also consider non-parametric learning-based approaches to detecting deviations from a baseline using sequential non-parametrics [31]. Since networks are inherently informationally decentralized systems we will consider multiple collection sites and we will investigate issues of communication and coordination among sites for real-time decentralized detection and localization of abnormal behavior. *However, due to space limitations, in the remainder of the white paper we restrict discussion to our basic strategy for centralized processing at a single data collection site. Below we discuss in more detail the modeling and data processing methodologies and their applications. REPLACE*.

3.b Change Detection Strategies

Sequential non-parametric approaches: Flesh this out....

Multivariate Time-Series Packet-level Models: Given a particular configuration of the network, the flow of packets through the network generates traffic which we characterize by a network of coupled multivariate time series models. These spatio-temporal traffic models capture the fast dynamics of local traffic flows as related to flows of traffic elsewhere in the network and exogenous inputs due to probing and ambient traffic sources. The parameters of the traffic model change according to the slower state transitions of the network which are determined by the DES model which (see description below). While we will investigate several classes of multivariate time series models including multiscale and multifractal traffic models [29], our focus will be on fractional autoregressive integrated moving average (FARIMA) models. Like multiscale models FARIMA models been shown to reliably capture both the long term dependency as well as short range dependency of single stream IP traffic flows [3, 26]. However several factors favor the FARIMA model for this application: 1) our research can leverage on the large body of research on identification of ARIMA models developed over the past 50 years; 2) decentralized and multivariate extensions of recursive FARIMA model identification algorithms appears to be more straightforward; 3) as FARIMA models are defined explicitly as temporal state recursions, they would appear naturally suited to detecting transient disruptions; 4) causal and time recursive state and parameter estimation algorithms are more easily developed.

To refine the multivariate traffic model and improve parameter identification performance we will explicitly account for any known physical traffic or transport characteristics. For example, *a priori* terminal-type information can help determine the appropriate traffic source models to use in the system identification algorithm. As another example, some of our most critical networks incorporate wireless components, e.g. the Airphone network or the police-band wireless data network. The effects of fading in the channel, co-channel interference, and transmitter modulation type provide characteristic signatures for these links which impact the structure of the packet loss, delay and link bandwidth distributions.

Discrete-Event Models of Network Behavior: We propose to use logical and stochastic DES models that will work in tandem with the multivariate spatio-temporal traffic models described above for the detection of component failures, attacks, and other anomalies in the network. The DES models will capture the network traffic behavior at a higher level of abstraction and on a different time scale than the spatio-temporal models. The events that will drive the DES models will comprise "observable" events that will be obtained by aggregating, quantizing, and filtering key variables from the spatio-temporal models as well as "unobservable" events that will capture special changes of the state of the overall network, possibly due to anomalies or attacks, that are not directly measured nor captured by the spatio-temporal models. The DES models, together with the sequences of observable events that will drive them online, will then be used to infer about the occurrence of the unobservable events and thus detect, if any, various anomalies in the behavior of the network.

The theoretical foundations for the task of discrete-event model-based inferencing lie in the failure diagnosis methodology for logical DES developed in our prior investigations [30, 10]. In fact, these works have been a major source of inspiration for related approaches for network fault detection and isolation in wireless LANs [9] and large-scale telecommunication networks [1, 27]. The methodology in [30, 10] will have to be significantly enhanced in order to address the objectives of this proposal. These enhancements include: (i) incorporation of nondeterministic and stochastic features in order to appropriately couple, in a hierarchical manner, the DES models with the spatio-temporal models; (ii) development of modular and decentralized algorithmic implementations in order to address the scalability requirement; and (iii) incorporation of event-driven models of DDoS and spoofing attacks and development of distributed and asynchronous algorithms for the detection of such attacks.

3.c. Data Processing Strategy

Anomalous Event Detection: The end goal is to perform event detection in real time by identifying anomalous sequences of state transitions in the DES network model and anomalous parameters of the multivariate time series traffic model. The detection of anomalous events from edge-measurements will consist of two steps: 1) estimation of a baseline of normal activity; and 2) detection and localization in time and space of deviations from the baseline. First we associate different labels, describing the "status" of an underlying subnetwork (e.g. normal, congested, under attack, etc.), with different sets of estimated DES states and traffic model parameters. Then system identification algorithms will be developed for classification of baseline states of the DES. Finally, sequential change point detection algorithms will be specified to detect deviations from the baseline as quickly as possible and with a prescribed false alarm rate. We will investigate both centralized methods, where all monitoring sites relay all information to a central collection site for processing, and decentralized methods, where clusters of data collection sites only exchange limited amounts of information. Previously developed centralized [33, 2, 25] and decentralized [34, 28, 8] sequential change point detection algorithms are inapplicable to our network model. We will investigate fully Bayesian decentralized detection methods using the sequential Monte Carlo Markov Chain framework such as that used in [6]. The focus will be on detection of the occurrence of a change point in traffic flows with penalties on the false signal rate and the temporal spatial localization error.

Network Inference and Tomography: Measurements from only a few distributed nodes will be used to do change point detection, localization, and model identificaton as discussed above. These measurements will be composed of a mix of ambient traffic measurements and active probe measurements. This gives our system the flexibility of nominally monitoring a only a few ambient flows for anomalies while bringing in more intrusive methods of active tomography only when these nominal measurements leads one to suspect that some significant deviation from baseline has occurred. Several members of our team have shown that unicast active probing can provide accurate estimates of link delay and loss characteristics from a few edge measurements without any special cooperation of routers in the network [4, 5, 32, 7]. Recently these methods have been extended to cases where traffic patterns may be changing over the probing period [6]. The challenge will be to integrate tomography methods into the multivariate time series and DES modeling framework that we propose. Since these models introduce structured temporal and spatial dependencies into delay and loss statistics we must account for these in the tomographic reconstruction algorithm.

4. Collaborative Data Collection (Leader: Prakash)

- Development and dissemination of software (Prakash)
- Multi-site information aggregation (Teneketzis)
- Distributed computation/compression/complexity (Pradhan)
- Data collection from wireless basestations (Liu,Nobel,Wallach)? Still not sure how to fit wireless into story.
- Data collector authentication (Prakash)
- Packet Vault deployment (Nobel, Prakash)

(*Prakash, Liu, Nobel, Wallach, Pradhan*)

We will develop an open collaborative infrastructure for data collection and distributed pattern recognition for monitoring the global network. This will involve free and wide dissemination of open-source software tools for data collection, computation, and transmission to Rice and UM. SImilarly to the SETIhome program the software would take advantage of a participant's free cpu cycles to perform data collection and computations. One possible scheme is to allow the user to specify the level of collaboration byt allocating varying amount of resources (disk space) through specification of a finite memory window over which data is stored, processed and forwarded. While any particular site may not collect data continuously over any long time interval the aggregate of all collaborating sites will produce a patchwork of asynchronous snapshots of the network at different sites and different times that will be sewn together and analyzed at UM and Rice. all collaborating sites will time stamp their data according to NIST/GMT time-clock information which will be downloaded periodically. Need more here....

5. Applications (Leader: Liu)

- Distributed denial of service (DDoS) attacks (Liu)
- Multi-site intrusions (Wallach,Nobel)
- Robot larceny (Reidi)
- Wireless?
- Service monitoring/verification (Coates, Nowak)

(*Lafortune, Reidi, Wallach, Nobel, Coates, Nowak, Hero*)

4.a DDoS attacks

A major current focus in networking research is to address the threat of distributed denial of service (DDoS) attacks. A large fraction of DARPA's Fault Tolerant Networking funded projects and commercial products from both established companies, such as Cisco Systems, and a flock of startup companies, such as Arbor Networks, Asta, Mazu, Reactive, etc. are proposing deployment of Internet-wide infrastructure to combat DDoS attacks. To a large extent, all these approaches rely on distributed traffic correlation

capabilities to detect anomalies. The basic architecture proposed invariably involves distributed monitors, some form of distributed traffic correlators fed by these monitors, and installation of traffic filters on detection of traffic anomalies. The research issues in building such an architecture include: How to detect attacks with minimal false positives? How to mitigate the attacks? And how to do both in a scalable and timely manner? The contributions of this project to combating DDoS are three fold:

1) The novel joint application of multivariate time series and DES models for fast change point detection in traffic behavior. We will evaluate the sensitivity and timeliness of these models in their application to detecting DDoS attacks. Should these models prove useful, they can be retrofitted to any global infrastructure for combatting DDoS attacks.

2) Instead of attempting to minimize false positive in DDoS detection, we are designing an early detection and mitigation mechanism that provides degraded service when DDoS attack is suspected but is otherwise more forgiving to false positives. The main idea is to detect DDoS attack not based on global traffic patterns but on traffic destinations' service characteristics. The proposed mechanism will employ the change-point detection models to make local DDoS mitigation decisions based only on locally available information. To tolerate a high rate of false positives, the mitigation mechanism will first *delay* traffic towards selected destinations, dropping them only when necessary. We will investigate the effectiveness of such local decisions in destroying the global correlation of a DDoS attack. At the minimum, such mechanism may allow more time for a global DDoS detection and mitigation mechanism to react.

3) Finally, our change-point detection models can be used as a benchmark to verify the effectiveness of DDoS architectures proposed by researchers and companies in the field.

5. Education (Leader: Nobel)

- Global network security curriculum (Nobel)
- Networking security internships (Liu)
- Networking laboratory (Liu, Nobel, Wallach, Prakash)
- Webcast seminars and workshops (Wallach)
- Summer UG and high school program (Hero)
- Network gaming competition (Prakash)

(*Liu, Prakash, Wallach, Noble, Baraniuk, Hero*)

The scope of this effort will provide many opportunities for undergraduate and graduate students to be involved in research. We also plan to include precollege students in this project through a summer internship program with private and public schools in Houston and Metro Detroit.

We are committed to making an impact on education for which we propose the following.

1. College students will be involved in developing a software tool for visualization of parameters in the spatio-temporal model and connectivity from the network tomography software.
2. A summer internship program will be developed to expose high school students to the areas of computer security and signal processing. This program will be coordinated with the existing high school program called Camp CAEN (Computer Aided Engineering Network) at UM.
3. One of the co-PIs has developed a random topology generator called Inet [24] that has been used by researchers in the networking field to generate realistic inter-domain topologies of the Internet. College students will be involved in testing out our anomaly detection algorithms on simulated networks.
4. We will develop a networking laboratory at Rice and UM which will serve the dual purpose of educating students in network security and providing data traces for our research. For more details on this lab see the budget justification. In this lab students will learn about attack strategies and

mitigation (on a scaled down "private" network emulation) and will also generate attack scenarios for testing. A competition will be held between Rice University and University of Michigan during the summer of each year with prizes going to those students who develop the best attack strategy and the best thwarting strategy. Judges will be drawn from the team of co-PI's and collaborators. High school students will also participate in this competition through the summer internship program.

5. Undergraduate students will help acquire and analyze real network data at Rice and UM in the context of classroom instructional laboratories and independent study projects.

The proposed research will also help us improve our curriculum in networking and security. We have recently made senior design a requirement for our degree programs. The proposed research should help us define senior design projects in networking and security courses at the undergraduate level. At the graduate level, co-PI Prakash is planning to teach a pilot course on network security in Winter 2002, which we plan to make a regular course. Intel has recently donated 25 laptops to co-PI Prakash to allow ad hoc networks to be set up for projects related to networking and security graduate curriculum.

References

[1] A. Aghasaryan, E. Fabre, A. Benveniste, R. Boubour, and C. Jard, "Fault detection and diagnosis in distributed systems: An approach by partially stochastic Petri nets," *Discrete Event Dynamic Systems: Theory and Applications*, vol. 8, no. 2, pp. 203–231, June 1998.

[2] M. Basseville and A. Benveniste, editors, *Detection of abrupt changes in signals and dynamical systems*, Springer Lecture Notes in Control and Information Sciences, 1986.

[3] R. J. Beran, *Statistics for Long-Memory Processes*, Chapman & Hall, 1994.

[4] M. Coates and R. Nowak. *Network Loss Inference using Unicast End-to-end measurement*, Sep. 2000.

[5] M. Coates and R. Nowak, "Network delay distribution inference from end-to-end unicast measurement," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Proc.*, May 2001.

[6] M. Coates and R. Nowak, "Sequential Monte Carlo inference of internal delays in nonstationary communication networks," to appear in *IEEE Trans. Signal Processing, Special Issue on Monte Carlo Methods for Statistical Signal Processing*, 2002.

[7] M. Coates, A. Hero, R. Nowak, and B. Yu, "Large scale inference and tomography for network monitoring and diagnosis," *IEEE Signal Processing Magazine*, vol. to appear, , May 2002. `http://www.eecs.umich.edu/~hero/comm.html`.

[8] R. W. Crow, "Quickest detection for sequential decentralized decision systems," *IEEE on Aerospace Electronics Systems*, vol. AES-32, no. 1, pp. 267–283, Jan. 1996.

[9] H. T. Şimşek, R. Sengupta, S. Yovine, and F. Eskafi, "Fault diagnosis for intra-platoon communication," in *Proc. 38th IEEE Conf. on Decision and Control*, December 1999.

[10] R. Debouk, S. Lafortune, and D. Teneketzis, "Coordinated decentralized protocols for failure diagnosis of discrete-event systems," *Discrete Event Dynamic Systems: Theory and Applications*, vol. 10, no. 1/2, pp. 33–86, January 2000.

[11] J. A. Fessler and A. O. Hero, "New complete-data spaces and faster algorithms for penalized-likelihood emission tomography," in *Proc. of IEEE Nuclear Science Symposium and Medical Imaging Conf*, pp. 1897–1901, San Francisco, November 1993.

[12] J. A. Fessler and A. O. Hero, "Complete data spaces and generalized EM algorithms," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, pp. IV.1–IV–4, Minneapolis, MN, April 1993.

[13] J. A. Fessler and A. Hero, "Penalized maximum likelihood image reconstruction using space alternating generalized EM algorithms," *IEEE Transactions on Image Processing*, vol. 4, no. 10, , October 1995.

[14] J. A. Fessler and A. O. H. III, "Space alternating generalized expectation- maximization algorithm," *IEEE Transactions on Signal Processing*, vol. 42, no. 10, pp. 2664–2677, October 1994.

[15] J. Fessler and A. Hero, "Cramer-Rao bounds for biased estimators in image restoration," in *Proc. of 36th IEEE Midwest Symposium on Circuits and Systems*, Detroit, MI, Aug. 1993.

[16] A. O. Hero, "The influence of the choice of complete data on convergence of EM-type algorithms," in *Proc. of the IEEE Workshop on Statistical Signal and Array Processing*, pp. 74–77, Victoria, Oct. 1992.

[17] A. O. Hero, "Theoretical limits for optical position estimation using imaging arrays," in *Actes du Colloque GRETSI*, pp. 793–796, Juan-les-Pins, France, Sept. 1991.

[18] A. O. Hero and J. A. Fessler, "Convergence in norm for alternating expectation-maximization (EM) type algorithms," *Statistica Sinica*, vol. 5, no. 1, pp. 41–54, 1995.

[19] A. O. Hero and J. A. Fessler, "A recursive algorithm for computing CR-type bounds on estimator covariance," *IEEE Trans. on Inform. Theory*, vol. 40, pp. 1205–1210, July 1994.

[20] A. O. Hero and J. A. Fessler, "A fast recursive algorithm for computing CR-type bounds for image reconstruction problems," in *Proc. of IEEE Nuclear Science Symposium*, pp. 1188–1190, Orlando, FA, Oct. 1992.

[21] A. O. Hero, J. A. Fessler, and M. Usman, "Exploring estimator bias-variance tradeoffs using the uniform CR bound," *IEEE Trans. on Signal Processing*, vol. 44, pp. 2026–2042, Aug. 1996. http://www.eecs.umich.edu/~hero/det_est.html.

[22] A. O. Hero, Y. Zhang, and W. L. Rogers, "Consistency set estimation for PET reconstruction," in *Proc. of Conference on Information Science and Systems*, Johns Hopkins Univ., MD, March. 1993.

[23] A. Hero, M. Usman, A. Sauve, and J. Fessler, "Recursive algorithms for computing the Cramer-Rao bound," *IEEE Trans. on Signal Processing*, vol. SP-45, no. 3, pp. 803–807, 1997.

[24] C. Jin, Q. Chen, and S. Jamin, "Inet: Internet topology generator," Technical Report CSE-TR-433-00, University of Michigan, EECS Dept., 2000. http://topology.eecs.umich.edu/inet.

[25] G. V. Moustakides, "Quickest detection of abrupt changes for a class of random processes," *IEEE Trans. on Inform. Theory*, vol. IT-44, no. 5, pp. 1965–1968, Sept. 1998.

[26] K. Nagarajan and T. Zhou, "A new resource allocation scheme for VBR video sources," in *Proc. Asilomar Conf. on Signals, Systems, and Computers (ASILOMAR)*, Oct. 2000.

[27] Y. Pencolé, "Decentralized diagnoser approach: Application to telecommunication neworks," in *Proc. DX'00: Eleventh International Workshop on Principles of Diagnosis*, A. Darwiche and G. Provan, editors, pp. 185–192, June 2000.

[28] A. R. Reibman and L. W. Nolte, "Optimal design and performance of distributed signal detection systems with faults," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 10, pp. 1771, 1990.

[29] V. Ribeiro, M. Coates, R. Riedi, S. Sarvotham, B. Hendricks, and R. Baraniuk, "Multifractal cross-traffic estimation," in *Proceedings of the ITC Specialist Seminar on IP Traffic Measurement, Modeling and Management*, Monterey, CA, Sept. 18-20 2000.

[30] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis, "Diagnosability of discrete event systems," *IEEE Trans. Automatic Control*, vol. 40, no. 9, pp. 1555–1575, September 1995.

[31] P. Sen, *Sequential non-parametrics: invariance principles and statistical inference*, Wiley, New York, 1981.

[32] M.-F. Shih and A. O. Hero, "Unicast inference of network link delay distributions from edge measurements," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, Salt Lake City, UT, May 2001. http://www.eecs.umich.edu/~hero/comm.html.

[33] A. Shirayev, *Optimal Stopping Rules*, Springer, 1978.

[34] D. Teneketzis and P. Varaiya, "The decentralized quickest detection problem," *IEEE Trans. Automatic Control*, vol. AC-29, no. 7, pp. 641–644, 1984.

[35] M. Usman and A. Hero, "Recursive CR-bounds: algebraic and statistical acceleration," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, Adelaide, Australia, April 1994.

[36] M. Usman, A. Hero, and J. A. Fessler, "Uniform CR bound: implementation issues and applications to image reconstruction," in *Proc. of IEEE Nuclear Science Symposium and Medical Imaging Conf*, pp. 533–537, Virginia Beach, November 1994.

[37] M. Usman, A. Hero, J. A. Fessler, and W. Rogers, "Bias-variance tradeoffs analysis using uniform CR bound for a SPECT system," in *Proc. of IEEE Nuclear Science Symposium and Medical Imaging Conf*, pp. 1463–1467, San Francisco, November 1993.

[38] M. Usman, A. Hero, and W. L. Rogers, "Performance gain analysis for adding a vertex view to standard SPECT," in *Proc. of 36th IEEE Midwest Symposium on Circuits and Systems*, August 1993.

[39] Y. Zhang, A. O. Hero, and W. L. Rogers, "Simultaneous confidence intervals for image reconstruction problems," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, pp. V.317–320, Adelaide, April 1994.

[40] Y. Zhang, A. O. Hero, and W. L. Rogers, "A bounded error estimation approach to image reconstruction," in *Proc. of IEEE Nuclear Science Symposium*, pp. 966–968, Orlando, FL, Oct. 1992.

5. Management Plan

Next page

To accomplish the research and education aims of this project requires a *focused large scale and multi-disciplinary effort*. The coordination of co-PI's, collaborators, and students from three colleges (UM, Rice, and McGill), several high schools, and a government agency (NASA Johnson Space Center) requires a tight management plan. Central to this plan will be to channel the activities of individual investigators into broad and productive collaborations that cross traditional boundaries which have separated networking and computer security from signal processing and statistics. To facilitate and enhance such collaborations all supported students will have at least two co-PI's in their thesis committees and supported UM students will be co-supervised. In addition, we will develop and team-teach courses combining trusted computing, signal processing, and network security. Furthermore, the focus on real networking problems and solutions provided by our non-academic collaborators at NASA and elsewhere will be used to integrate these diverse research acivities.

Team Management: The first year we will have a face-to-face kick off meeting involving all co-PI's and collaborators. Industry and government representatives will be invited to attend this meeting. In addition to these larger meetings there will be several meetings over each year of the project. These will include monthly meetings via webcast/videoconference/teleconference to assess progress and explore new ideas for education and research. Several co-PI's will visit Rice to meet with co-PI Rob Nowak and NASA collaborators twice a year to help specify suitable simulations on portions of the NASA network for testing our algorithms. One or two day year-end meetings of all coPI's and collaborator will take place in Ann Arbor. At this meeting co-PI's and collaborators will present previous year's research and education results. Industry and government representatives will be also invited to participate. At these meetings we will set or refine specific goals for the next year.

Electronic Dissemination: A website will be created to archive research reports and articles, sample data traces, interactive software, course materials, and announcements. This website will be accessible to the public. We will hire students to create an informative and appealing website with the help of coPI's. A graduate student will be appointed as webmaster for web maintenance and administration.

Summer Internship Program: A summer program for high school students will be organized every year. This program will consist of two parts: 1) an intensive two week camp where they will learn about computer networks in a hands on educational and recreational environment; 2) participation in the attack simulation activities, including the UM vs. Rice competition (see Education Section of white paper). We will seek out kids from a variety of backgrounds, including under-represented socio-economic groups in Houston and Detroit, to participate. We will work closely with Camp CAEN (`http://campcaen.engin.umich.edu/`), a computer exploration summer camp at the College of Engineering at UM, to recruit high school students and to develop a network security curriculum.

Continuing Education: Every year we will organize a sequence of short courses. The short courses will be run during 4 weeks during the summer session at the University of Michigan and will cover aspects of networking such as security, network tomography, sensor networks, and traffic modeling. These courses will be aimed at networking professionals. The courses will be offered in cooperation with UM's existing summer short course program in order to benefit from existing infrastructure.

Workshops: We will organize an annual or biennial workshop on Signal Processing for Networks which will have keynote speakers, special invited sessions, and contributed sessions. As contrasted to the short courses, which are aimed at continuing education of networking professionals, the workshops will be aimed toward the academic community. The workshops will take place over three successive days. They will have a strong education component involving tutorials on network traffic measurement, network security, and network modeling. At each workshop there will be a session on novel classroom teaching methods for lower level signal processing and networking courses. We will also have sessions featuring papers presented by students (undergraduate and precollege) on networking projects completed over the previous year in connection with this grant.

Budget Justification

1. We are including Dr. Mark Coates on this grant as international collaboration. Dr. Coates will start as an Assistant Professor at McGill University in Montréal Canada in Winter 2002. He has been a research collaborator with Profs. Nowak and Hero and his expertise in stochastic optimization using Monte Carlo Markov Chain will be crucial to our research aims. No ITR funds are requested for Dr. Coates. However, if this ITR is funded we will apply for additional funding in the NSF Office of International Programs.

2. Funding for two post-docs are requested for this project. This will provide post-graduate education and specialized research assistance to co-PI's.

3. $500K is requested for equipment to set up the isolated network environments at Rice and UM to be used for educational of college and precollege students on network security, generate attack scenarios in the context of the summer gaming competition, and serve as a small scale testbed for co-PI's and their graduate students. The networks will consist of two LAN's each consisting of 50-75 PC terminals and routers (both wireless and wired) at Rice and UM. Communications between the LAN's at Rice and UM will be through the Internet using a pair of dedicated DNS servers. More details will be given in the full proposal.

4. $500K is requested for two full time lab adminstrators to help develop and maintain the lab environments at Rice and UM over the five years of the grant.

5. $30K per year is requested to fund the summer internship program for high school students. This will be enough to set up the program each year (recruiting, publicity, etc) and provide scholarships to 20-30 students to attend Camp CAEN from Rice and UM.

6. Given the scope of the educational aims of the project we request $30K per year for funding undergraduate students to help develop the network laboratory environment, to help develop web tools, and to collect data traces. Supplementary funding for undergraduate research projects will also be requested from NSF if we are funded.