

# ITR: Modeling and Prediction of Network Behavior

Alfred Hero(PI), Mark Coates(co-PI), Sugih Jamin(co-PI),  
Stéphane Lafortune(co-PI), Mingyan Liu(co-PI), George Michailidis(co-PI),  
Robert Nowak(co-PI), Sandeep Pradhan(co-PI), Atul Prakash(co-PI), Demosthenis Teneketzis(co-PI)

October 30, 2001

## 1 Summary

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 understand so as to predict the aggregate behavior of such large heterogeneous networks. The sheer size and complexity of our critical data and sensor networks makes the modeling and prediction problem extremely challenging. However, a solution to this problem will be essential for detecting anomalous traffic patterns and mitigating potentially debilitating congestion or attacks. Better prediction models will lead to improved offline and on-line algorithms for identifying anomalous/malicious behavior, for implementing dynamic routing and flow control for congestion, for network provisioning, and for remote monitoring and service verification. Our approach is a potent combination of emerging techniques in signal processing and networking research including: novel probing methods to perform network measurement and tomography; distributed data compression and coding for complexity reduction; and change detection and localization performed with hierarchical modeling of traffic statistics using a spatio-temporal traffic model whose parameters can change according to state changes in a discrete event system model.

Crucial to predicting aggregate behavior of networks will be our ability to determine traffic flow statistics throughout the network and to characterize what constitutes a significant change in flow patterns. This will require measurements of internal link traffic and a reliable dynamical model to establish a 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 and model-identification presents extremely challenging 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 nodes, called edge nodes, of the network. Events which cause changes in parameters of a multivariate spatio-temporal traffic flow model will then be detected and isolated according to a deterministic structured dynamical discrete event systems (DES) model. DES models have been widely applied to networks for fault detection and isolation. This is the first time that these models have been proposed for detection of anomalous statistical behavior of traffic in large scale networks.

Several important applications will be considered including: verification of evolutionary models of distributed denial-of-service (DDoS) attacks; development of a large scale heterogeneous network simulation tool; real-time anomaly detection in multimedia surveillance networks; detection/localization of congestion and queuing delays; bilateral (supply/demand) verification of service; remote detection of ephemeral network disruptions, e.g. due to link failures or attacks. **should we include some more applications or delete some?** Our methodology and algorithms will be validated by running experiments on real large scale networks such as the Internet and from real private networks at the NASA Johnson Space Center. This research will be coupled to a vigorous education component which will involve high school and college students in data collection; network gaming for simulation of attack scenarios; and

distance learning.

## 2 Proposal Overview

We are a multi-disciplinary team of researchers in the fields of adaptive signal processing, Bayesian networks and optimization, network traffic measurement, discrete event systems, multivariate statistics, tomography, distributed data compression, sensor networks and stochastic control. We propose to investigate techniques for rapid detection and localization of the onset of spatio-temporal changes in global network traffic which can be used for predicting 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 *temporally and spatially* dependent traffic models, inference of characteristics of internal links that transport the traffic, and measurements that can be acquired 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 new approaches. Our approach is systematic and integrated strategy of modeling and measurement that overcomes the limitations of previous approaches.

Specific innovations of this proposal are:

1. Use of network tomography methods to reduce the number of monitoring sites required to accurately estimate internal traffic patterns and link statistics.
2. Development of multi-channel traffic models which incorporate physical layer information, account for short-range and long-range spatio-temporal dynamics across the network, and can be adaptively identified from incomplete measurements via tomographic techniques.
3. Integration of discrete event system models into the spatio-temporal traffic model for classification of detectible sequences of changes in network transport characteristics.
4. Application of distributed detection algorithms to the above models to tracking changes in network connectivity, changes in loss and delay distributions, changes in traffic correlation patterns, and other potential anomalies.
5. Experimental validation with real multivariate data traces from selected sites on the Internet and from the large private network at NASA JSC (Letter will accompany full proposal).
6. Creation of software tools for multi-layer dynamic visualization of changes in network connectivity, delay and loss rates, and traffic.
7. Development of an instructional laboratory on network security using an isolated IP network of terminals running the visualization software to teach students about security through a multiterminal computer game. In this game students will match wits against each other on developing and mitigating attack strategies.

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 failures; 2) an integrated and flexible multiple time series model 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 high school interns and 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 Statistics, Networking, Signal Processing, and Software.

### 3 Background on Group

**This should be seamless paragraph describing the synergy of our group for this project.** Many of the co-PI's on this project have been leaders covering a range of networking research areas including: software and Internet measurements [9, 19, ?], network tomography [7, 26], traffic modeling [?], discrete event network failure models [?], etc... Many of us have a history of fruitful collaboration on these problems [?, ?, ?] and this project represents a natural merging of these complementary activities for the purpose of prediction of network behavior, detection of anomalous events, and improving education in network security. **Need add more material here**

### 4 Methods

The research approach in this proposal represents a dramatic departure from existing activities in traffic analysis and network monitoring. The traffic models will be learned recursively over space-and-time from a few monitoring sites at the edge of the network. After achieving steady state, the residual model errors will be used to track temporal and spatial changes which might be indicative of a disruption in traffic. To reduce the intrinsic communication complexity of distributing statistical traffic information between the monitoring nodes we will investigate methods of distributed data compression. Optimal quickest detection algorithms will be investigated for timely detection of significant changes in traffic or topology of the network. Detectable changes will be classified based on a discrete event model which accounts for probable sequences of events leading to link failures due to natural congestion or attacks. We describe elements of this approach in more detail below.

**Spatio-Temporal Traffic Models:** The fractional autoregressive integrated moving average (FARIMA) model has been previously applied to single stream network traffic and has been shown to reliably capture both the long term dependency of a process as well as its short range dependency [4, 22]. We will extend the FARIMA model to a new multiple stream spatio-temporal FARIMA model defined on a random graph determined by network connectivity. The application of this model will require: i) model identification from incomplete measurements, i.e. determining the fractional order, ARMA coefficients, and spatial dependency parameters; ii) assessing the performance (goodness-of-fit) of the model. Carrying out these two basic tasks for large data sets will require advances both at the analytical and at the algorithmic level. Once developed the model identification algorithms will be applied to: (i) generate synthetic traces and validating them against experimental data in order to understand the evolution of different anomalous traffic scenarios; (ii) on-line implementations of quickest detection algorithms for discerning onset of such traffic patterns.

Increasingly, our most critical networks incorporate both wireless and wired components, e.g. the FAA airphone network, the Police-band wireless data network. or NASA's Mission Control Network. Thus by explicitly including physical layer interactions and terminal types, e.g. data or video sensor, into the model, more accurate predictions can be performed. Conversely, it could well be possible to classify the physical medium of unknown parts the network from estimates of the link model parameters. For example the effects of signal propagation and fading, channel bandwidth, and transmitter modulation provide characteristic signatures for different links which determine single link packet loss, delay and link bandwidth distributions. Furthermore, wireless links generate co-channel multi-user interference (MUI) reducing SNR (signal-to-noise ratio) and significantly correlating link packet losses and delay between close links.

#### **Discrete Event Models of Link Failure**

*Demos and Stephane's input*

**Model selection:** The spatio-temporal model represents a summary statistic of the status of the network under different conditions and at different time and spatial scales. The big advantage is that with a small number of parameters one can describe and localize different normal and anomalous situations characterizing the network. Therefore, we can associate different labels, describing the "status" of an

underlying subnetwork (e.g. normal, congested, under attack, etc.), with different estimated values of the model parameters. The problem then becomes to predict the correct status label given a set of model parameters estimated from traffic data. This represents a novel application of the traditional model selection approach in statistics and machine learning [24], where one tries to design efficient decision rules that assign objects to a fixed number of prespecified classes based on their measurements on a number of attributes. In our application the goal becomes to assign models to a fixed number of network states based on their estimated parameters. To enhance classification performance we will use methods of aggregating classifiers via *boosting* [14, 25] and other randomized tree voting methods [1].

*George - we need combine these two. Also need to describe how the DES will be implemented to narrow down the trajectories of the model parameters*

**Quickest Detection:** The goal will be to detect as early as possible the occurrence of the change point, but at the same time controlling the false signal rate [27, 2, 21]. While a limited literature on distributed sequential change point detection schemes exists, e.g. [23, 8], the current state-of-the-art falls short of providing adequate solutions demanded by the needs of our problem. It will be necessary to develop new methodologies for dealing with distributed change point detection within the dynamical spatio-temporal modeling framework.

**Network Tomography and Inference:** Measurements from only a few distributed nodes will be used to do change point detection, localization, and model identification as discussed above. Passive and active probing using multicast methods [30, 6, 10] and unicast methods [7, 26] have been developed to estimate link delay and loss characteristics from such edge measurements. As buffer delay is directly related to traffic rate on the link, the spatio-temporal model can be used to extend our tomography methods to spatially dependent traffic. Using the bounded error tomography methods of [18] we can generate  $(1 - \alpha)100\%$  simultaneous confidence regions for the link parameters. These will be used to specify control levels for detecting incipient changes.

*Rob some more input? Add or delete as you wish here*

**Feasibility Analysis:** We will derive both parametric and non-parametric lower bounds on intrinsic estimation, prediction and change detection performance for our spatio-temporal network models. These types of bounds are extremely useful for evaluating mission feasibility [3, 16, 15, 17]. Such an analysis can be used to investigate the minimum quantity/quality of measurement data required to estimate a parameter to a given accuracy. Such analysis will give insights into the minimum discernible changes in delay statistics, traffic rates, and other parameters relevant to detection of attacks and other events.

## 5 Impact on Education

The scope of this effort will provide many opportunities for undergraduate and graduate students to be involved in research. We propose to develop innovative educational tools through this research which might include the following.

1. We can propose to host a yearly workshop to disseminate the results and get feedback from others in the field? If so we can ask for funds to invite some outside speakers.
2. Undergraduate 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. This tool will be written in matlab using the GUI toolbox. XXXX will supervise these project in the context of the senior level design course EECSXXX in the Computer Science and Engineering Division.
3. Is there good software out there for emulating a fairly large representative network backbone? If so we could have students involved in testing out our algorithms on simulated networks. What do you think Sugih?
4. How about incorporating some of our simulation software into a computer game where students try to construct attacks on each others simulated networks? The students would learn about attack

strategies and mitigation (on a scaled down “private” network emulation) and we would get a bunch of scenarios to test out of this. We could introduce a contest for junior and senior students where the winners would develop the best attack strategy and the best thwarting strategy. If this could fly we might even involve high school students in this gaming enterprise...

5. Undergraduate students will help acquire and analyze network data at Rice and UM.
6. Rob: perhaps we could develop some of this stuff on the web under the aegis of Connexions? If so we can involve some Connexions UG’s and link the proposal to the educational activities already going on at Rice.
7. George: any creative ideas on including statistics students?
8. Sandeep: if you put together a sensor network you can have some students involved in setting it up? Any other ideas?

*Atul: include some stuff on software development, inclusion of software in a course on network security, any thing else*

## 6 Focus Applications

### DDoS Applications:

*Sugih’s stuff here*

#### Network Provisioning Applications

*Sugih?*

Based on detected deviation from the baseline traffic model, a network administrator can be warned of shifts in traffic patterns such that appropriate action, such as increased provisioning of network bandwidth can be instituted before overloading sets in. Conversely, by detecting deviation from normative traffic pattern, a customer may be able to verify that it is receiving the level of service contractually guaranteed by its provider. When congestion or faults occur on the network, we may be able to isolate their location by comparing measured traffic against the baseline model at various points on the network. A software tool for topology generation and simulation will be a specific byproduct of our development of a spatio-temporal traffic model, discussed in more detail below.

#### Simulation Applications

*Sugih?*

The need for realistic random topologies in simulations has long been recognized by researchers working on routing and multicast protocols, e.g. [5, 29, 28]; more recently, the need for realistic random topologies has also been voiced by researchers studying traffic dynamics and protocol behavior [20, 13, 12]. In recognition of this, several topology generators have been proposed in the literature. The most recent one, proposed by one of the co-PI’s [19], is called *Inet* which takes advantage of power-law relationships [11] in its construction of random topologies. Unfortunately all extant topology generators, including our own, model only connectivity between nodes but not the properties of the connections, e.g. the generated links are not assigned realistic bandwidth, propagation delay, and loss rate. Specifically, we propose to capture invariants of Internet link properties, such as distribution of link bandwidth at different parts of the Internet and distributions of propagation delay and loss rate characteristics. Due to the large number of links on the Internet and hesitancy of some network owners in allowing measurement and characterization of their networks, we do not propose to measure the whole Internet, but only a sampling thereof. Our generator will be used in conjunction with network layer routing protocols and detailed physical layer models for single link performance to build both router-level and AS-level maps of heterogeneous networks.

#### Multimedia Surveillance Network Applications:

Several important issues will be studied including: distributed data compression; traffic modeling for mixed flows of data and multimedia streaming; quickest detection from incomplete or degraded image/audio or other measurements;.....

*Mingyan and Sandeep's stuff here*

**Network Tomography Applications**

*Rob's stuff*

## 7 Management Plan

The involvement of co-PI's and students from three schools (UM Engineering, UM LSA, and Rice) and Govt Agencies (NASA and perhaps DOE) and the substantial educational activities of this project require a tight management plan. The management aspects of the project will be coordinated by the PI (A. Hero) with the help of an administrative secretary. While all the coPI's know each other well and have a history of collaborations in research and education, we will have regular research meetings to ensure timely progress of the project. The first year we will have a face-to-face kick off meeting involving all coPI's in order to set the goals for the rest of the year. Industry and government collaborators will be invited to attend this meeting. The last year of the proposal we will host a workshop on Mathematical Network Modeling and Analysis which will take place over two successive days and have a mix of contributed and invited speakers.

In addition to these larger meetings there will be several meetings over each year of the proposal including:

- We will have monthly meetings via webcast/videoconference/teleconference to assess progress.
- The PI will visit Rice to meet with coPI Rob Nowak and XXX at NASA Johnson Space Center twice a year.
- Yearly meetings of all coPI's will take place in Ann Arbor which will be run as a mini-symposia with presentations of previous year's research results.



## References

- [1] Y. Amit and D. Geman, “Shape quantization and recognition with randomized trees,” *Neural Computation*, vol. 9, pp. 1545–1588, 1997.
- [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] B. Baygun and A. O. Hero, “Optimal simultaneous detection and estimation under a false alarm constraint,” *IEEE Trans. on Inform. Theory*, vol. 41, no. 3, pp. 688–703, 1995.
- [4] R. J. Beran, *Statistics for Long-Memory Processes*, Chapman & Hall, 1994.
- [5] L. Breslau and D. Estrin, “Design of inter-administrative domain routing protocols,” *Proc. of ACM SIGCOMM '90*, pp. 231–241, Sep. 1990.
- [6] R. Caceres, N. G. Duffield, J. Horowitz, and D. Towsley, “Multicast-based inference of network internal loss characteristics,” *IEEE Trans. on Inform. Theory*, vol. IT-45, pp. 2462–2480, Nov. 1999.
- [7] M. Coates and R. Nowak, “Network loss inference using unicast end-to-end measurement,” Technical report, Dept. ECE, Rice University, Mar 2000.
- [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] P. Danzig, S. Jamin, R. Cacerés, D. Mitzel, and D. Estrin, “An empirical workload model for driving wide-area TCP/IP network simulations,” *Journal of Internetworking: Research and Experience*, vol. 3, , 1992. **URL** <http://netweb.usc.edu/jamin/traffic/>.
- [10] N. G. Duffield and F. LoPresti, “Multicast-based inference of packet delay variance at interior network links,” in *Proc. of IEEE INFOCOM*, Tel Aviv, Mar. 1999.
- [11] M. Faloutsos, P. Faloutsos, and C. Faloutsos, “On Power-Law Relationships of the Internet Topology,” *Proc. of ACM SIGCOMM '99*, pp. 251–262, Aug. 1999.
- [12] A. Feldman et al., “Netscope: Traffic engineering for ip networks,” *IEEE Network Magazine*, 2000.
- [13] A. Feldman, A. Gilbert, P. Huang, and W. Willinger, “Dynamics of ip traffic: A study of the role of variability and the impact of control,” *Proc. of ACM SIGCOMM '99*, pp. 301–313, Aug. 1999.
- [14] Y. Freund and R. Schapire, “A decision theoretic generalization of online learning and an application to boosting,” *Journal of Computer and System Sciences*, vol. 55, pp. 119–139, 1997.
- [15] J. D. Gorman and A. O. Hero, “Lower bounds for parametric estimation with constraints,” *IEEE Trans. on Inform. Theory*, vol. IT-36, pp. 1285–1301, Nov. 1990.
- [16] A. O. Hero, “Lower bounds on estimator performance for energy invariant parameters of multi-dimensional Poisson processes,” *IEEE Trans. on Inform. Theory*, vol. 35, pp. 843–858, July 1989.
- [17] 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**.
- [18] A. O. Hero, Y. Zhang, and W. L. Rogers, “Tomographic feature detection using parallelotope bounded error algorithm,” in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, volume 4, pp. 2849–2852, Munich, April 1997. **http://www.eecs.umich.edu/~hero/tom\_imaging.html**.
- [19] 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**.

- [20] D. Mitzel and S. Shenker, "Asymptotic resource consumption in multicast reservation styles," *Proc. of ACM SIGCOMM '94*, 1994. URL <http://netweb.usc.edu/mitzel/Sigcomm94/sigcomm94.ps.Z>.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] B. D. Ripley, *Pattern recognition and neural networks*, Cambridge U. Press, 1996.
- [25] R. Schapire, "The strength of weak learnability," *Machine Learning*, vol. 5, pp. 197–227, 1990.
- [26] 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>.
- [27] A. Shirayev, *Optimal Stopping Rules*, Springer, 1978.
- [28] L. Wei and D. Estrin, "The trade-offs of multicast trees and algorithms," *Int'l Conf. on Computer Communications and Networks*, 1994.
- [29] W. Zaumen and J. Garcia-Luna Aceves, "Dynamics of distributed shortest-path routing algorithms," *Proc. of ACM SIGCOMM '91*, pp. 31–42, Sep. 1991.
- [30] A.-G. Ziotopoulos, A. O. Hero, and K. Wasserman, "Estimation of network link loss rates via chaining in multicast trees," in *Proc. IEEE Int. Conf. Acoust., Speech, and Sig. Proc.*, Salt Lake City, UT, May 2001. <http://www.eecs.umich.edu/~hero/comm.html>.