

**Title:** Early Warning of Network Catastrophes

**Lead Laboratory:** ITL

**Proposal Champions:** Kevin Mills/772, Chris Dabrowski/772, Jim Filliben/776, Fern Hunt/771, Bert Rust/771

**Proposal Abstract:** Society grows increasingly dependent on networked systems, such as the Internet, computational clouds, and content distribution networks, in which interactions among millions of components lead to dynamic global patterns that cannot be predicted by analyzing behavior of individuals. Such patterns include catastrophic events [1-10] that arise from cascading failures, malicious attacks, and sudden load increases, which can drive systems from normal operating ranges to congested states, leading to widespread outages for significant periods. To date, no practical methods exist to provide early warning of such catastrophes. Thus, our nation's growing dependence on networked systems will continue to be plagued with costs from adverse consequences of large-scale failures.

Over the past decade, researchers in physics and other disciplines have devised a mathematically based theory to explain processes leading to catastrophes in networks. Related studies [11-19] demonstrate that such events are preceded by widespread changes in self-organized patterns, arising from increasing spatiotemporal correlation. Such patterns manifest ultimately as *phase transitions*, coincident with sudden changes in selected (so-called “*order*”) parameters, and foreshadowed by increases in variance and long-range dependence. While this theory could form a basis for early warning of network catastrophes, two barriers exist. First, the theory relies on abstract models that use topologies and behaviors unrepresentative of real networks. We will evaluate and validate the theory in realistic network models. Second, the theory assumes measurement methods impractical in real networks. We will develop measurement methods that can be deployed in real networks to give early warning of network catastrophes.

If successful, this project will provide measurement methods to enhance commercial network control and management systems [20]. The project “...will create a strong foundation of system measurement that has not existed before...[to] help avoid...real-life network failures...” [21], improving resilience of networked systems in many applications [22].

**Context of the proposed research:** Commerce, government, utility grids, and even modern social interactions depend increasingly on large information systems, based on Internet technology. Disruptions of such systems, which appear likely to increase in scale and duration, incur significant costs. For example, Table 1 summarizes results from a study [23] of the cost of network failures for six companies in different industry sectors. As shown, companies incurred losses costing as much as \$100K/hour from either complete outages or periods of degraded performance. Extrapolating to multiple companies, affected for extended periods by failures within large Internet service providers, implies staggering costs in aggregate. For example, another study [24] of outages across companies and industries estimated overall costs averaging about 3.6% of annual revenues. Such outages continue to occur [1-10], and the scope and cost will increase in the future, as the nation's reliance on networked systems grows.

Over the past decade, academic researchers have devised a mathematically based theory to explain processes that lead to catastrophic events in networked systems. Related studies demonstrate that onset of catastrophic events is preceded by widespread changes in self-organized patterns, arising from increasing spatiotemporal correlation. Such patterns manifest

ultimately as phase transitions, coincident with sudden changes in selected order parameters, and foreshadowed by increases in variance and long-range dependence, as a system approaches a critical point. Table 2 provides a summary of nine such studies [11-19], which demonstrate phase transitions and identify various precursor signals that arise when a network approaches a critical point. We will use the studies by Sarkar et al. [17] and Ryalova et al. [19] to introduce some key elements of the theory of network phase transitions.


Table 1. Summary of Infonetics study of network downtime costs in six companies; the table also gives the percentage of incidents due to complete outages vs. performance degradations

Sector	Revenue/Year	Downtime Cost	Cost/Hour	Outages	Degradations
Energy	\$6.75 Billion	\$4.3 Million	\$1624	72%	28%
High Tech	\$1.3 Billion	\$10.2 Million	\$4167	15%	85%
Health Care	\$44 Billion	\$74.6 Million	\$96,632	33%	67%
Travel	\$850 Million	\$2.4 Million	\$38,710	56%	44%
Finance	\$4 Billion	\$10.6 Million	\$28,342	53%	47%

Table 2. Summary of selected theoretical research related to network phase transitions

Year	Researchers	Location	Model	Metrics	Precursor Signal
2001	Sole & Valverde [11]	Spain & USA (SFI)	2D Lattice	Packet Delay, Queue Length, Throughput	Self-similarity in log-log plot of power vs. freq.
2002	Woolf et al. [12]	UK	2D Lattice	Packet Delay, Queue Length, Throughput	Long-Range Dependence (LRD) in time-series autocorrelation
2004	Arrowsmith et al. [13]	UK	Triangular & Hexagonal Lattice	Packet Delay, Queue Length, Throughput	LRD shown with Hurst parameter increases from rescaled range statistical (R-S) analysis
2005	Mukherjee & Manna [14]	India	2D Lattice	Packet Delay, Queue Length, Load per Node	Self-similarity in log-log plot of power vs. freq.
2007	Lawniczak et al. [15]	Canada	2D Lattice	Packets in Flight	LRD shown with Hurst parameter increases from R-S analysis
2007	Tadic et al. [16]	Slovenia, Austria, UK	Generated SF & UH	Packet Delay, Queue Length, Network Load	Systemic changes in network-load time series
2009	Sarkar et al. [17]	USA	2D Lattice	Packet Delay, Queue Length	Order parameter becomes positive
2009	Wang et al. [18]	China	Generated ER, WS, HK	Packets in Flight/Injected	Order parameter becomes positive
2010	Rykalova et al. [19]	USA	1D Ring & 2D Lattice	Packet Delay, Queue Length, Network Load	Increasing amplitude fluctuation in metrics

Topology Key: SF = Scale-Free UH = Uncorrelated Homogeneous ER = Erdos-Reyni Random  
 WS = Watts-Strogatz Small World HK = Holme-Kim variant of Preferential Attachment



(a) Average Transit Delay ( $D'$ ) vs. Network Load ( $\lambda$ ) (b) Transit Delay vs. Time for Four Network Loads ( $\lambda$ )

Figure 2. Illustrating a critical point in network load [17]

Figure 2 (a), from Sarkar et al. [17], plots the relationship between average transit delay ( $D'$ ) for delivered packets and increasing network load ( $\lambda$ ). Transit delays remain small until the network load reaches a critical point ( $\lambda_c$ ), where a phase transition occurs, leading to a steep increase in packet transmission delays. Figure 2(b) demonstrates time series for packet delays at various load values: below ( $\lambda = 0.25$ ), near ( $\lambda = 0.30$ ) and above ( $\lambda = 0.35$  and  $\lambda = 0.40$ ) the critical point. The plots suggest that increases in the slope of time series of selected measured variables could signal crossing of a critical point, allowing network managers to be alerted prior to complete network collapse.

Other researchers have identified increasing systemic correlations emerging as a network nears a critical point. Such correlations manifest as long-range dependence and increasing amplitude fluctuation in various measured variables. For example, Figure 3, from a study by Rykalova et al. [19], shows fluctuations over time in the number of messages in a network at two load levels—(a)  $\lambda = 0.15$  and (b)  $\lambda = 0.1996$ —for a network where the critical point is  $\lambda_c = 0.2$ . Notice that amplitude fluctuations increase by two orders of magnitude, and are highly temporally correlated, as the critical point is approached.




Figure 3. Number of messages vs. time for two load levels: (a)  $\lambda=0.15$  and (b)  $\lambda=0.1996$

As shown in Table 2, other studies demonstrate similar phase-transition phenomena for abstract network models—models that lack some essential characteristics of real networks. The proposed ITL team has significant experience building computationally tractable, realistic network models [25], and applying such models to study [26] macroscopic behavior arising from congestion. From that previous work, the ITL team has identified key network characteristics (e.g., topology, relative forwarding speeds, source and receiver distributions, user behavior, and protocol functions) that must be included in realistic network models. Further, the ITL team has access to a network emulation facility that can be used to validate simulation results against laboratory networks, and to evaluate measurement methods, and related software, in operational settings. Using this knowledge, experience, and capability, the ITL team is uniquely positioned to investigate and leverage the theory of network phase transitions in real networks, and to create practical measurement methods based on that theory.

**Technical plan:** We propose to build on a promising theory to establish practical, measurement-based methods to provide early warning of impending catastrophes in networked systems. Our target milestone will be to develop measurement methods and tools to monitor spatiotemporal behavior in deployed networks, and predict impending catastrophes. First, we will validate existing theory by simulating realistic networks in order to characterize conditions, e.g., topologies, queuing disciplines, loads, congestion-control regimes, and user behaviors, where phase-transitions occur, and conditions where they do not. Second, we will confirm our simulated findings in laboratory networks, and ultimately real networks. Third, we will design, develop, evaluate, and apply measurement and analysis methods in laboratory networks to warn of impending catastrophes. Fourth, we will establish collaborations with network operators to evaluate our measurement methods in real networks. Technical risk arises from two main sources: (1) phase-transition behaviors demonstrated in abstract network models may fail to appear in real networks; and (2) online measurement methods might not scale sufficiently to provide effective alerts for nationwide networks. While the first risk is inherent in the proposed research, we have a plan (see *Task 4* below) to mitigate the second risk.

Our project deliverables will provide: (1) deepened understanding of the applicability of an existing phase-transition theory to real networks; (2) practical measurement and analysis methods, and related software, to monitor real networks for impending system collapse; and (3) assessment of the utility of measurement-based methods for early warning of catastrophes in real networks. As we elaborate below, the project plan (as shown in the attached project task schedule) consists of three, eighteen-month phases, encompassing eight tasks. Major decision points exist after each of the first two phases.

**Phase I: Validation**—We will determine conditions under which the phase-transition behaviors demonstrated by theory, and related studies of abstract network models, also appear (or not) in realistic network models, and in laboratory networks. The phase requires three tasks, each lasting six months.

In *Task 1, Modeling*, we will inject realistic network characteristics from our previous work [25-26] into an abstract model taken from academe. For example, Echenique, Gomez-Gardenes and Moreno (EGM) developed an abstract model [27] based on an 11,174-node topology taken from the Internet autonomous system map, circa 2001. While the resulting topology is somewhat realistic, the EGM model lacks realism in node characteristics, as well as in source and receiver characteristics, in user behaviors, and in protocol procedures. We will replace the large EGM topology with a realistic, but smaller, more computationally tractable, topology, representing a U.S. network. We will then inject varied router speeds, finite buffers,

and propagation delays into the EGM model. We will also modify the EGM model to ensure that sources and receivers attach with bounded interface speeds, and only at access routers. We will allow sources and receivers to be distributed non-uniformly across access routers in a topology. We will model users as cyclic on-off processes that can transfer varied file sizes, and that exhibit limited patience when transfers take too long, or rates are too slow. Finally, we will inject into the EGM model the connection-establishment and congestion-control procedures associated with the transmission-control protocol (TCP), which is used for over 90% of the data transferred across the Internet.

In *Task 2, Characterization*, we will assess the sensitivity (see Chapter 4 in [26]) of our modified EGM model to the realistic factors we injected. Subsequently, we will attempt to demonstrate the phase-transition behaviors shown in nine previous research projects (see Table 2), and we will evaluate the ability of various precursor signals (e.g., changes in the slope of time series for order parameters and increasing amplitude fluctuations and long-range dependence in key metrics) to warn of impending phase transitions. We will determine if the onset of phase transitions can be predicted based on observable changes in only a subset of nodes, and, if so, how such subsets relate to network topology. At the same time, we will characterize relationships between false positives and the selection of measurement intervals and alerting thresholds. We will establish whether phase-transition behaviors seen in abstract network models also appear in realistic models, and under what circumstances. We will evaluate which precursor signals appear most effective in predicting onset of phase transitions, and under what constraints. If phase-transition behaviors do not appear in realistic network models, then this task will provide explanations, which can serve to refocus ongoing academic research. Results from this task may also inform ongoing academic research that aims at predicting phase transitions in other types of networks, such as electrical grids and cyber-physical systems.

In *Task 3, Emulation*, we will replicate phase-transitions demonstrated under Task 2, but in the context of the ITL Emulab network laboratory, which currently has 100 physical nodes that can emulate about 1500 virtual nodes. (With IE funds from this project, we intend to double the capacity of the ITL Emulab in order to match the scale of our simulations.) Demonstrating phase transitions in emulated networks will validate simulated phase transitions found in Task 2.

**Phase II: Measurement Methods**— Existing phase-transition studies either sample measurements at one node or examine measurements across the entire network. Sampling at one node is not informative in real networks, which are composed of heterogeneous components (e.g., nodes with varying topological placements, capacities, and functionalities). Measuring across an entire network is infeasible for deployments of realistic scale, because the volume of measurement data would be difficult to collect, convey, and analyze. We will design, develop, and test pragmatic measurement and analysis methods that can be inserted into operational networks to provide effective sampling and signaling. This phase requires three tasks, each lasting six months.

In *Task 4, Design*, we will explore techniques to select a limited number of measurement locations to achieve sufficient detection of precursor signals, as identified and selected from Task 2. We suspect that network topology plays a crucial role in selection of effective measurement locations. In previous work [28] with a heterogeneous topology, we found that measuring at about 10% of the possible observation points was sufficient to detect signals from a range of distributed denial of service attacks. Further, we are collaborating with Dr. Yan Wan from the University of North Texas, who is investigating cascading failures in network models, using measurement methods based on sampling concepts [29-31] taken from control engineering,

algebraic graph theory, and information theory. In the context of simulation models, we will investigate combinations of such sampling techniques, and determine their ability to detect precursor signals, as well as to quantify associated uncertainty.

In *Task 5, Development*, we will construct measurement and analysis software that can be deployed in real networks. The measurement software, which leverages existing network performance monitoring software known as perfSONAR [32], will collect requisite measurements and convey them to an analysis point. We will also construct analysis software that can process collected measurements to produce precursor signals, and after applying appropriate filters and thresholds, as determined during Task 2, generate early warnings of catastrophic events. The early warnings will include uncertainty estimates.

In *Task 6, Testing*, we will deploy our measurement and analysis software in ITL's network emulation laboratory and determine its ability to provide early warnings for phase transitions demonstrated in Task 3. This will establish the practicality and effectiveness of our measurement and analysis methods, and software.

**Phase III: Technology Transfer and Further Evaluation**—We will package and release our software for general use, and will establish partnerships to evaluate our methods and software in commercial settings. The phase requires two tasks, the first lasting six months, and the second lasting twelve.

In *Task 7, Software Release*, we will package our software into public-release form, which can be downloaded and used by any organization, and then distribute the software through perfSONAR, Source Forge [33], or similar software-distribution channels.

In *Task 8, Evaluation*, we will establish partnerships with network operators to deploy and test our measurement and analysis methods and software. Basing our software on perfSONAR will provide immediate access to an existing consortium of (mainly) government and academic organizations seeking to build interoperable network performance middleware. To expand the scope of our partnerships, we will also invite interested commercial collaborators to enter a CRADA directly with NIST. The main goals of this task are: (1) to evaluate and improve our measurement and analysis software based on deployment in real networks; and (2) to interest network operators in deploying our measurement and analysis methods in national networks.

**Potential impacts:** If completely successful, this project will transfer a promising theory into practice, thus improving the resilience of the nation's networks and systems based on those networks, and ultimately reducing costs to the U.S. economy arising from widespread network outages and degradations. Further, the knowledge and measurement and analysis methods, developed by this project may find application in additional systems of national importance, systems such as computational clouds, the electrical grid and cyber-physical systems. For NIST, and particularly ITL, the project will open new research vistas into measurement and analysis methods aimed at characterizing behavior and dynamics in large-scale networks. Such measurement and analysis methods will be among the most important contributions that NIST can make to the future of measurement science for information technology. Even at its least successful, the project will characterize and explain conditions under which phase transitions will fail to materialize in real networks, thus raising the nation's confidence in existing networked systems.

Each phase of the project will generate unique scientific contributions that have independent value beyond the cumulative goal of the project. For example, independent assessment of the theory of network phase transitions will either reinforce existing directions, or refocus ongoing research into alternate paths.

**Qualifications of research team:** The multidisciplinary team comprises NIST's leading experts in modeling of network-based systems, statistical analysis and experiment design, time series analysis, and mathematical analysis. In addition, over 80% of the IMS labor funds will be invested in significantly expanding ITL's expertise in network measurement and data analysis. Through recruiting a new FTE and two new postdocs with proven research records in the design and analysis of systems for network measurement, we will position NIST to become a technical leader in subsequent research in these areas.

**Kevin Mills** (772.04) ST-1550 Senior Computer Scientist—Dr. Mills has 30+ years of experience in developing Internet models, encompassing network structure and characteristics, transport layer protocols, and application behavior, and applying those models to study congestion behavior.

**Chris Dabrowski** (772.04) ZP/1550/IV Computer Scientist—Mr. Dabrowski has created Markov chain models and simulations representing large distributed systems, including infrastructure clouds, computational grids, and communication networks. He recently developed methods to predict causes and expected patterns of performance degradation in large distributed systems.

**Jim Filliben** (776.04) ZP/1529/V Supervisory Mathematical Statistician—Dr. Filliben has 40+ years of experience in leading NIST research and application of advanced statistical modeling and analysis techniques, including five years focused on experiment design for communications networks.

**Fern Hunt** (771.01) ZP/1529/V Mathematician—Dr. Hunt has developed proofs of mathematical results in areas such as dynamical systems, probability, information theory, and complex systems and has also built models of failure in networks. She is currently working on near-optimal sets for the spread of information in networks.

**Bert Rust** (771.01) ZP/1529/V Mathematician—Dr. Rust is a world-class expert in developing mathematical models to characterize time series data, and has applied that expertise to model data ranging from climate change measurements to traffic measurements from the Internet.

**New Hire FTE** (772) ZP/IV Computer Scientist / Mathematician – Hire or contract recent Ph.D. with direct research experience in the measurement and analysis of network behavior.

**New Postdocs (3)** – Statistician / Mathematician / Computer Scientist - Recruit three new postdocs with direct research experience in the design and evaluation of network measurement software and systems, and data analysis.

**New Guest Researcher** – Recruit new guest researcher with experience in network data collection and analysis.

**Resources required:**

Budget Resources (\$K)										
	STRS					Invested Equipment (IE)				
Div/Grp	FY15	FY16	FY17	FY18	FY19	FY15	FY16	FY17	FY18	FY19
772/04	\$810	\$810	\$810	\$810	\$810	\$425	\$375			
776/04	\$230	\$230	\$230	\$230	\$230					
771/01	\$172	\$172	\$172	\$172	\$172					
Totals	\$1,212*	\$1,212	\$1,212	\$1,212	\$1,212	\$425	\$375			

\*83% of labor spending for one new FTE, four new Post-Docs and one new Guest Researcher

Staffing Resources			
Div/Grp	NIST Employee Names	# of NIST FTEs	# of NIST Associates
772/04	New FTE, New Postdocs, New Guest Researcher	1.0	3.0
772/04	Kevin Mills and Chris Dabrowski (each 25% IMS funded)	0.50	
776/04	Jim Filliben	0.25	
776/04	New Postdoc		1.0
771/01	Fern Hunt and Bert Rust (each 25% IMS funded)	0.50	
Totals		2.25	4.0

Invested Equipment (IE) Planned Purchases *	
	Equipment Description & Estimated Cost
FY2015	<ul style="list-style-type: none"> <li>\$375K – 50 additional Emulab nodes + software licenses</li> <li>\$50K – 6 additional network switches + cables</li> </ul>
FY2016	<ul style="list-style-type: none"> <li>\$375K – 50 additional Emulab nodes + software licenses</li> </ul>

\*Totals referenced must tie to requested amounts in "Budget Resources" table.

**Project Task Schedule:**

Phase	ID	Task	Year	Qtr	2015				2016				2017				2018				2019			
					4	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3				
Validation	1	Modeling			■	■																		
	2	Characterization				■	■																	
	3	Emulation						■	■															
Measurement Methods	4	Design							■	■														
	5	Development									■	■												
	6	Testing											■	■										
Technology Transfer	7	Software Release												■	■									
	8	Evaluation														■	■	■	■					



## References:

1. S. Floyd and V. Jacobson. The Synchronization of Periodic Routing Messages. *IEEE/ACM Transactions on Networking*, (2) 2, 122-136, 1994.
2. C-F Liew. Taiwan Earthquakes Interrupts Telecommunications Across Asia. *Tech News*, Dec. 27, 2006. <http://www.liewcf.com/taiwan-earthquakes-interrupts-telecommunications-across-asia-2718/>
3. J. Duffy. Cisco routers caused major outage in Japan: report. *Network World*, May 16, 2007. <http://www.networkworld.com/news/2007/051607-cisco-routers-major-outage-japan.html>
4. A. Brones. Internet failure hits several continents. *Gadling*, January 31, 2008. <http://www.gadling.com/2008/01/31/internet-failure-hits-several-continents/>
5. L. Whitney. Amazon EC2 cloud service hit by botnet, outage. *CNET*, December 11, 2009. [http://news.cnet.com/8301-1009\\_3-10413951-83.html](http://news.cnet.com/8301-1009_3-10413951-83.html)
6. J. Reed. 15 Percent of All Internet Traffic Secretly Rerouted Through China. *Defense Tech*, November 16, 2010. <http://defensetech.org/2010/11/16/15-percent-of-all-internet-traffic-secretly-rerouted-through-china/>
7. C. Miller. Amazon Cloud Failure Takes Down Web Sites. *New York Times Bits Blog*, April 21, 2011. [http://bits.blogs.nytimes.com/2011/04/21/amazon-cloud-failure-takes-down-web-sites/?\\_php=true&\\_type=blogs&\\_php=true&\\_type=blogs&\\_r=1](http://bits.blogs.nytimes.com/2011/04/21/amazon-cloud-failure-takes-down-web-sites/?_php=true&_type=blogs&_php=true&_type=blogs&_r=1)
8. C. Shu. Updated: Netflix Crippled On Christmas Even By AWS Outages. *Tech Crunch*, December 24, 2012. <http://techcrunch.com/2012/12/24/netflix-crippled-on-christmas-eve-by-aws-outages/>
9. D. Fitzgerald. East Coast Internet Traffic is Restored. *Wall Street Journal*, Oct. 20, 2013. <http://online.wsj.com/news/articles/SB10001424052702304864504579145813698584246>
10. S. Madrak. Massive Internet Outage Today. *Crooks and Liars*, Jan. 27, 2014. <http://crooksandliars.com/2014/01/massive-internet-outage-today>
11. R. Solé and S. Valverde. Information transfer and phase transitions in a model of internet traffic. *Physica A*, Vol. 289, pp. 295-605, 2001.
12. M. Woolf, D. Arrowsmith, R. Mondragon, and J. Pitts. Optimization and phase transitions in a chaotic model of data traffic. *Phys Rev E*, Vol. 66. 046106, 2002.
13. D. Arrowsmith, R. Mondragon, J. Pitts, and M. Woolf. Phase Transitions in Packet Traffic on Regular Networks a comparison of source types and topologies. Report 08, Institut Mittag-Leffler, 2004.
14. G. Mukherjee and S. Manna. Phase transition in a directed traffic flow network. *Phys Rev E*, (71) 6, 2005.
15. A. Lawniczak, P. Lio, S. Xie, and J. Xu. Study of Packet Traffic Fluctuations Near Phase Transition Point From Free Flow to Congestion in Data Network Model. *Proceedings of Canadian Conference on Electrical and Computer Engineering*, 360-363, April 22-26, 2007.

16. B. Tadic, G. Rodgers, and S. Thurner. Transport on Complex Networks: Flow, Jamming and Optimization. *International Journal of Bifurcation and Chaos*. (17) 7, 2363-2385, 2007.
17. S. Sarkar, K. Mukherjee, A. Srivastav and A. Ray. Understanding phase transition in communication networks to enable robust and resilient control. *Proceedings of the American Control Conference*, 1549-1554, June 10-12, 2009.
18. D. Wang, N. Cai, Y. Jing, and S. Zhang. Phase Transition in Complex Networks. *Proceedings of the American Control Conference*, pp. 3310-3313, June 10-12, 2009
19. Y. Rykalo, L. Levitan, R. Browe. Critical phenomena in discrete-time interconnection networks. *Physica A*, Vol. 389, 5259-5278, 2010.
20. I. Sainee. Appendix A. Letter of support from Iraj Sainee, Head of the Mathematics of Networks & Communications Department of Bell Labs.
21. H. Lambert. Appendix B. Letter of support from David Lambert, President and CEO of Internet 2.
22. C. Lee. Appendix C. Letter of support from Craig Lee, Senior Scientist, The Aerospace Corporation.
23. J. Wilson, "The cost of network downtime 2003," Infonetics Research, 2003.
24. J. Wilson, "The cost of network downtime 2005," Infonetics Research, 2005.
25. K. Mills, E. Schwartz and J. Yuan, "How to Model a TCP/IP Network using only 20 Parameters," *Proceedings of the 2010 Winter Simulation Conference (WSC 2010)*, Dec. 5-8, Baltimore, MD, 2010.
26. K. Mills, J. Filliben, D. Cho, E. Schwartz and D. Genin, Study of Proposed Internet Congestion Control Algorithms, NIST Special Publication 500-282, May 2010, 534 pages.
27. P. Echenique, J. Gomez-Gardenes, and Y. Moreno, "Dynamics of Jamming Transitions in Complex Networks," *Europhysics Letters*, 71, 325, 2005.
28. J. Yuan and K. Mills, "Monitoring the Macroscopic Effect of DDoS Flooding Attacks," *IEEE Transactions on Dependable and Secure Computing*, 2/4, 324-335, 2005.
29. A. Rai, M. Xue, E. Yeung, S. Roy, Y. Wan, A. Saberi, and S. Warnick, "Initial-condition estimation in network synchronization processes: graphical characterizations of estimator structure and performance," *Complex Systems*, 21/4, 2013.
30. Y. Wan and S. Roy, "On inference of network time constants from impulse response data: graph-theoretic Cramer-Rao bounds," *Proceedings of 48th IEEE Conference on Decision and Control*, Shanghai, P.R. China, pp. 4111-4116, Dec. 16-18, 2009.
31. G. Movva, Y. Wan, S. Fu, and H. Wu, "Optimal sensor placement for structural health monitoring: a comparative study between the control engineering and civil engineering approaches," *SPIE NDE Smart Structure*, San Diego, California, March 2013.
32. B. Tierney et al., Instantiating a Global Network Measurement Framework, LBNL Technical Report LBNL-1452E, January 2009.
33. <http://sourceforge.net>

**Appendix A. Letter of Support from Iraj Saniee,  
Head of the Mathematics of Networks & Communications Department of  
Bell Labs**

-----Original Message-----

From: Saniee, Iraj (Iraj) [mailto:iis@research.bell-labs.com]  
Sent: Thursday, December 20, 2012 7:27 PM  
To: Dabrowski, Christopher E.; Mills, Kevin L.  
Subject: Your proposal

Dear Drs. Dabrowski and Mills,

Thank you very much for sharing with me the abstract of your proposal, which I enjoyed reading.

The study of critical phenomena and methods for predicting the onset of "turbulence" has a very respectable history, originally in physics and statistics. Kolmogorov's classic paper on this topic is a very nice early example. The distributed nature of today's information systems and infra-structure, as you observe, adds to the potential criticality of this complex system, making prediction of large-scale systemic failures even harder to categorize or detect.

Measurements, and in particular parsimonious measurements if and when proven effective in facilitating prediction of large-scale failures, are the missing elixir. There is currently no network management system I'm aware of which can correlate disparate events (across network layers and components) which may give rise to catastrophic failures of the information networks. This is despite the fact that the Internet was designed to be inherently survivable with respect to multiple and distributed physical interruptions.

Again, as you observe in your abstract, cyber-attacks and other higher-layer protocol failures could still contribute to large-scale failures of the information infra-structure, in parallel with environmental and electrical systems, for example. Your emphasis on measurement and your track record of research on distributed information systems focus this complex problem on a very plausible solution methodology. In my opinion, the success of the proposed research would help fill a vacuum in commercial network control and management systems which are primarily designed to serve myopic owner-client needs rather than global infrastructure availability concerns.

I wish you success in your proposed research and I look forward to your results if and when available to the general public.

Best wishes,

Iraj Saniee, PhD  
Head, Mathematics of Networks & Communications Dept.  
Bell Labs - Research, Alcatel-Lucent  
Room 2C-326, 600 Mountain Avenue  
Murray Hill, NJ 07974  
U S A  
(908) 582 6410  
<http://ect.bell-labs.com/who/iis/>

**Appendix B. Letter of Support from H. David Lambert,  
President and CEO, Internet2**



January 11, 2013

Christopher Dabrowski  
National Institute of Standards  
And Technology  
U.S. Department of Commerce  
100 Bureau Drive, Stop 8920  
Gaithersburg, MD 20899-8920

Dear Mr. Dabrowski:

On behalf of Internet2, we are pleased to write in support of your proposal entitled "Measurement Science to Predict Catastrophic Events in Global Communication Networks."

This proposal is likely to contribute to our ability to manage dynamic global communication networks, particularly in the accurate prediction of system-wide catastrophic events, and the avoidance of major network collapse. The project's findings will address processes involved in malicious network attacks, cascading failures, and sudden load increases, while also enhancing our ability to predict network failure onset. By developing large-scale models of Internet topologies and protocols, it will measure precursor data as signals of pending catastrophes and build measurable precursor behaviors. This will create a strong foundation of system measurement that has not existed before that is likely to help avoid potentially debilitating real-life network failures and their scientific and economic consequences. In the process, it will also contribute significantly to global technological and economic development.

Internet2 has a long history of successful impact in network performance and network security issues, including many collaborations with members of this proposal team, including Abdella Battou. Internet2 is a consortium of advanced networking institutions in the United States that has been led and governed by the research and education community (including over 200 university members) since 1996. Internet2 promotes the missions of its members by providing pioneering network capabilities and unique partnership opportunities that facilitate the development, deployment, and use of advanced networking capabilities.

We are pleased to write in support of this innovative proposal.

Sincerely,

A handwritten signature in black ink, appearing to read "H. David Lambert". The signature is fluid and cursive, with a long horizontal stroke at the end.

H. David Lambert  
President and CEO, Internet2

Office of the President & CEO  
1150 18<sup>th</sup> Street, NW, Suite 900  
Washington, DC 20036

**Appendix C. Letter of Support from Craig Lee,  
Senior Scientist, The Aerospace Corporation**



The Aerospace Corporation  
2310 E. El Segundo Blvd.  
El Segundo, CA 90245-4609  
310.336.5000

MAILING ADDRESS  
P.O. Box 92957  
Los Angeles, CA 90009-2957  
[www.aero.org](http://www.aero.org)

December 28, 2012

National Institute of Standards and Technology  
100 Bureau Drive Stop 8920  
Bldg. 222/B218  
Gaithersburg, MD 20899-8920

Dear Drs. Dabrowski and Mills,

It is with great interest that I support the "Measurement Science To Predict Catastrophic Events In Global Communication Networks" proposal. It is both timely and critical as we deploy -- and rely upon -- larger and larger distributed computing infrastructures for our national infrastructures.

In my professional capacity as a Senior Scientist at The Aerospace Corporation, in the field of parallel and distributed computing, and as a Past President of the Open Grid Forum, I directly advise on issues relating to technology adoption for satellite ground systems. This is for both command and control, and also for the dissemination of data products once they are on the ground. As you may imagine, these systems are becoming larger and more integrated. There is also a great push for migration to on-demand, distributed cloud architectures that must host multiple satellite missions, while providing greater business value.

Our national security and economic interests depend on these systems being stable and robust. Two key examples are the Global Positioning System and the Joint Polar (weather) Satellite System. However, as these systems get larger, more highly distributed, more inter-related, and cloud-based, my concern is that their reliability and availability will suffer. Perhaps more importantly, there might be emergent behaviors that are completely unexpected and catastrophic in nature.

For these reasons, I have been following your work for several years now. I strongly support your efforts to understand how these massive systems might fail and to prototype a measurement science as perhaps an "early warning system". This line of work must be pursued, and its results used to shape satellite ground systems of the future.

Best regards,

A handwritten signature in black ink, appearing to read 'Craig A. Lee'.

Dr. Craig A. Lee  
Senior Scientist, The Aerospace Corporation  
Board of Directors, Past President, Open Grid Forum ]