Predicting the Unpredictable in Complex Information Systems

Kevin Mills
NIST

Joint work with Chris Dabrowski, Jim Filliben and Sandy Ressler of NIST and Yan Wan of University of North Texas

College Park, Maryland
January 24, 2014

Past Research (2006-2011)

How can we understand the influence of distributed control algorithms on global system behavior and user experience?

- Mills, Filliben, Cho, Schwartz and Genin, Study of Proposed Internet Congestion Control Mechanisms, **NIST SP 500-282** (2010).
- Mills and Filliben, "Comparison of Two Dimension-Reduction Methods for Network Simulation Models", *Journal of NIST Research* **116-5**, 771-783 (2011).
- Mills, Schwartz and Yuan, "How to Model a TCP/IP Network using only 20 Parameters", *Proceedings of the Winter Simulation Conference* (2010).
- Mills, Filliben, Cho and Schwartz, "Predicting Macroscopic Dynamics in Large Distributed Systems", *Proceedings of ASME* (2011).

INTERNET

- Mills, Filliben and Dabrowski, "An Efficient Sensitivity Analysis Method for Large Cloud Simulations", *Proceedings of the 4th International Cloud Computing Conference*, IEEE (2011).
- Mills, Filliben and Dabrowski, "Comparing VM-Placement Algorithms for On-Demand Clouds", *Proceedings of IEEE CloudCom*, 91-98 (2011).

CLOUD

For more see: http://www.nist.gov/itl/antd/emergent_behavior.cfm

What to measure

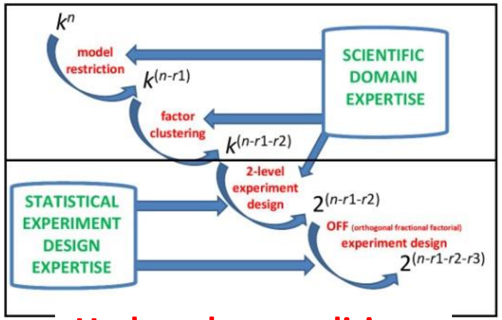

Under what conditions

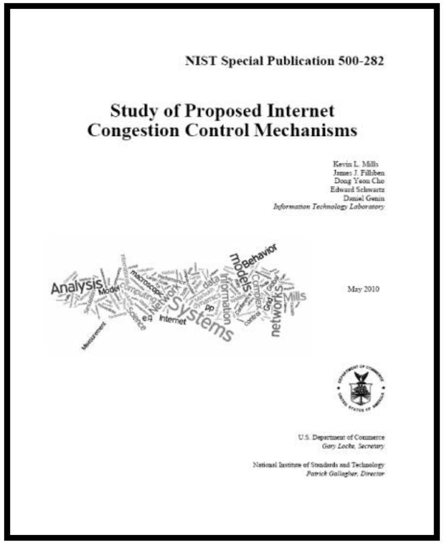
NIST Special Publication 500-282

Study of Proposed Internet Congestion Control Mechanisms

Kevin L. Mills
James J. Filliben
Dong Yeon Cho
Edward Schwartz
Daniel Genin
Information Technology Laboratory

May 2010

U.S. Department of Commerce
Gary Locke, Secretary

National Institute of Standards and Technology
Patrick Gallagher, Director

http://www.nist.gov/itl/antd/Congestion_Control_Study.cfm

At an affordable cost

Ongoing Research (2011-on)

How can we increase the reliability of complex information systems?

Research Goals: (1) develop ***design-time methods*** that system engineers can use to detect existence and causes of costly failure regimes prior to system deployment and (2) develop ***run-time methods*** that system managers can use to detect onset of costly failure regimes in deployed systems, prior to collapse.

Recent Investigation of Design-Time Methods:

- **State-space reduction techniques** (transferred from previous research)
- **Markov chains** + **cut-set analysis** + **perturbation analysis** *(SEE RELATED PAPERS)*
- **Anti-optimization** + **genetic algorithm** (*SEE RELATED PAPER*)

Ongoing Investigation of Run-Time Methods:

- Techniques (e.g., autocorrelation analysis) to **measure critical slowing down**, which may provide early warning signals for critical transitions in large systems (e.g., Scheffer et al., “Early-warning signals for critical transitions”, *NATURE*, **461**, 53-59, 2009)

Ongoing Run-Time Work

Papers published in the physics literature report and demonstrate transitions from uncongested to congested phases in simulated, abstract network models

FOUR RESEARCH QUESTIONS

(1) Are congestion-based phase transitions, as seen in abstract network models, associated with a percolation process?

(2) Do congestion-based phase transitions occur in realistic network models, and, if so, under what conditions? If not, then why not?

(3) If congestion-based phase transitions do occur in realistic network models, are there precursor patterns of behavior that can be used to signal incipient transitions so that network managers can take remedial actions to avoid an undesired transition; for example, from a free-flowing state to a jammed state?

(4) If precursor behavior patterns do exist, are there implementable detection methods that could be inserted into operational networks to provide effective signaling?

EGM Example: P. Echenique, J. Gomez-Gardenes, and Y. Moreno, "Dynamics of Jamming Transitions in Complex Networks", *Europhysics Letters*, 71, 325 (2005)

Simulations based on 11174-node scale free network*, $P_k \sim k^{-\gamma}$ & γ=2.2, taken from a 2001 snapshot of the Internet Autonomous System (AS) topology collected by the Oregon Router Server

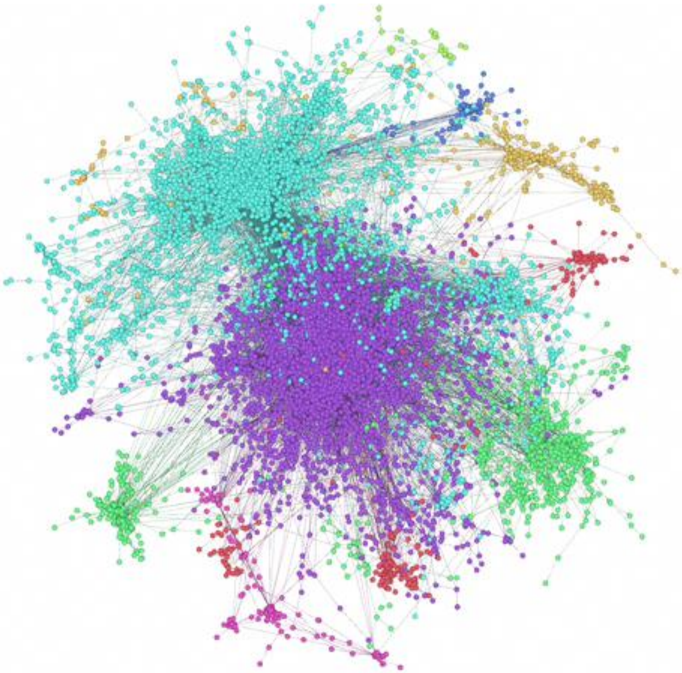

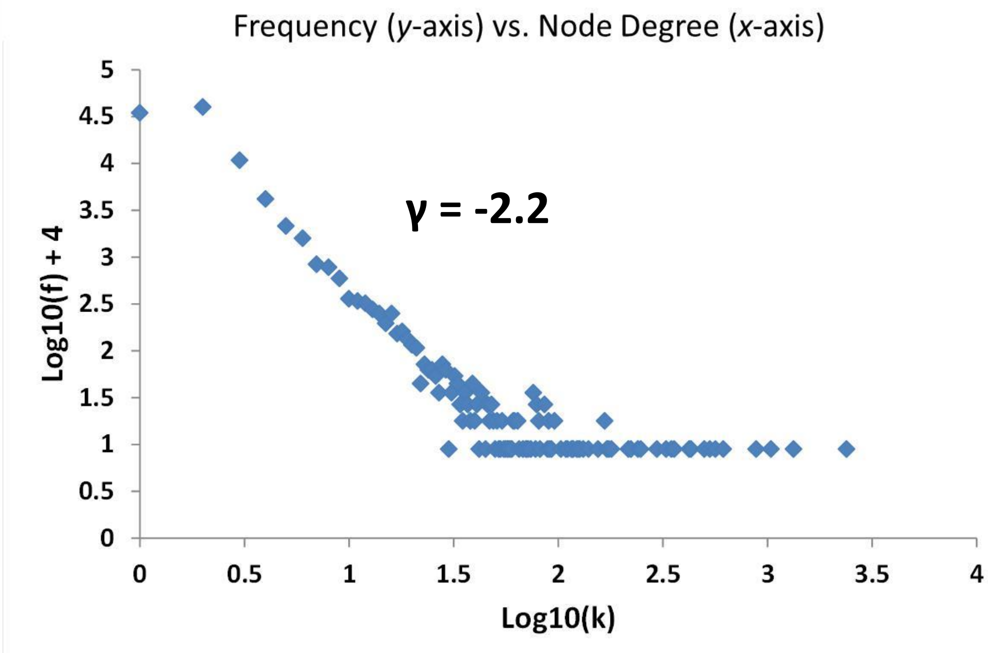

EGM simulation on the 11174-node scale free network topology

Node Buffer Size: ≤ ***q*** (∞ for EGM) packets buffered, excess packets are dropped
Injection Rate: ***p*** packets injected at random nodes (uniform) at each time step
Destination Node: chose randomly (uniform) for each packet
Forwarding Rate: 1 packet per node at each time step
Routing Algorithm: If node is destination, remove packet; Otherwise select next-hop as neighboring node *i* with minimum δ_i
System Response: proportion ρ of injected packets queued in the network

Computing δ_i

h is a ***traffic awareness*** parameter, whose value 0 ... 1.

$$\delta_i = hd_i + (1-h)c_i,$$

where *i* is the index of a node's neighbor, d_i is minimum #hops to destination via neighbor *i*, and c_i is the queue length of *i*.

***h* = 1 is shortest path**

Measuring ρ

$$\rho = \lim_{t\to\infty} \frac{A(t+\tau) - A(t)}{\tau p}$$

A = aggregate number of packets
t = time
τ = measurement interval size
p = packet inject rate

EGM simulation results + Our replication of EGM simulation

q= ∞

EGM simulation results

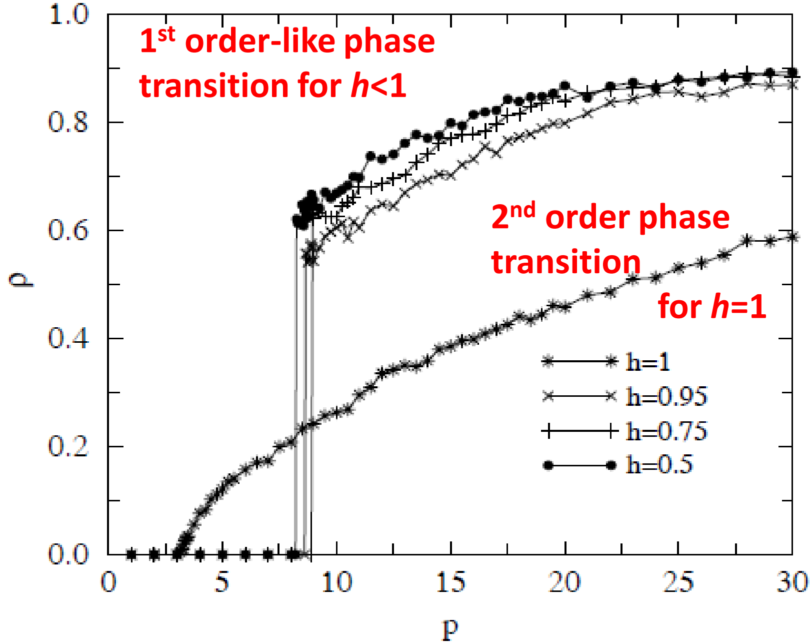

Increasing packet injection rate

Our simulation results

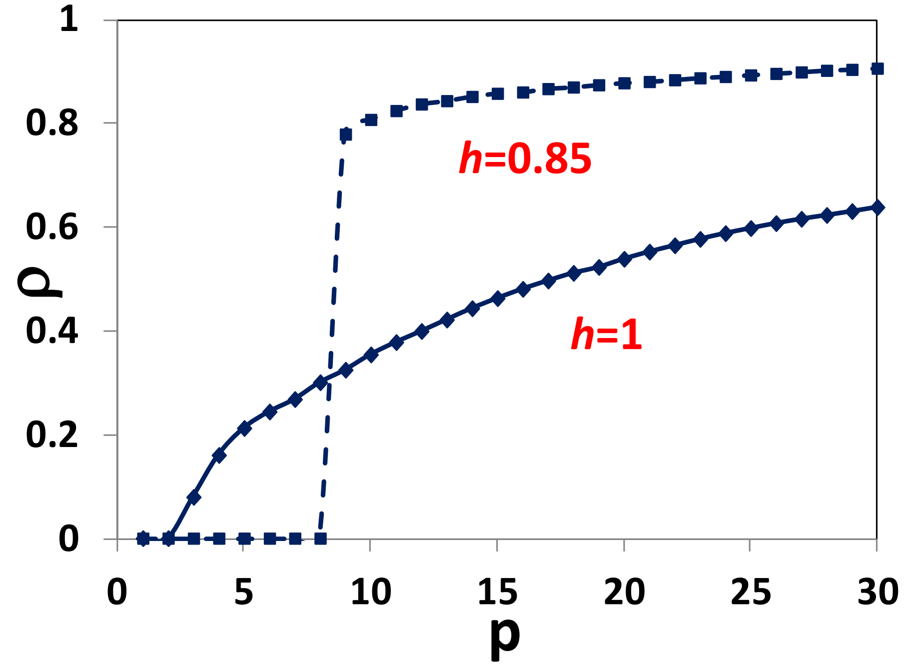

Do similar phase transitions occur in smaller networks?

q= ∞

Three-tiered, 218-node ISP Network

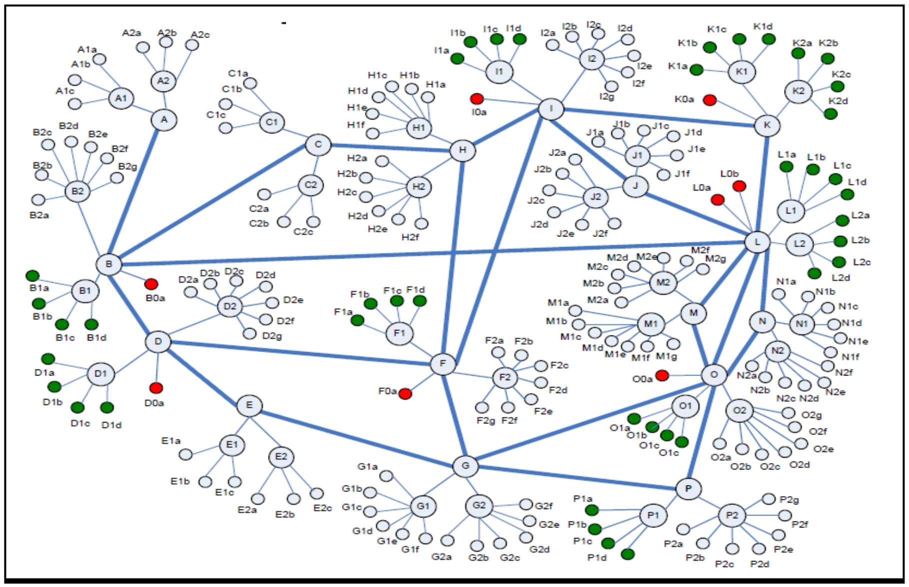

Our simulation results

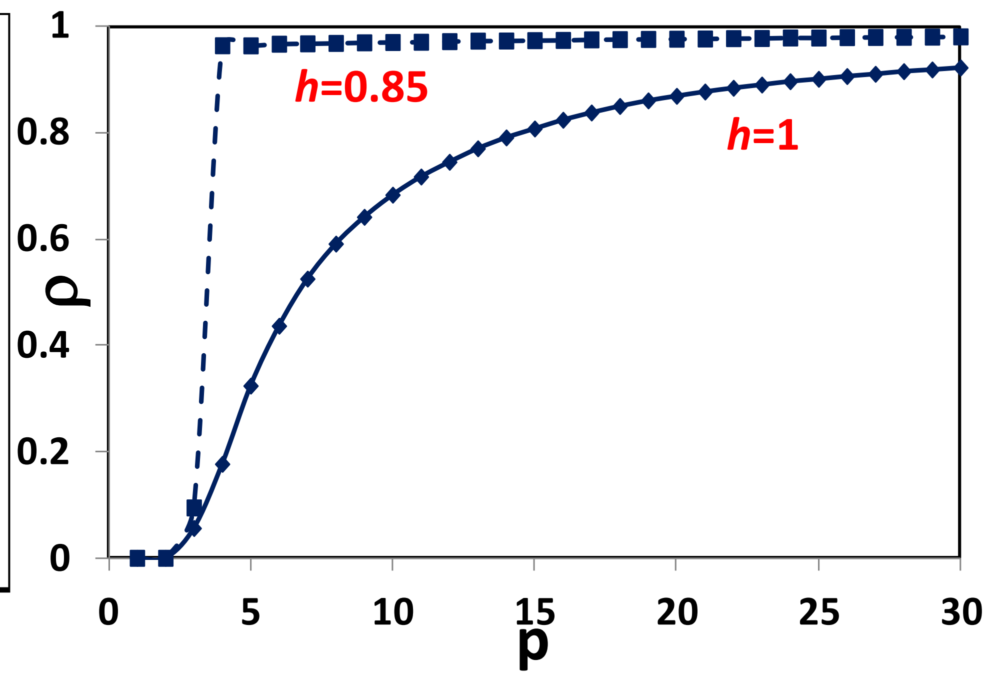

Are these phase transitions associated with a spreading process, i.e., a percolation process?

A Brief Tutorial on Percolation Theory

- **Percolation → spread of some property** of interest in a graph leads to the formation of a *giant connected component* (GCC), as measured by P_∞ - the proportion of nodes included in the GCC
- Nodes have **property of interest** with probability *p*. If $p > p_c$ (***percolation threshold***), a *percolation transition* occurs, in which a GCC forms consisting of connected sites that possess the property.
- **Order Parameter:** $P_\infty \sim (p - p_c)^\beta$, where β is known as a *critical exponent.*

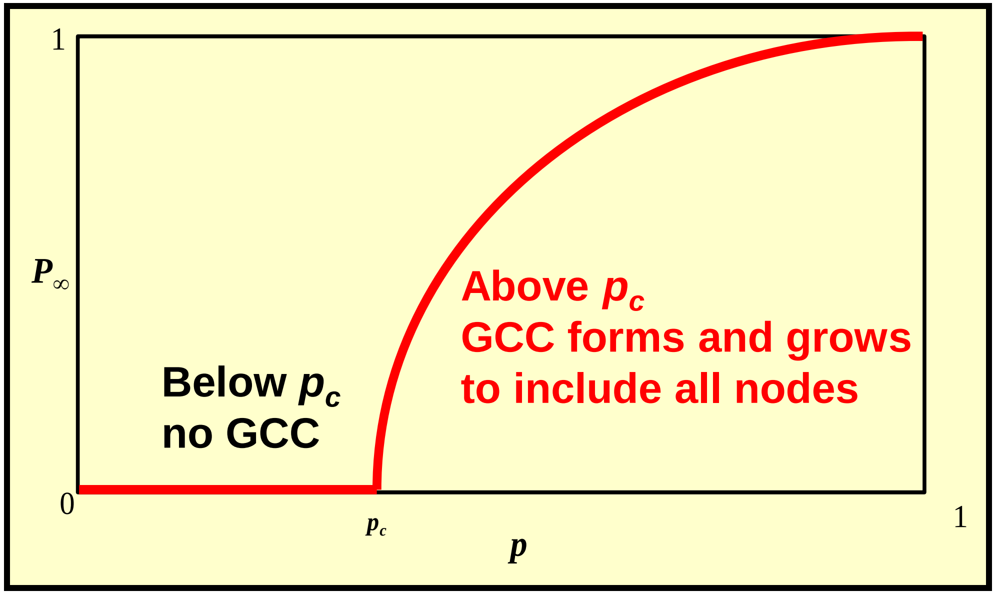

How Can We Use This Theory?

Let congestion be the property of interest

Let *p* be the packet injection rate

Let P_∞ be the proportion of congested nodes in the GCC

Is congestion spread associated with a percolation process?

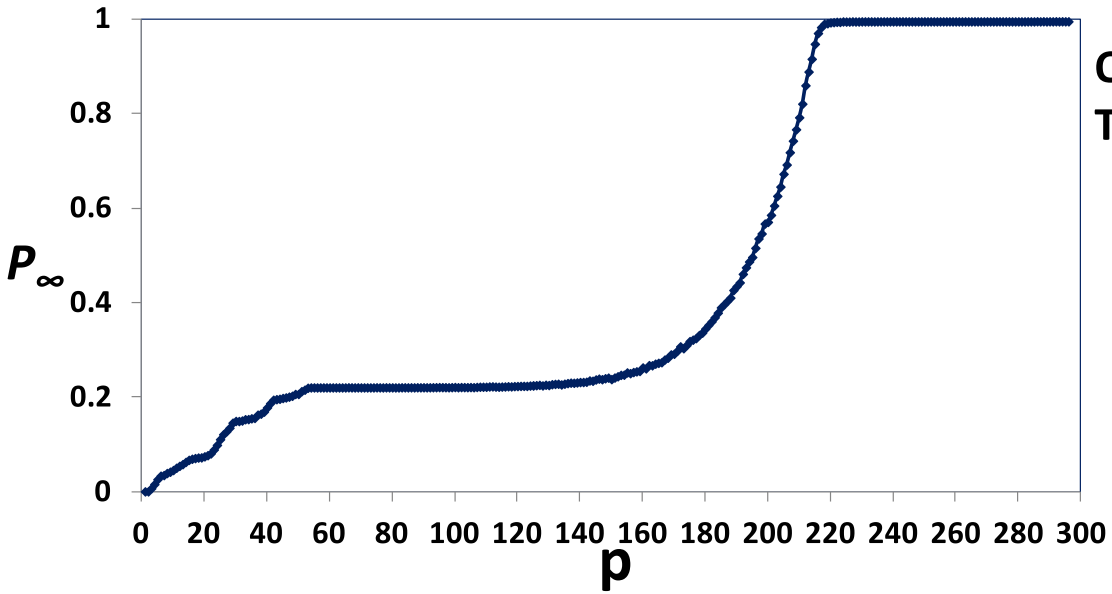

Our simulation results for Three-tiered, 218-node ISP Network

$h=1$

$q=\infty$

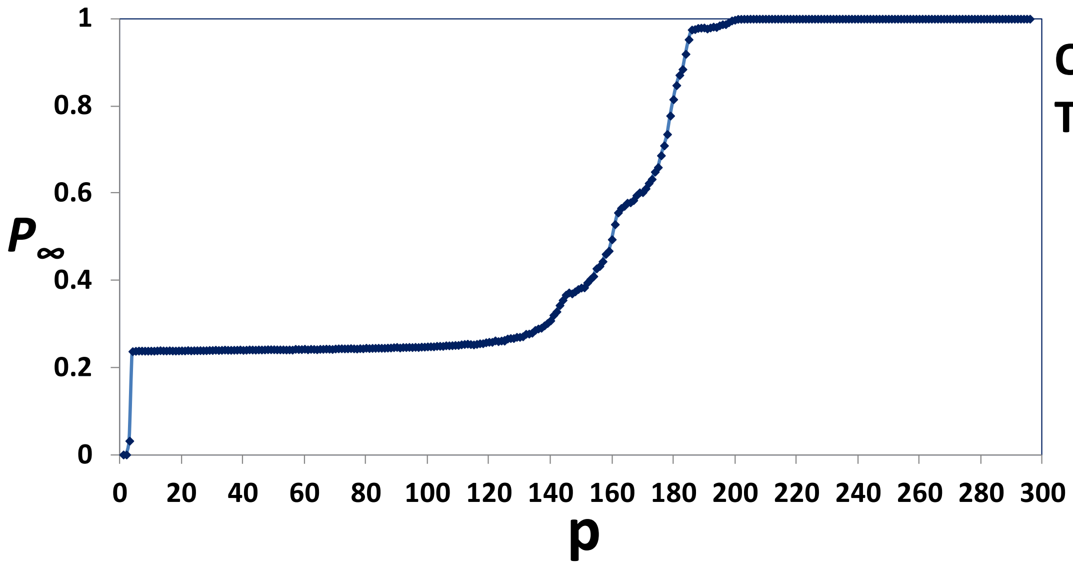

Our simulation results for Three-tiered, 218-node ISP Network

$h=0.85$

$q=\infty$

Does percolation occur when queue lengths are bounded?

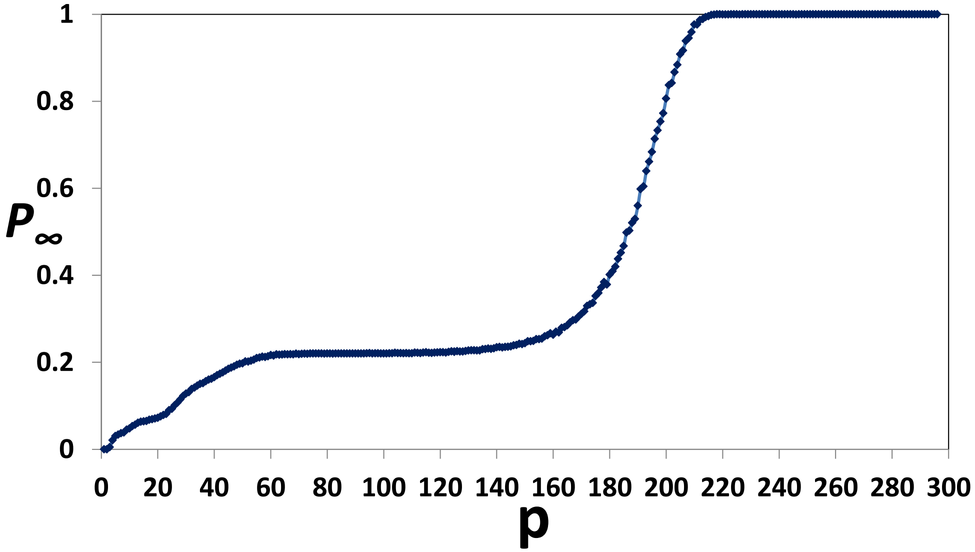

Our simulation results for Three-tiered, 218-node ISP Network

***h*=0.85**

***q*=10**

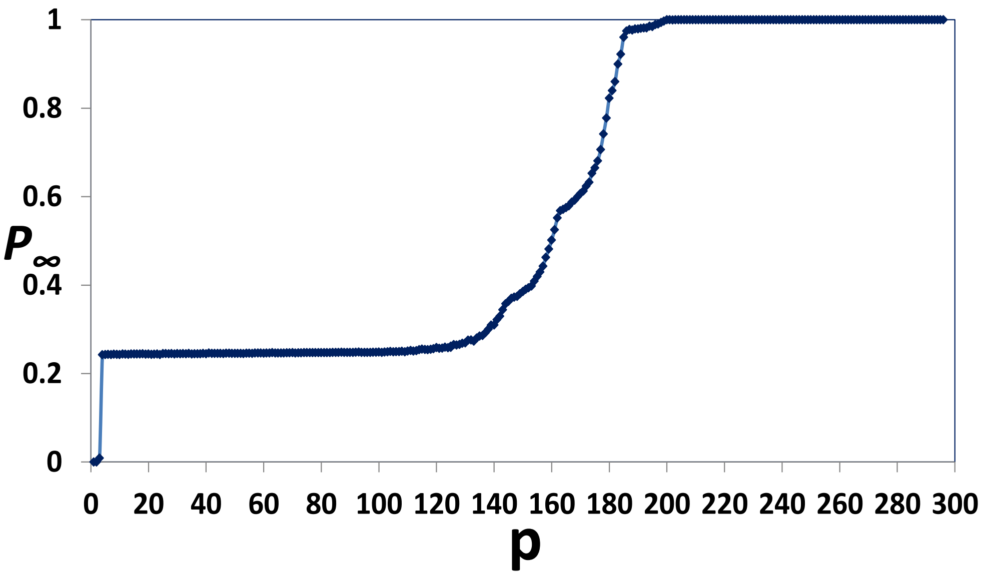

Our simulation results for Three-tiered, 218-node ISP Network

***h*=0.85**

***q*=100**

FUTURE WORK I: Do percolation and/or associated phase transitions occur when networks exhibit more realistic characteristics?

(1) **More complex topologies** – cross links between POP routers – as well as inter-AS topologies.
(2) **Varied router speeds** – engineered in a reasonable manner.
(3) **Propagation delays** – on transit links.
(4) **Sources and receivers attached to access routers** – only sources inject data packets.
(5) **Sources and receivers distributed** – around the network in a non-uniform pattern.
(6) **Sources and receivers have bounded interface speeds** – with limited variation.
(7) **Sources modeled explicitly** – as cyclic ON-OFF processes.
(8) **Sources show limited patience** – with prolonged or slow transfers.
(9) **Sources transfer randomly chosen file sizes** – from a variety of classes.
(10) **Sources use TCP procedures** – including connection establishment, slow start and congestion avoidance.

FUTURE WORK II: If percolation and/or associated phase transitions do occur in realistic network models, then:

(1) **Are there precursor patterns of behavior that can be used to signal incipient transitions?**

(2) **If precursor patterns exist, are there pragmatic, implementable detection methods that could be inserted into operational networks to provide effective signaling?**

Questions?

☺ *Thanks for Listening* ☺

Contact information about the NIST Complex Systems program:
{boisvert@nist.gov}

Contact information about studying Complex Information Systems:
{kmills, jfilliben, cdabrowski@nist.gov}

Contact information about Information Visualization for Complex Systems:
sressler@nist.gov

Contact information about related grant-based research, *Run-Time Failure Detection and Control in Complex Information Systems*:
Yan.Wan@unt.edu

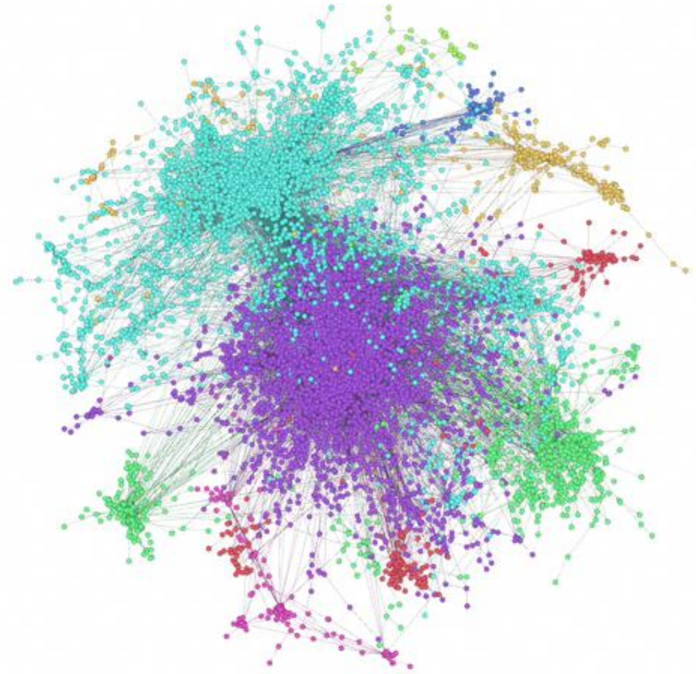

More @: http://www.nist.gov/itl/antd/emergent_behavior.cfm