Early Warning of Network Catastrophes

K. Mills, C. Dabrowski, J. Filliben, F. Hunt, and B. Rust

Information Technology Laboratory

August 28, 2014

Proposal in a Nutshell

Goal: Provide early warning of looming network catastrophes, so operators can take remedial actions

Approach: *Transform theory into practical measurement method, creating new capability to monitor and manage networks*

Impact:

- ✓ Improve resilience of nation's networks
- ✓ Reduce cost of widespread outages and degradations
- ✓ Transfer methods to other systems of national importance, e.g., computational clouds, cyber-physical systems, smart power grids

Argument in Brief

<u>Problem/Opportunity</u>

- Cost of outages/degradations in networks large and growing
- Network outages/degradations spread in space and time
- Spreading process modeled as percolation from statistical physics
- Near a critical point process exhibits measureable precursor signals
- To date, precursor signals shown in abstract network models only

<u>We propose to</u>

I. Validate precursor signals in realistic network models
II. Design, develop, and test measurement method and related monitoring & analysis software deployable in real networks
III. Evaluate monitoring & analysis software with partners

Network Outages and Degradations: Crippling, Costly, and Continuing

Across 5000 US companies the cost is about $85B/yr

Infonetics study of network downtime costs in **six selected companies**

Company/ 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%

- **Dec 2006**: Disaster On the Net: Internet Outages Caused by Taiwan Earthquakes
- **May 2007**: Cisco Routers Caused Major Outage in Japan
- **Jan 2008**: Internet Failure Hits Several Continents
- **Dec 2009**: Amazon EC2 Cloud Service Hit by Botnet Outage
- **Nov 2010**: 15% of All Internet Traffic Secretly Rerouted Through China
- **Apr 2011**: Amazon Cloud Failure Takes Down Web Sites
- **Dec 2012**: Netflix Crippled by AWS Outages on Christmas Eve
- **Feb 2013**: Food Stamp System in 17 States affected by Outage
- **May 2014**: Pentagon Police Hit By 'Catastrophic' Network Outage

Network outages and degradations can last for hours, days, or weeks

Costs will Grow as Network Dependence Increases due to Cloud Computing

EC2 OUTAGE REACTIONS SHOWCASE WIDESPREAD IGNORANCE REGARDING THE CLOUD Apr 22, 2011

Amazon EC2 Outage Explained and Lessons Learned Apr 29, 2011

BUSINESS

Microsoft's Azure Cloud Suffers Serious Outage Feb 29, 2012

(Real) Storm Crushes Amazon Cloud, Knocks out Netflix, Pinterest, Instagram Jun 30, 2012

Storms, leap second trigger weekend of outages Jul 2, 2012

How did Amazon have a cloud service outage that was caused by generator failure? Jul 3, 2012

AWS outages, bugs and bottlenecks explained by Amazon Jul 3, 2012

Never-before-seen software bug caused flood of requests creating a massive backlog in the system

Salesforce.com hit with second major outage in two weeks Jul 10, 2012

SalesForce outages show SaaS customers dependence on providers' DR plans Jul 12, 2012

What's happened to the cloud? Jul 13, 2012

Are major cloud outages in recent times denting confidence?

Google Talk, Twitter, Azure Outages: Bad Cloud Day Jul 26, 2012

Sampling of Cloud Outages over 15 Months: Apr 2011-July 2012

Networks: Graphical Topologies Spanning Space and Exhibiting Time-Varying Dynamics

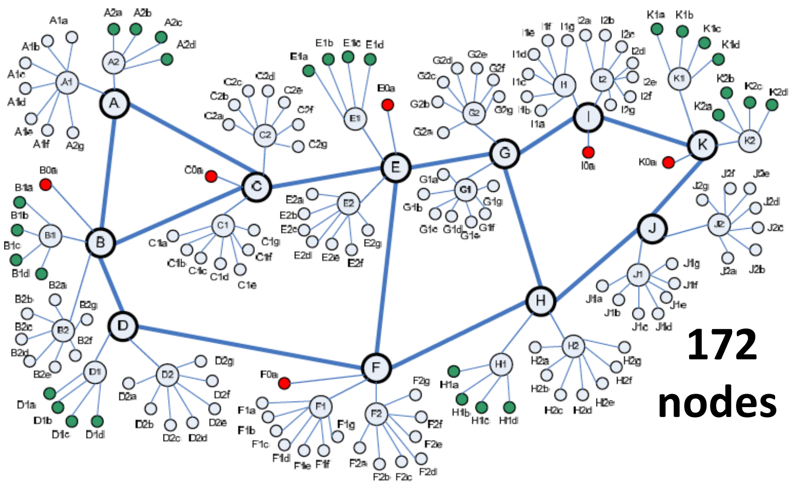

172 nodes

US *Internet2* Academic Network

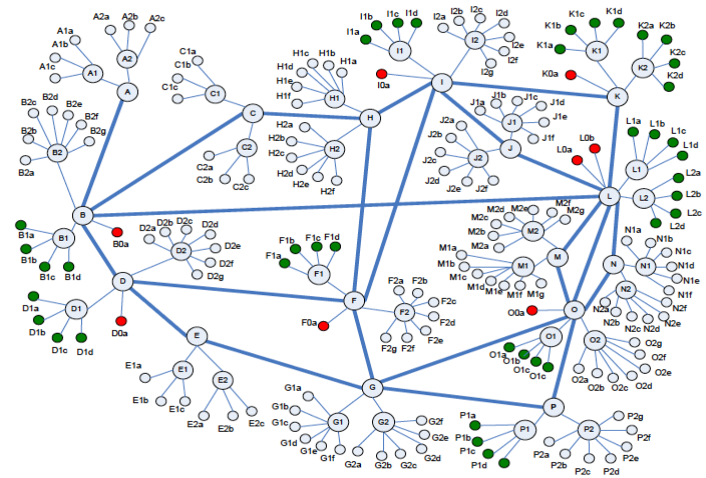

218 nodes

Commercial US Internet Service Provider Network

11,174 nodes

circa 2001

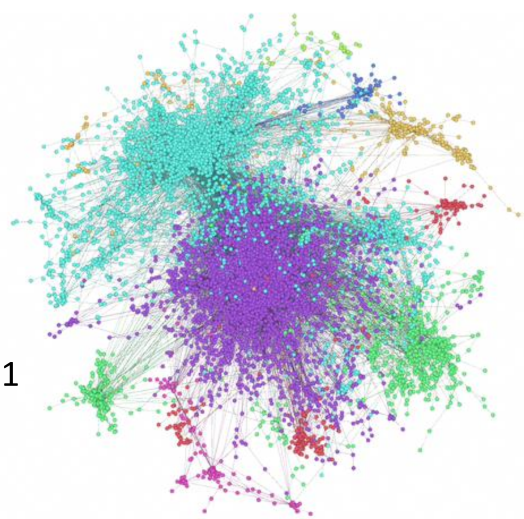

Global Internet Autonomous System Network

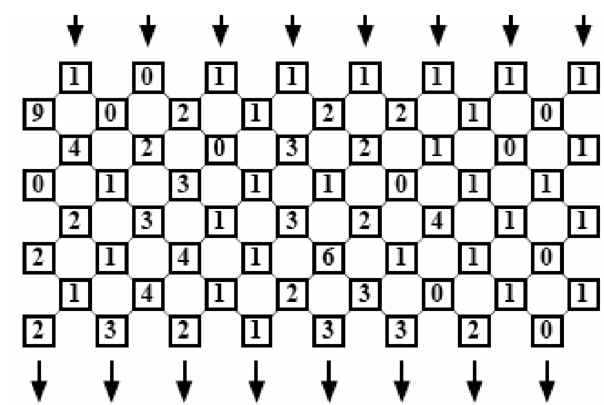

64 nodes

2D Lattice Network

Like Physical Systems, Networks Exhibit: Spatial Structure, Microscopic Behavior, and Macroscopic Spatiotemporal Dynamics

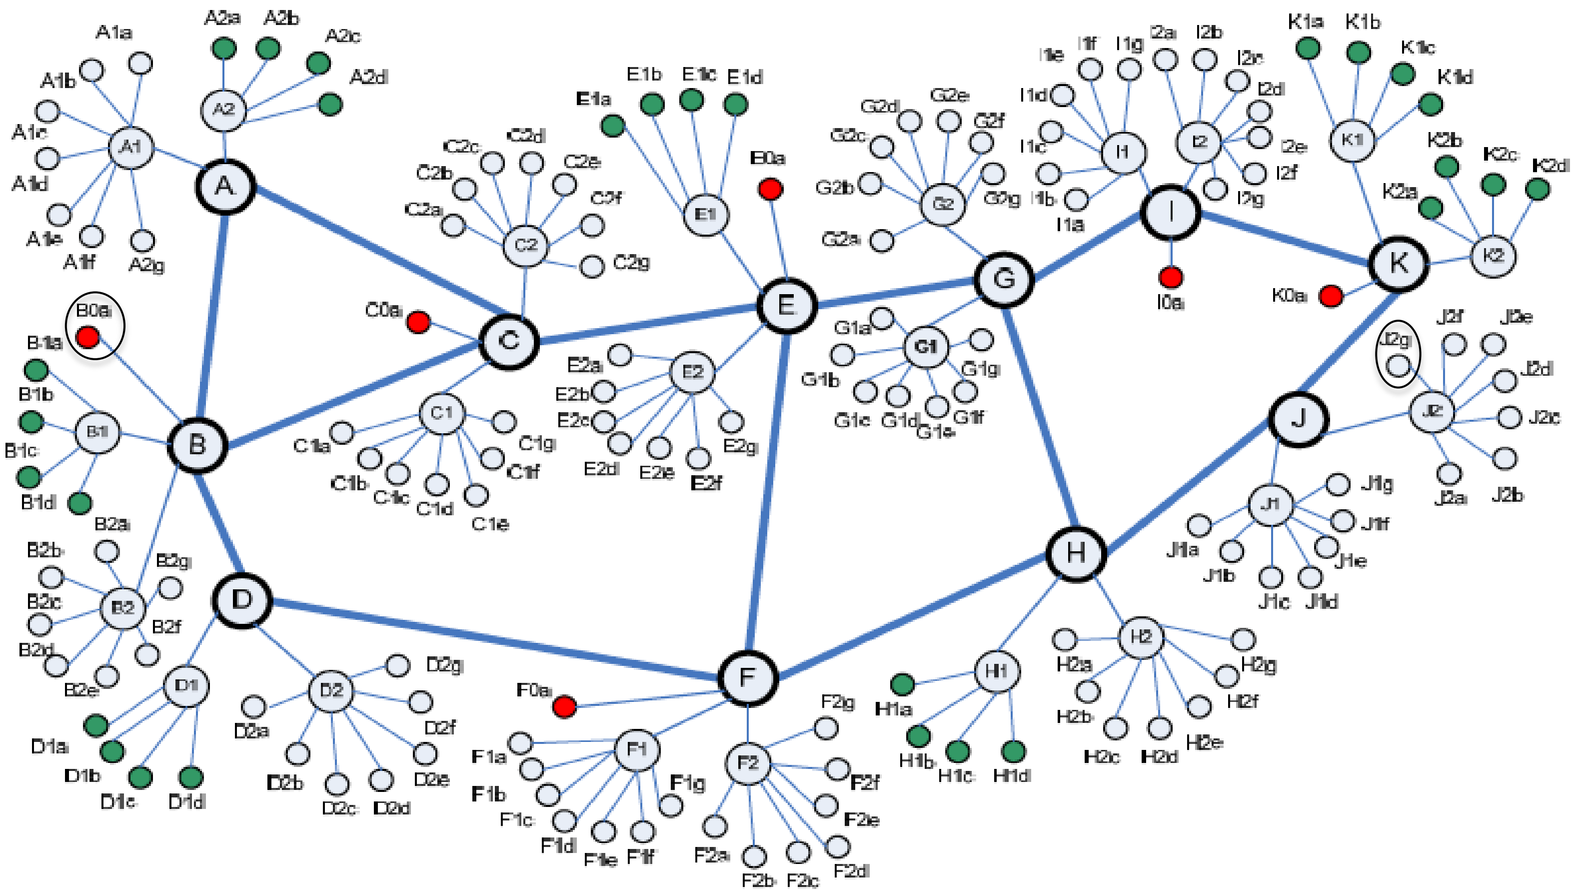

Explanatory Example: Data Stream between NIST and Google

Outages and Degradations Spread across Networks in Space and Time

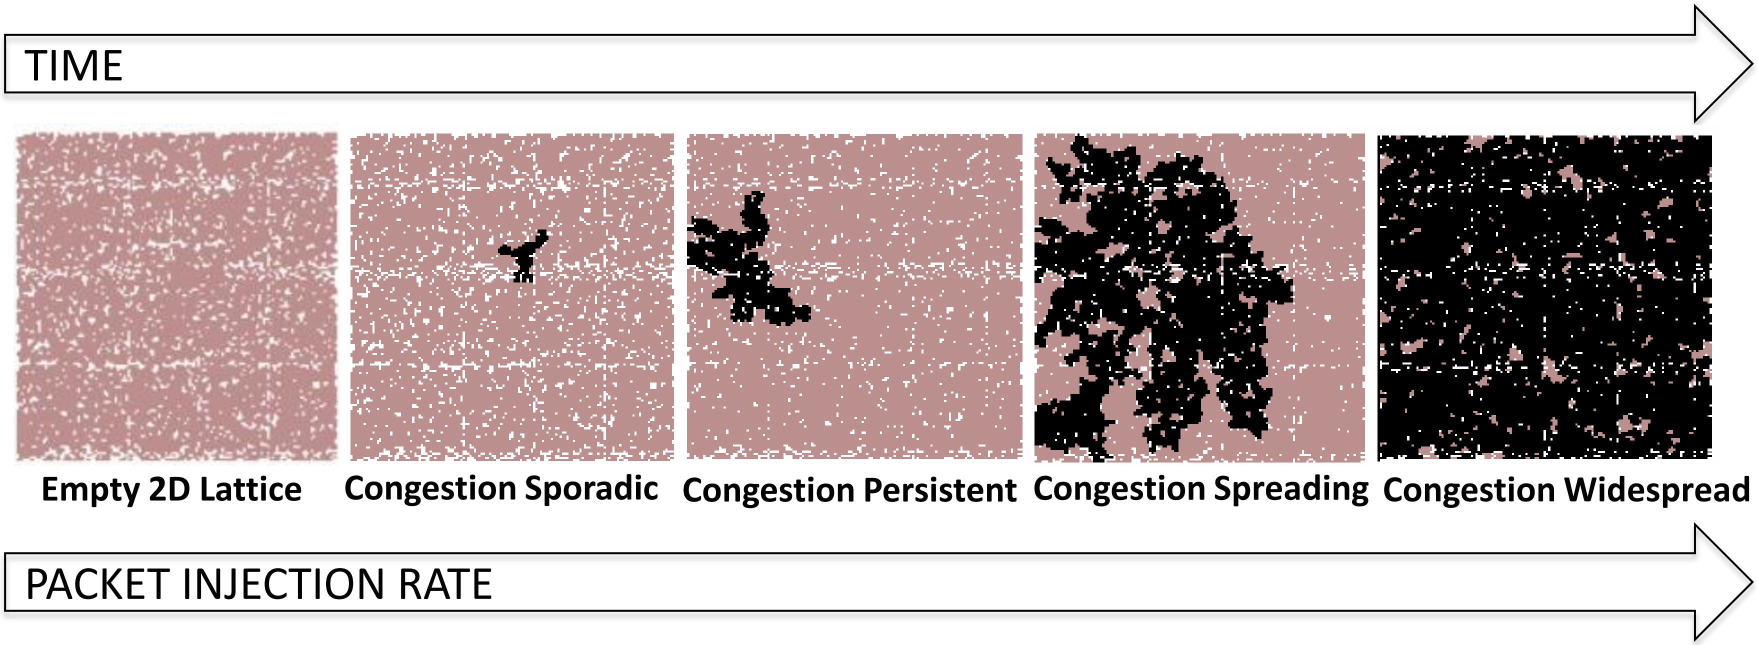

IMPLICATION: SHOULD BE POSSIBLE TO DETECT THE SPREADING PROCESS AND PROVIDE EARLY WARNING OF INCIPIENT CATASTROPHE

2D lattice animation taken from "Percolation Theory Lecture Notes", Dr. Kim Christensen, Imperial College London, October 9, 2002

Spreading Process often Modeled as Percolation from Statistical Physics

Percolation → spread of some property in a lattice (or graph) leading to the formation of a *giant connected component* (GCC), as measured by P_∞, the proportion of nodes encompassed by the GCC

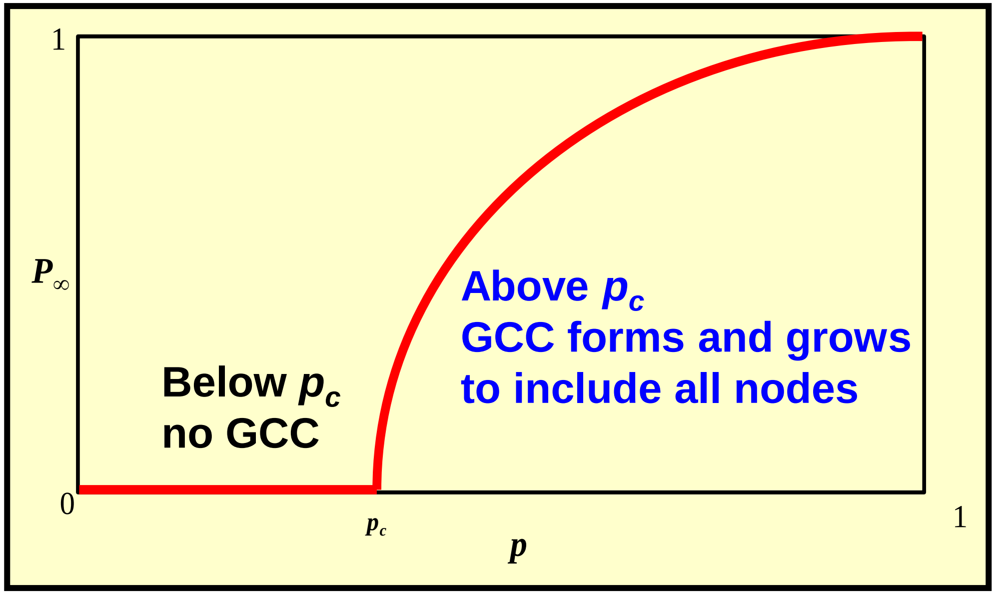

p is probability a node has property

p_c is known as the critical point

$p < p_c$ → no spread

$p = p_c$ → *percolation phase transition*

$p > p_c$ → spread occurs, and expands with increasing *p*

Near a critical point, the process exhibits precursor signals, typically attributable to increasing, systemic correlation

Academics Model Spreading Network Congestion as a Percolation Process

Year	Researchers	Location	Topology	Metrics	Precursor Signal
2001	**Sole & Valverde**	**Spain & USA (SFI)**	**2D Lattice**	**Packet Delay, Queue Length, Throughput**	**Self-similarity in log-log plot of power vs. freq.**
2002	**Woolf et al.**	**UK**	**2D Lattice**	**Packet Delay, Queue Length, Throughput**	**Long-Range Dependence (LRD) in time-series autocorrelation**
2004	**Arrowsmith et al.**	**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**	**India**	**2D Lattice**	**Packet Delay, Queue Length, Load per Node**	**Self-similarity in log-log plot of power vs. freq.**
2007	**Lawniczak et al.**	**Canada**	**2D Lattice**	**Packets in Flight**	**LRD shown with Hurst parameter increases from R-S analysis**
2007	**Tadic et al.**	**Slovenia, Austria, UK**	**Generated SF & UH**	**Packet Delay, Queue Length, Network Load**	**Systemic changes in network-load time series**
2009	**Sarkar et al.**	**USA**	**2D Lattice**	**Packet Delay, Queue Length**	**Order parameter becomes positive**
2009	**Wang et al.**	**China**	**Generated ER, WS, HK**	**Packets in Flight/Injected**	**Order parameter becomes positive**
2010	**Rykalova et al.**	**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

Precursor Signals Appear in Communication Network Models

Penn State Researchers (Sarkar, Mukherjee, Srivastav, and Ray 2009) find increasing transit delay as $p > p_c = 0.3$

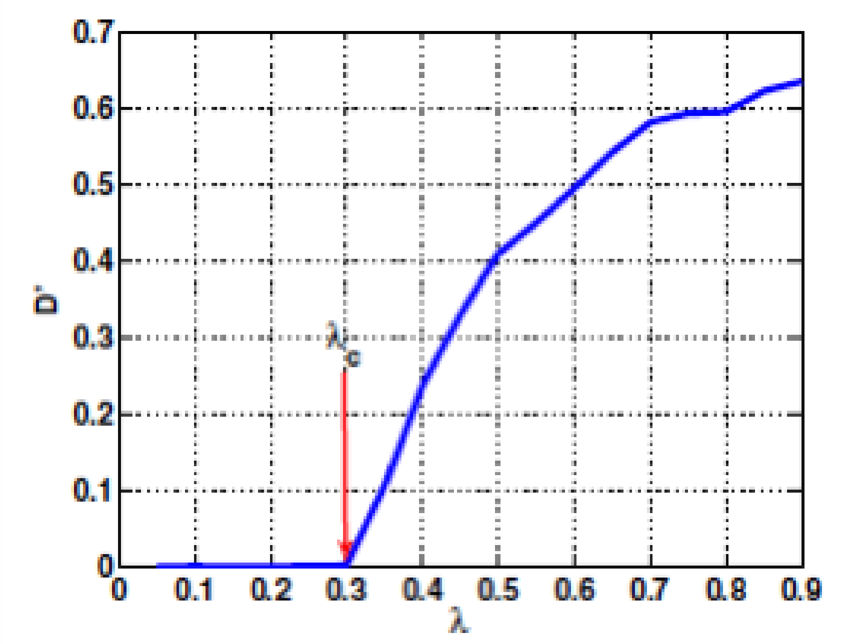

Aggregate Avg. Transit Delay (D*) vs. Network Load (λ)

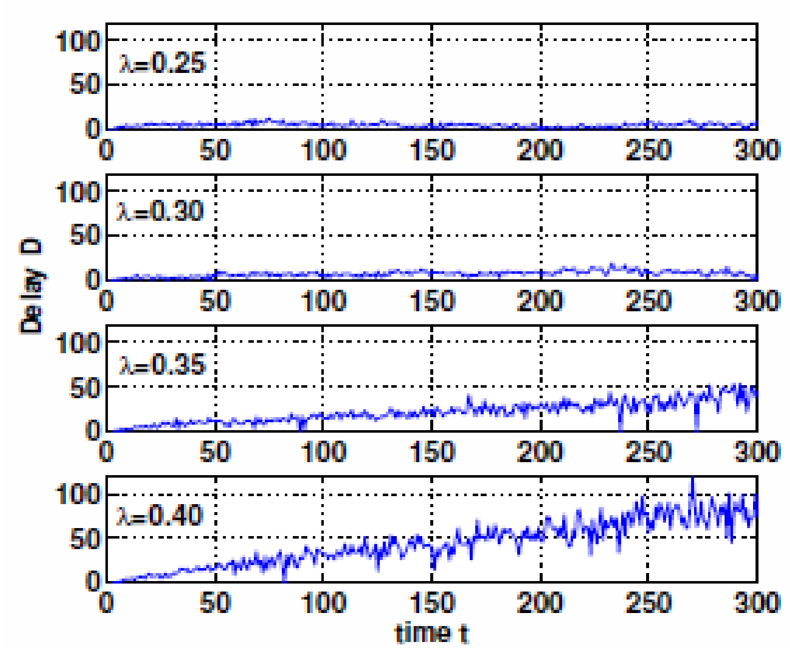

Sampled Avg. Transit Delays (D) for Four Network Loads (λ)

Increasing slope in time series of selected measured variables could signal crossing a critical point, allowing network managers to be alerted prior to network collapse

Boston University Researchers (Rykalova, Levitan, and Browe 2010) find increased correlation in time series of packets in transit as p nears $p_c = 0.2$

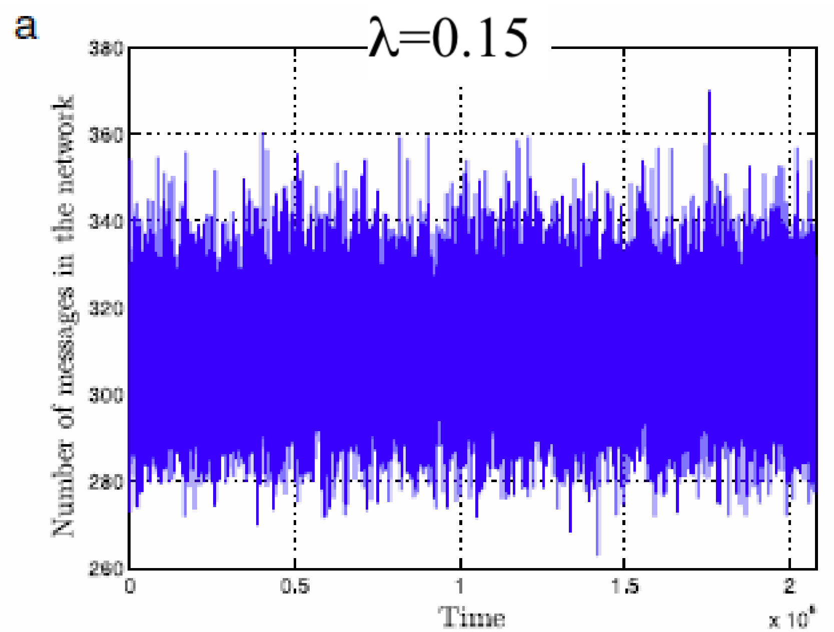

Packets fluctuate between 260 and 380

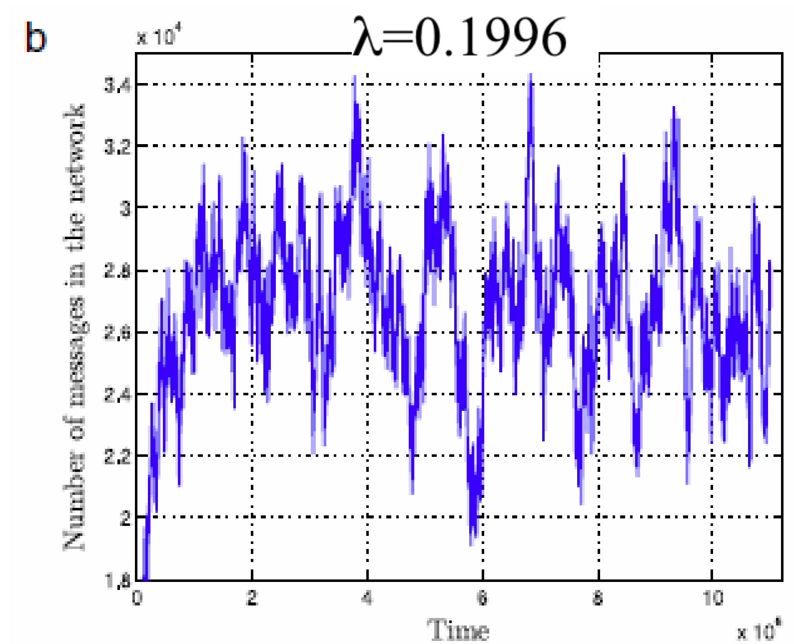

Packets fluctuate between 18,000 and 34,000

Increasing autocorrelation in time series could signal an approaching critical point, allowing network managers to be warned prior to network collapse

Why hasn't theory been transformed into practical measurement method?

<u>Theoretical Models Lack 10 Key Characteristics of Real Networks</u>

Network Factors

1. Human-engineered, tiered topologies
2. Router buffer sizes finite
3. Router speeds varied to meet demands, limit losses

Attachment Factors

4. Injection from sources and receivers only at lowest tier
5. Distribution of sources and receivers non-uniform
6. Connection of sources/receivers with few varied speeds

User Factors

7. Duty cycle of sources exhibits cyclic behavior
8. Human sources exhibit limited patience
9. Sources transfer flows of various sizes

Protocol Factor

10. Flows use the Transmission Control Protocol (TCP) to modulate injection rate based on measured congestion

Thus the applicability of the theory is unclear to industry

Project Aim & Plan

Aim: *Transform theory into practical measurement method, creating new capability to monitor and manage networks*

Phase I – Validate & characterize existing theory

Phase II – Develop & test measurement method

Phase III – Transfer technology & conduct further evaluation

Phase I – Validate & characterize existing theory

- **Task 1**: Validate theory in realistic simulations
  - How does realism influence percolation?
  - Which precursor signals are most effective?
- **Task 2**: Explore simulated measurement methods
  - Where to measure?
  - How do choice of measurement interval and alerting threshold influence false positives/negatives?
- **Task 3**: Verify findings with emulation (>> realism)
  - Deploy model scenarios in Emulab (http://emulab.antd.nist.gov)
  - Do simulated findings hold under emulation?

Phase II – Develop & test measurement method

- **Task 4**: Design measurement method
  - What proportion of nodes must be monitored?
  - What monitoring points most effective?
  - What measurement intervals and thresholds best?
- **Task 5**: Develop measurement & analysis software
  - Leverage existing perfSONAR platform
  - Construct analysis & alerting software
- **Task 6**: Test software using emulation
  - Deploy in Emulab
  - Test early warning for selected scenarios

Phase III – Technology transfer & further evaluation

- **Task 7**: Release measurement & analysis software
  - Package for public distribution
  - Distribute via perfSONAR and/or similar channels
- **Task 8**: Evaluate software in real deployments
  - Identify and provide any needed improvements
  - Leverage perfSONAR consortium
  - Expand into commercial partnerships

What Impact?

- **If completely successful**: transfer theory to practice
  - Improve resilience of nation's networks
  - Reduce cost of widespread outages and degradations
  - Transfer methods to other systems of national importance, e.g., computational clouds, cyber-physical systems, smart power grids
- **If minimally successful**: explain limits of theory
  - Increase confidence in the nation's networks
  - Stimulate new directions for ongoing academic research

Who Would Benefit?

Network Operators

David Lambert, *Internet2*: "…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."

Equipment Vendors

Iraj Saniee, Bell Labs: "…the proposed research would help fill a vacuum in commercial network control and management systems…"

Network Users

Craig Lee, Aerospace: "This line of work must be pursued, and its results used to shape … systems of the future."

Why NIST?

- National need (***Presidential Policy Directive 21* – Feb. 12, 2013**)
  - Network outages costly, crippling, and continuing
  - Growing national dependence on networks
- Incentives not available in academe or industry
  - ***Focus*** on measurement problems of national importance
  - ***Commitment*** to sufficient time horizon
  - ***Flexibility*** to produce diverse outputs as required
- Well positioned, qualified team
  - ***Deep knowledge*** modeling realistic networks
  - ***Decade of experience*** studying Internet congestion
- State-of-the-art laboratory facilities
  - Top-notch network simulation test bed
  - In-house Emulab mirrors top academic equivalents
- Ongoing industry contacts
  - Network operators
  - Equipment vendors
  - perfSONAR network measurement community

Research Team Qualifications

- **Kevin Mills** (PhD senior computer scientist)—experience developing realistic Internet models to study congestion behavior, and successfully leading this team in previous, related research (http://www.nist.gov/itl/antd/emergent_behavior.cfm)
- **Chris Dabrowski** (MS computer scientist)—experience developing methods to predict causes and patterns of performance degradation in networks
- **Jim Filliben** (PhD mathematical statistician)—leading expert in advanced statistical modeling and analysis techniques, including five years focused on network congestion
- **Fern Hunt** (PhD mathematician)—experience in mathematical modeling of dynamical systems, complex systems, and models of network failures
- **Bert Rust** (PhD mathematician)—world-class expert in developing mathematical models to characterize time series data, ranging from climate change measurements to traffic measurements from the Internet

Proposed Schedule

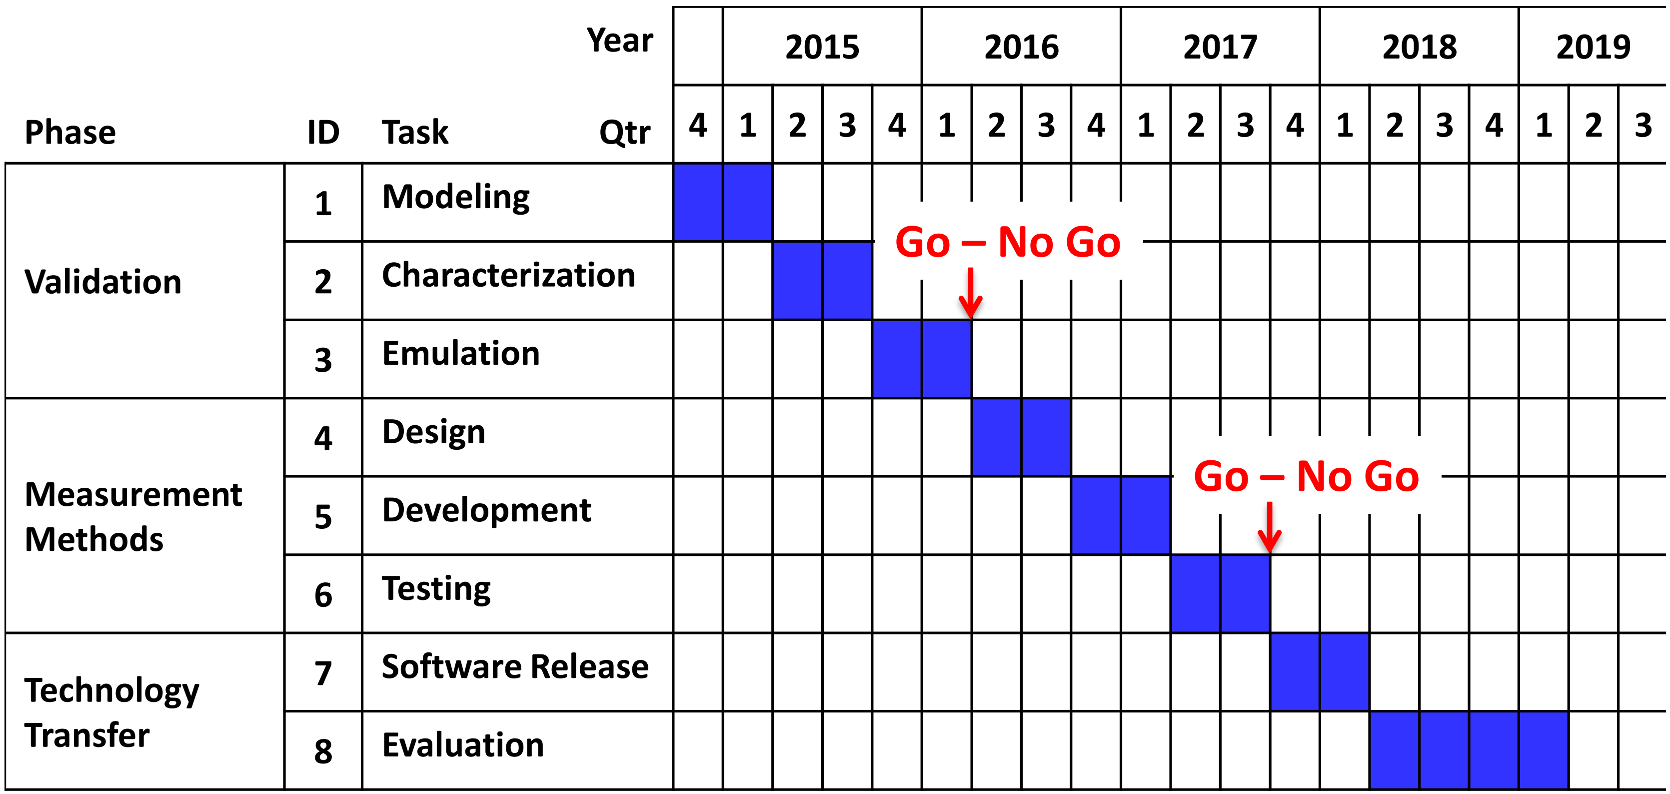

Total Resource Requirements

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	$405	$425	$375			
776/04	$230	$230	$230	$230	$115					
771/01	$172	$172	$172	$172	$136					
Totals	$1,212*	$1,212	$1,212	$1,212	$656	$425	$375			

?

***80% of labor spending for one new FTE, three new Post-Docs, and one new Guest Researcher**

Staffing & IE Budgeting Details

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	– $375K – 50 additional Emulab nodes + software licenses – $50K – 6 additional network switches + cables
FY2016	– $375K – 50 additional Emulab nodes + software licenses

Questions?
Discussion?