The Need for Realism when Simulating Network Congestion*

Kevin Mills and Chris Dabrowski
National Institute of Standards & Technology

Communications & Networking Symposium
April 3-6, 2016

*For more details see: NIST Technical Note 1905
http://dx.doi.org/10.6028/NIST.TN.1905

Information Technology Laboratory
NIST

Total talk is ≈ 16 slides

- Motivation – 2 slides
- Research Questions and Approach – 2 slides
- Models – 5 slides
- Experiment Design – 1 slide
- Results – 5 slides
- Findings – 1 slide

NIST Information Technology Laboratory

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

All Find that Signals Appear Near a Critical Point in Abstract Network Models

Information Technology Laboratory NIST

Abstract Models Lack Key Traits of Real Networks

Routers & Links

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

Computers

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

Users

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

Protocols

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

DOES LACK OF REALISM MATTER WHEN SIMULATATING NETWORK CONGESTION?

Information Technology Laboratory

Specific Research Questions

1. Does congestion spread in abstract models mirror spread in realistic models?
2. Are some elements of realism essential to capture when modeling network congestion?
3. Are some elements unnecessary?
4. What measures of congestion can be compared, and how, across diverse network models?

Information Technology Laboratory NIST

Research Approach

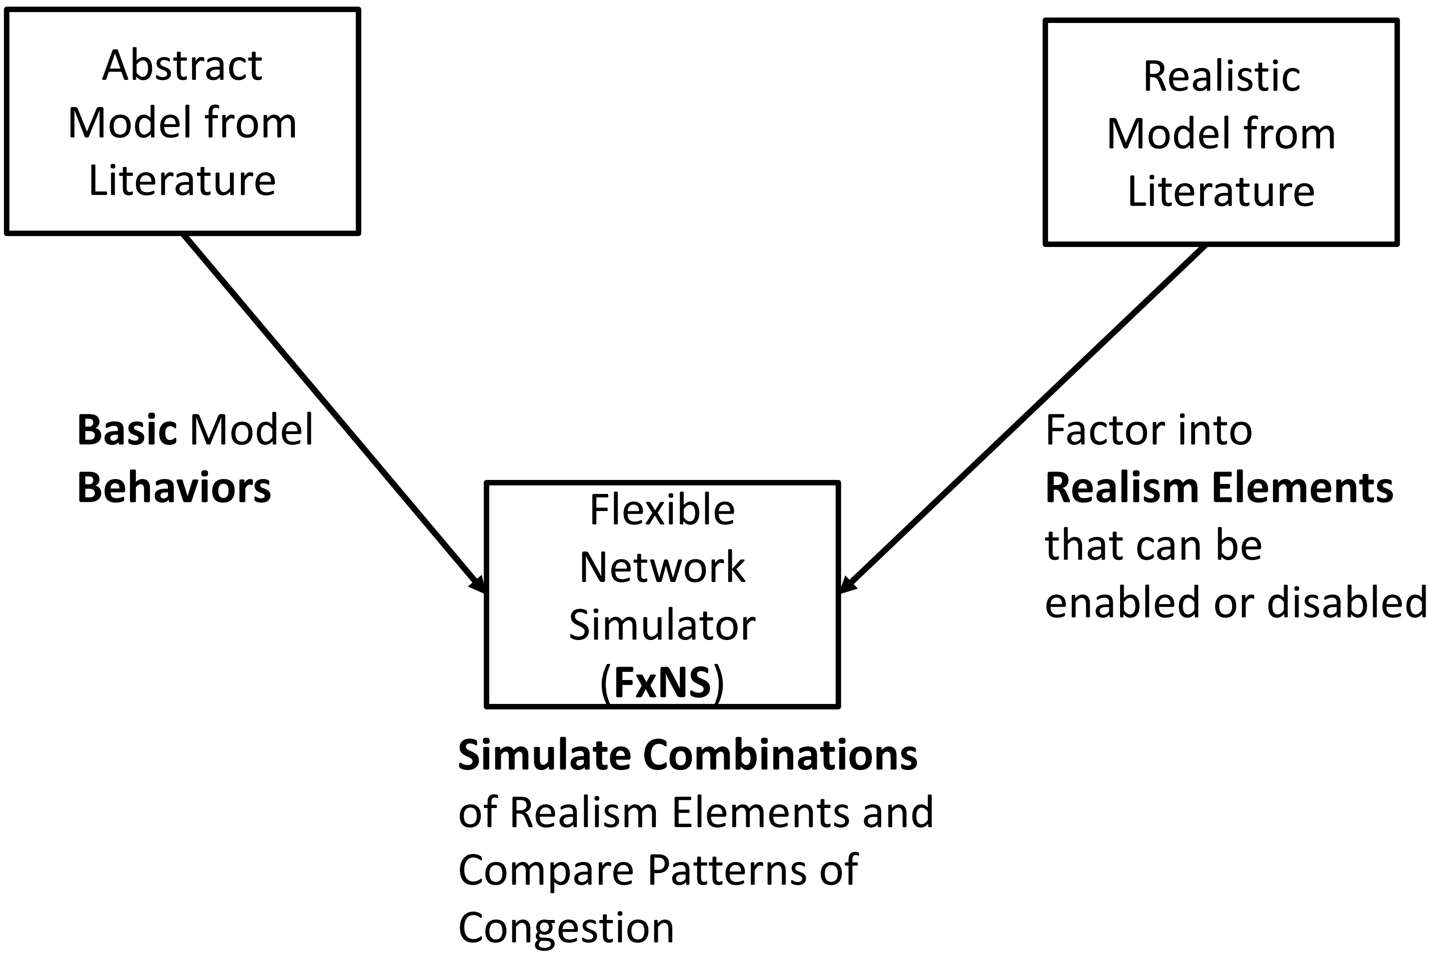

Information Technology Laboratory

Models

- Abstract **EGM** Model→high abstraction
- Realistic **MesoNet** Model→high realism
- Flexible **FxNS** Model→combinations of realism from low to high

Information Technology Laboratory NIST

The Abstract (EGM) Model

P. **E**chenique, J. **G**omez-Gardenes, and Y. **M**oreno, "Dynamics of Jamming Transitions in Complex Networks", *Europhysics Letters*, 71, 325 (2005)

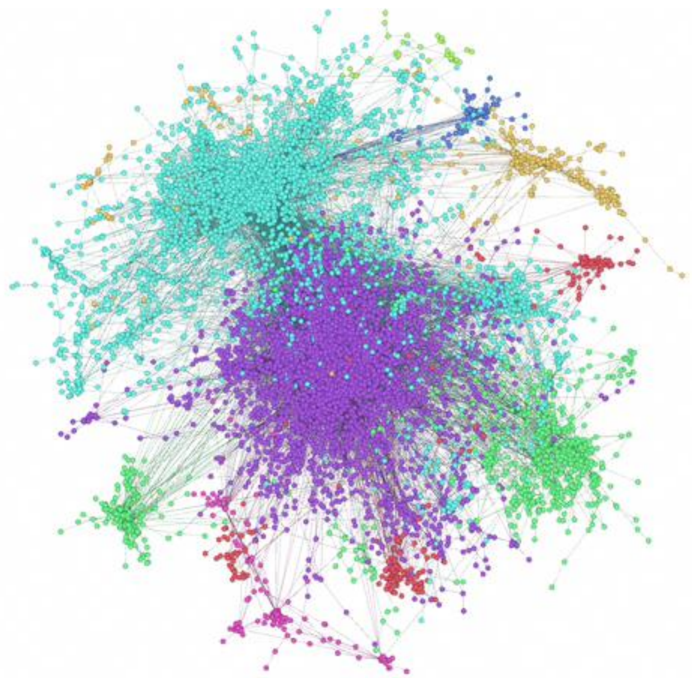

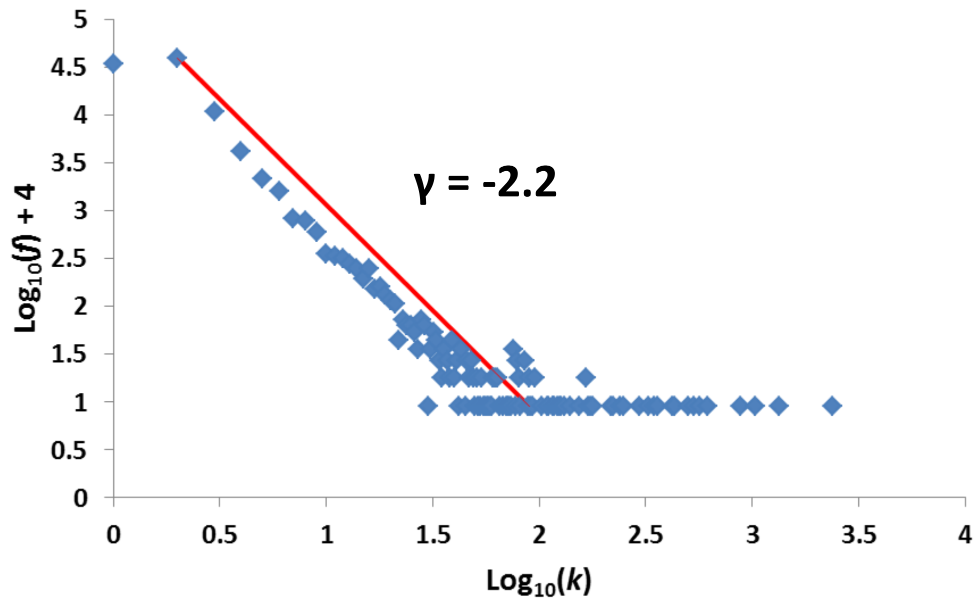

Simulations based on 11,174-node scale-free graph, $P_k \sim k^{-\gamma}$ & γ=2.2, taken from a 2001 snapshot of the Internet Autonomous System (AS) topology collected by the Oregon Router Server (image courtesy **Sandy Ressler**)

NIST Information Technology Laboratory

Details of the EGM Model

Node Buffer Size: **∞ for EGM, all packets buffered, no packets 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 (in hops)**

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**

Comparative Simulation Results

EGM Simulations

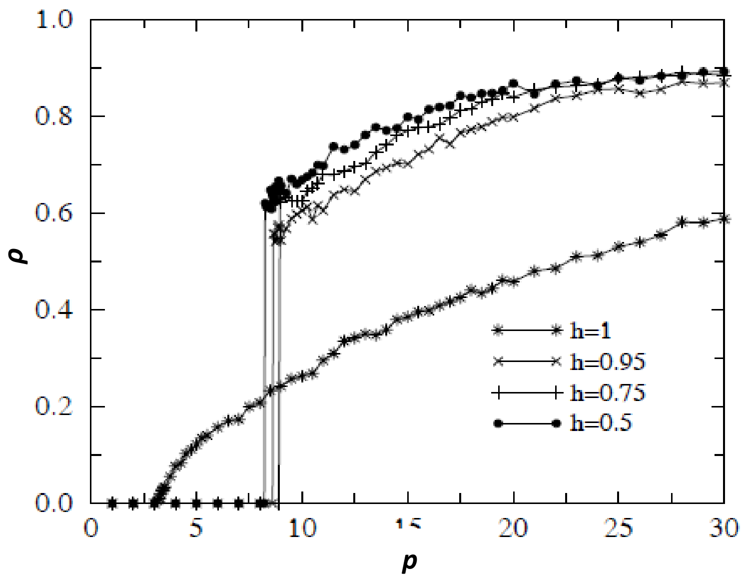

FxNS Simulations with
All Realism Elements Disabled

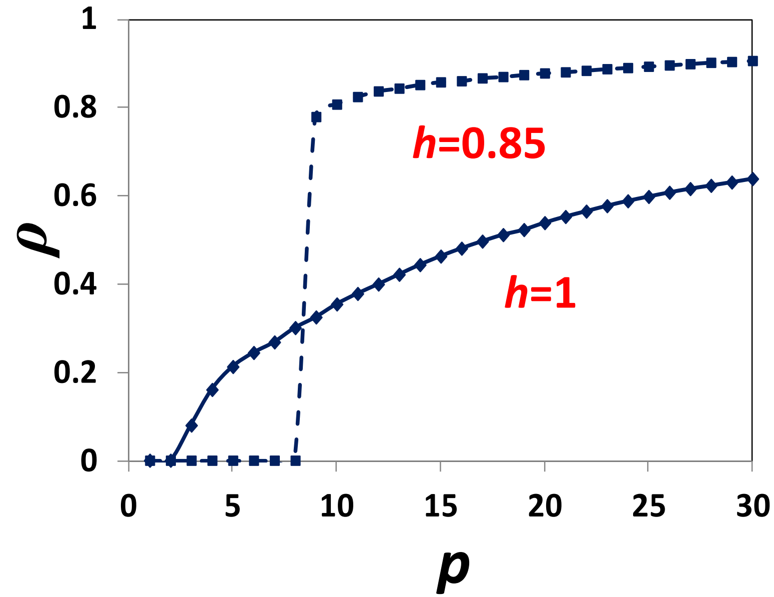

NIST
Information Technology Laboratory

The Realistic (MesoNet) Model

K. Mills, E. Schwartz, and J. Yuan, "How to Model a TCP/IP Network using only 20 Parameters", WSC 2010, Dec. 5-8, Baltimore, MD.

Category	ID	Name	FxNS
Network	x1	topology	NC
	x2	propagation delay	DE
	x3	network speed	VS
	x4	buffer provisioning	PD
Sources & Sinks	x5	number sources/sinks	SR
	x6	source distribution	
	x7	sink distribution	
	x8	source/sink speed	VS
Users	x9	think time	*p*
	x10	patience	n/a
	x11	web object file sizes	FL
	x12	larger file sizes	n/a
	x13	localized congestion	
	x14	long-lived flows	
Congestion Control	x15	control algorithm	TCP
	x16	initial *cwnd*	
	x17	Initial *sst*	
Simulation Control	x18	measurement interval	fixed
	x19	simulation duration	fixed
	x20	startup pattern	*p*

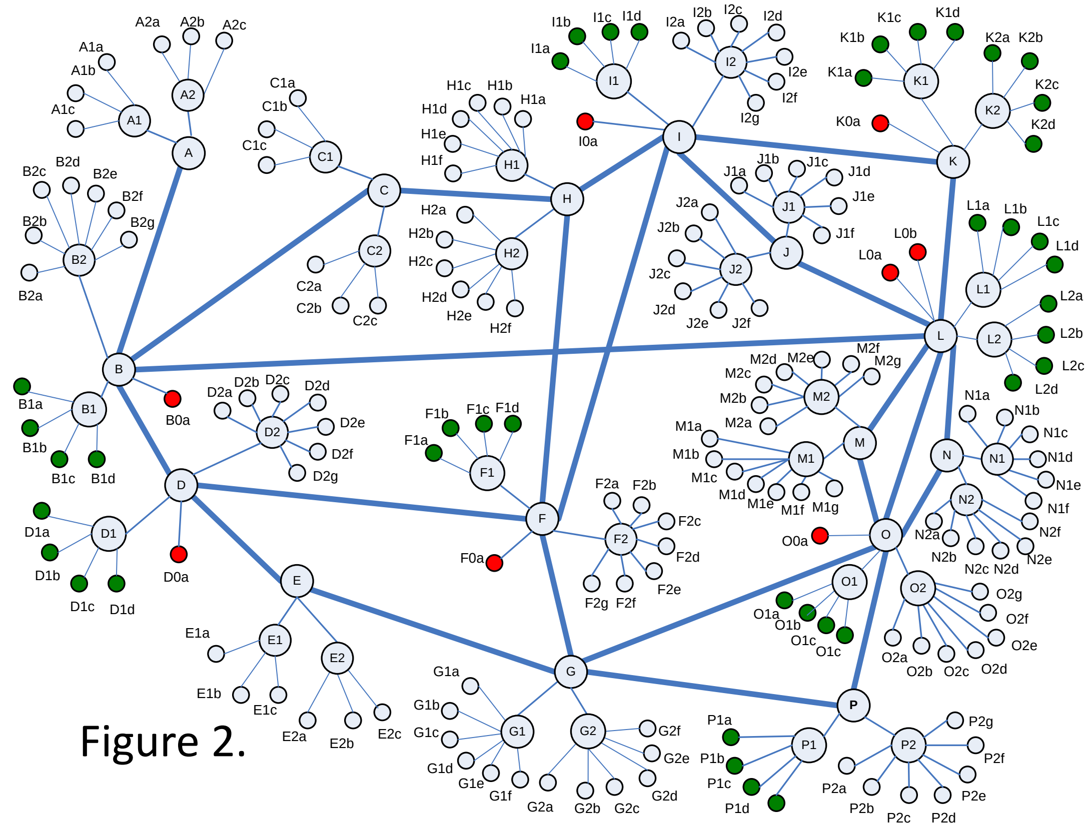

Figure 2.

Comparisons of MesoNet Simulations vs. FxNS Simulations (all realism elements enabled) for eight MesoNet responses are available in **NIST TN 1905 – Appendix A**

NIST Information Technology Laboratory

FxNS Combinations

7 **Realism Elements**

PD	**Packet Dropping**
NC	**Node Classes**
VS	**Variable Speeds**
DE	**Propagation Delay**
SR	**Sources and Receivers**
FL	**Flows**
TCP	**Transmission Control Protocol**

7 **Dependencies** among Realism Elements

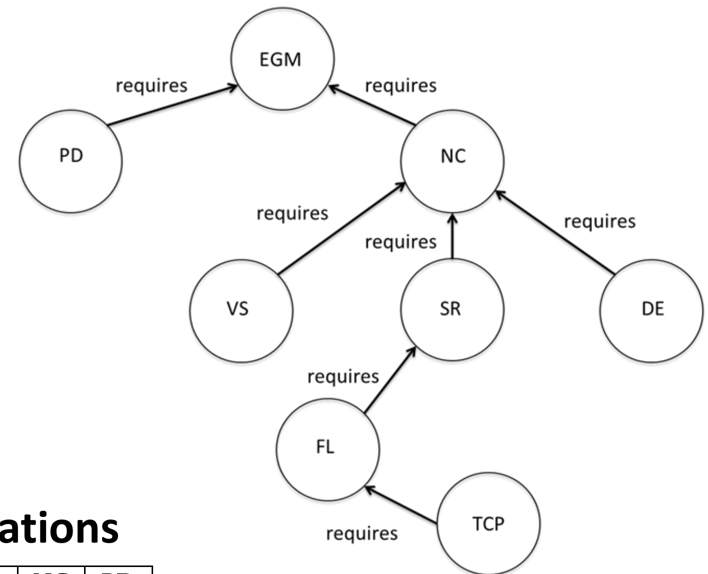

34 Valid FxNS **Combinations**

Seq	Cmb	TCP	FL	SR	DE	VS	NC	PD
1	c0	0	0	0	0	0	0	0
2	c1	0	0	0	0	0	0	1
3	c2	0	0	0	0	0	1	0

. . .

Seq	Cmb	TCP	FL	SR	DE	VS	NC	PD
32	c123	1	1	1	1	0	1	1
33	c126	1	1	1	1	1	1	0
34	c127	1	1	1	1	1	1	1

NIST Information Technology Laboratory

Experiment Design

	Enabled	Disabled
PD	buffers = 250×router speed	buffers = ∞
NC	3-tier 218-node topology as in Fig. 2 with routers labeled as core, PoP, D-class, F-class or N-class	flat 218-node topology as in Fig. 2 but with routers unlabeled
VS	core 80 p/ts; PoP 10 p/ts; D-class 10 p/ts; F-class 2 p/ts; N-class 1 p/ts; fast source/sink 2 p/ts; normal source/sink 0.2 p/ts	all routers and sources/sinks 9 p/ts
DE	core links have propagation delays	no propagation delays
SR	51,588 sources & 206,352 sinks deployed uniformly below access routers	no sources or sinks deployed
FL	transfers are packet streams: sized randomly from Pareto distribution (mean 350, shape 1.5) - streams set up with TCP connection procedures	transfers are individual packets
TCP	packet transmission regulated by TCP congestion-control including slow-start (initial *cwnd* = 2 *sst* = $2^{30}/2$) and congestion avoidance	packet transmissions not regulated by congestion-control

FIXED PARAMETERS

- 218-Router Topology (Fig. 2)
- Routing (SPF propagation delay)
- Duration (200,000 ts per *p*)

VARIABLE PARAMETERS

- Packet-Injection Rate *p* (up to 2500)
- FxNS Combination

RESPONSES

- Congestion Spread $\chi=|G_\chi|/|G_N|$
- Connectivity Breakdown $\alpha=|G_\alpha|/|G_N|$
- Proportion of Packets Delivered π
- Scaled (0..1) Latency of Delivered Packets δ

Only concepts in common among all 34 combinations: graph and packet

Results[1,2]

[1] **136 *xy*-plots** (34 FxNS combinations × 4 responses) are available at:
http://tinyurl.com/poylful

[2] Related FxNS simulation data can be explored interactively using a **multidimensional visualization** created by **Phillip Gough** of CSIRO:
http://tinyurl.com/payglq6

NIST Information Technology Laboratory

Results I – Abstract (C0) vs. Realistic (C127)

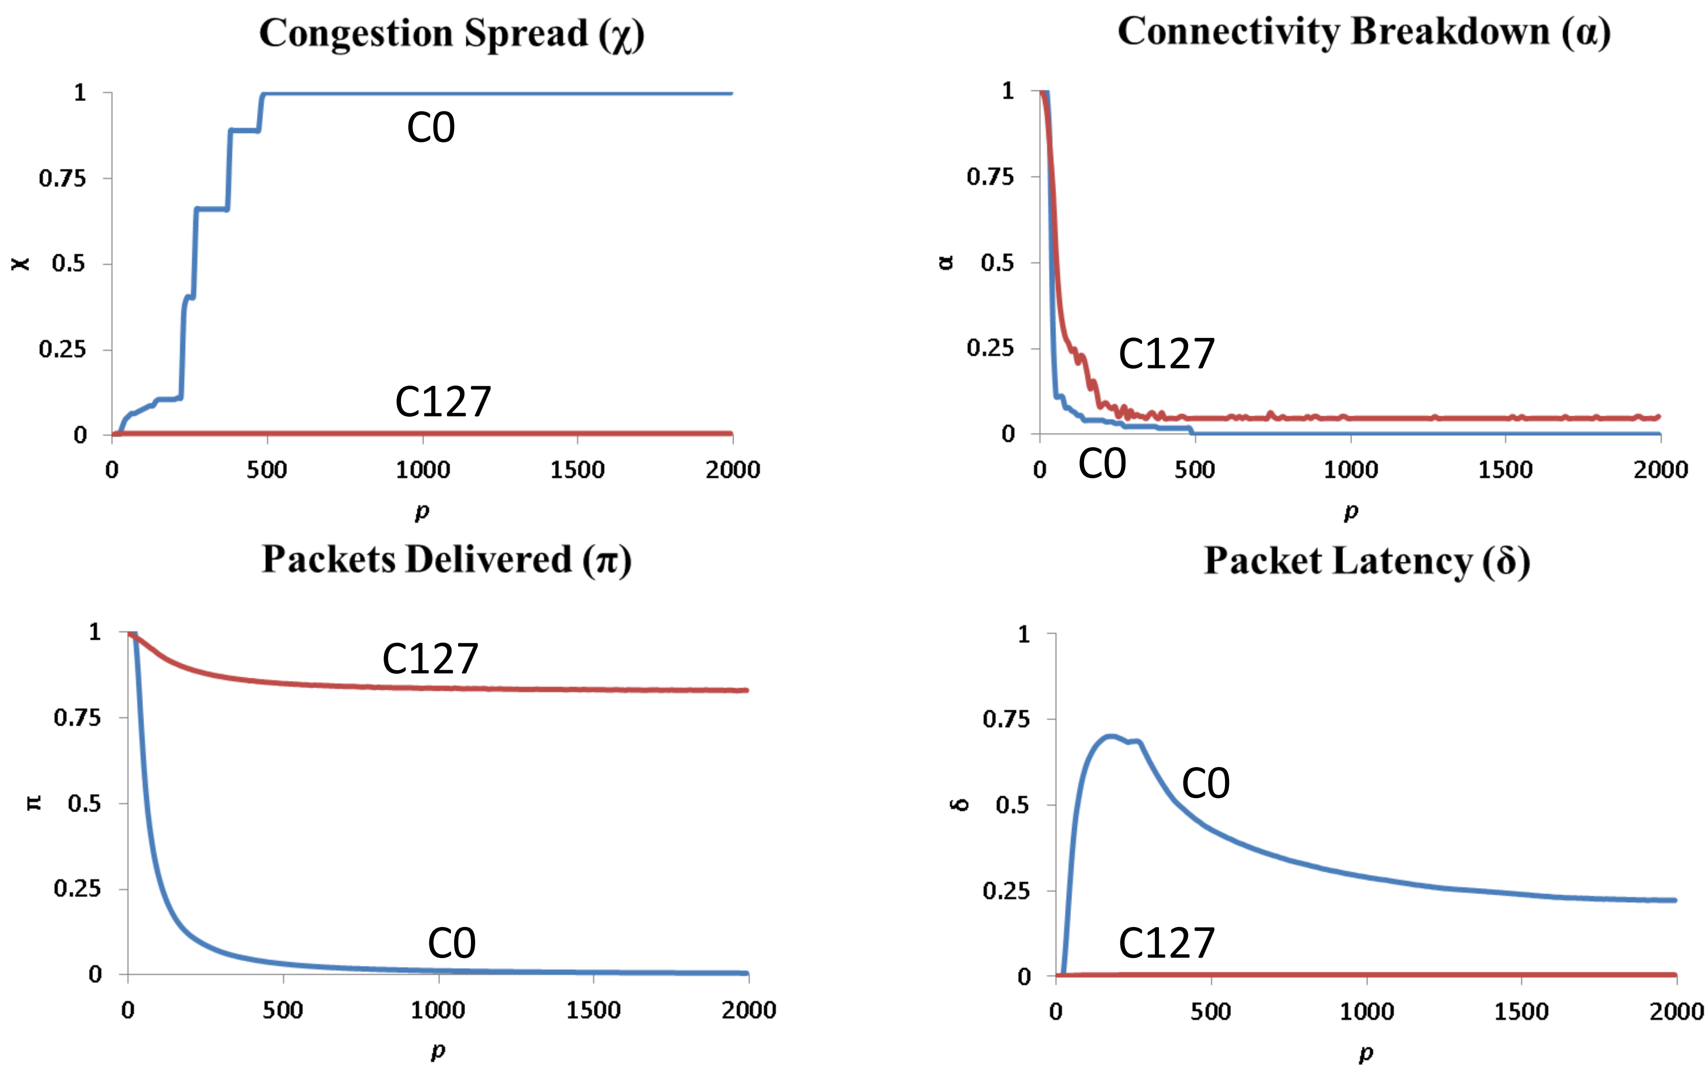

Plots for all responses and all 34 combinations available: http://tinyurl.com/poylful

Results II – Congestion Spread χ All Combinations

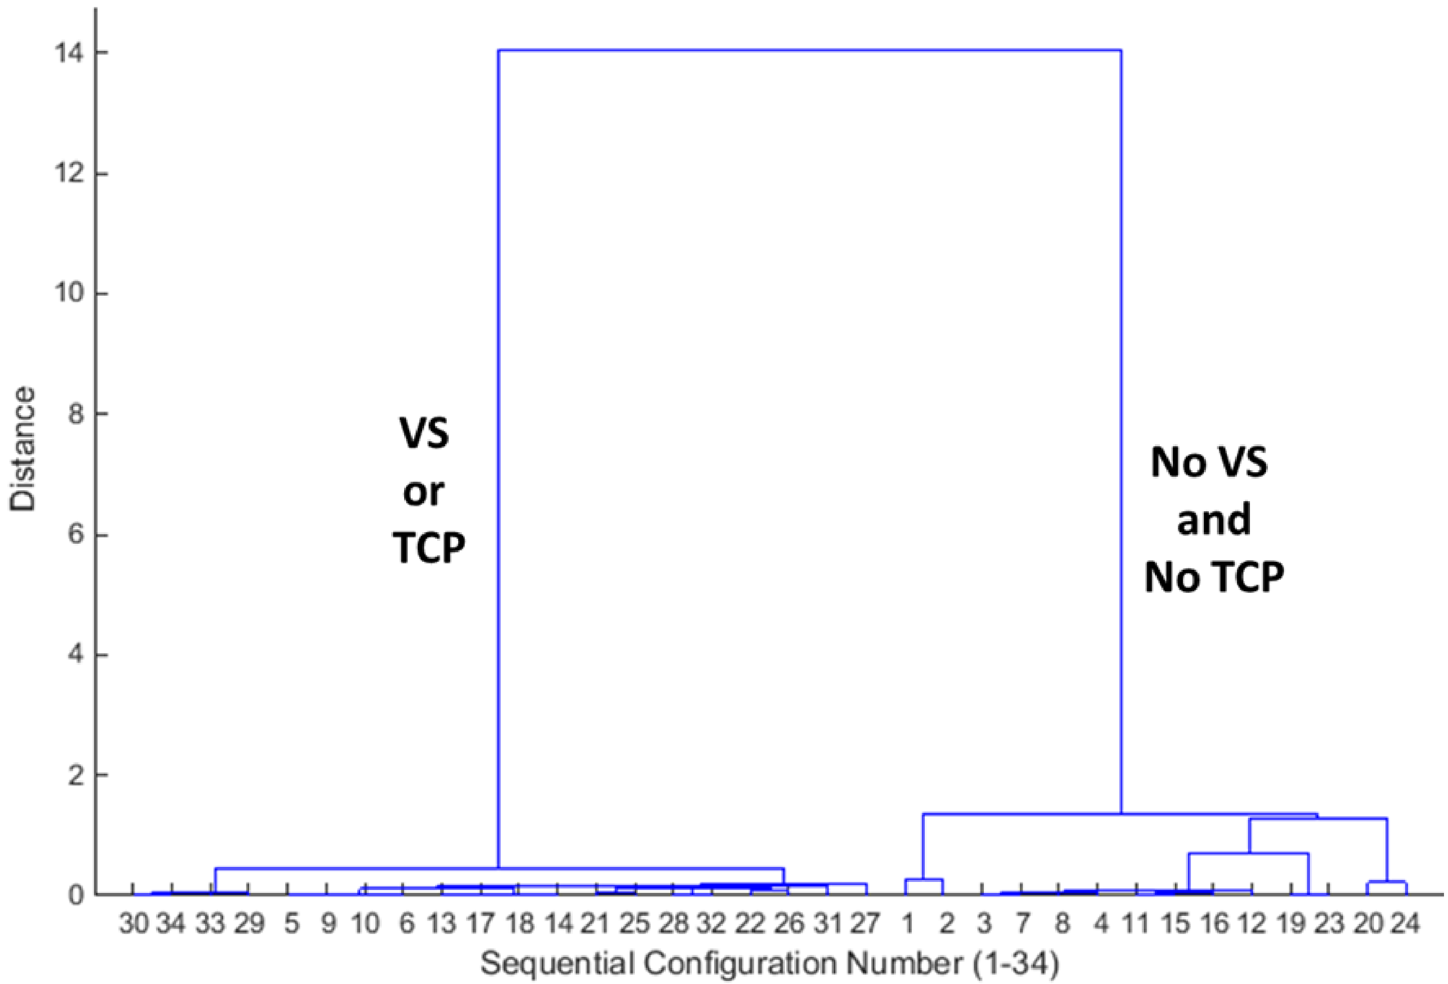

Information Technology Laboratory
NIST

Results III – Connectivity Breakdown α All Combinations

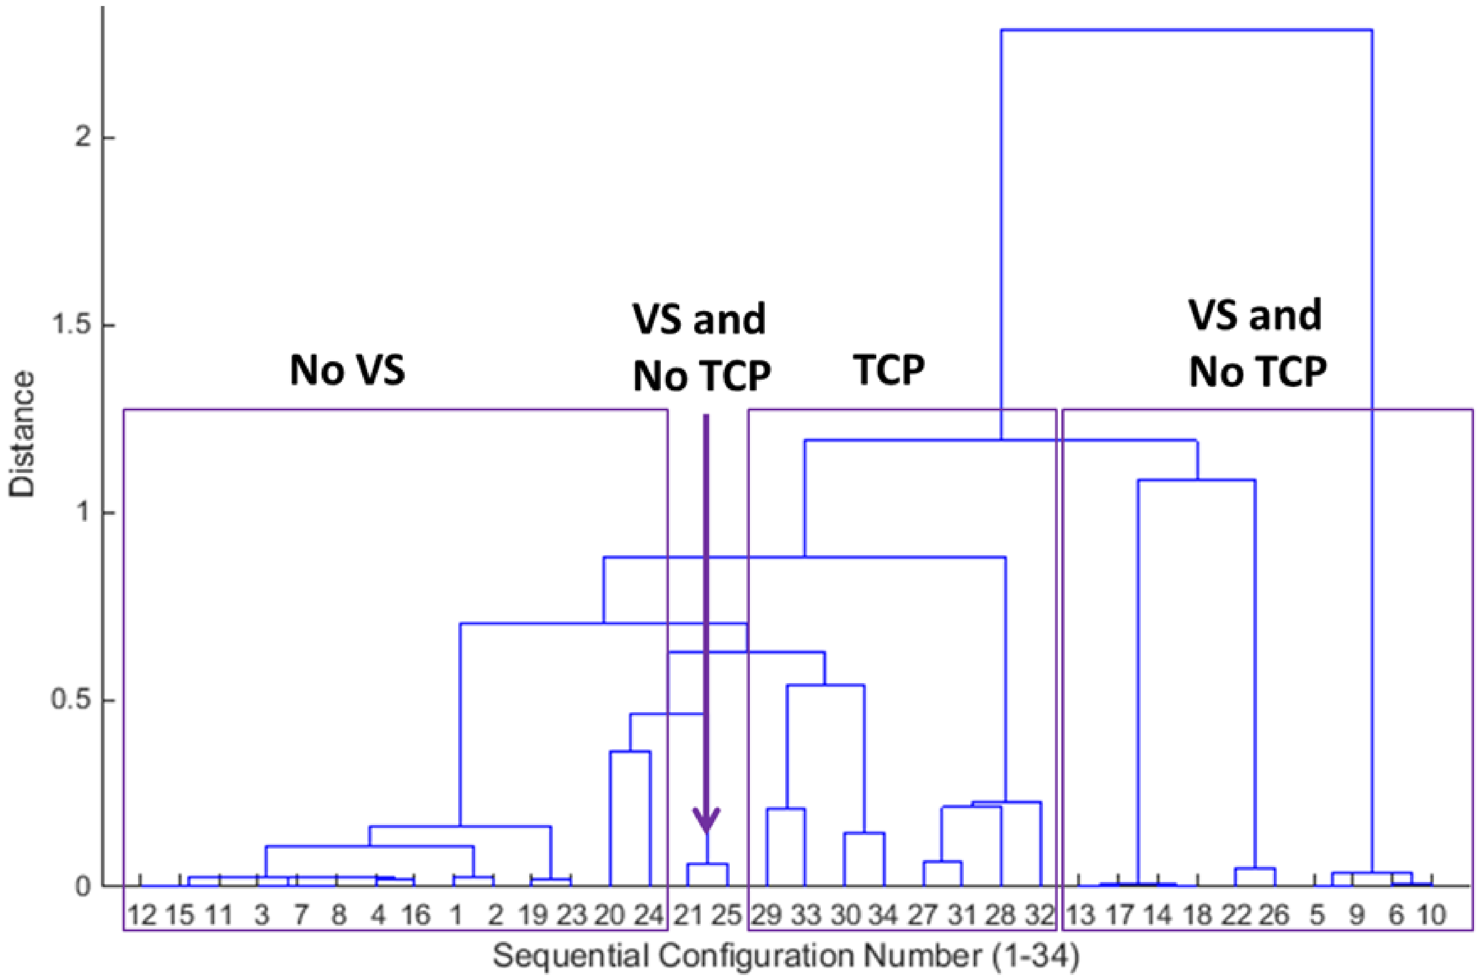

Results IV – Packet Delivery π All Combinations

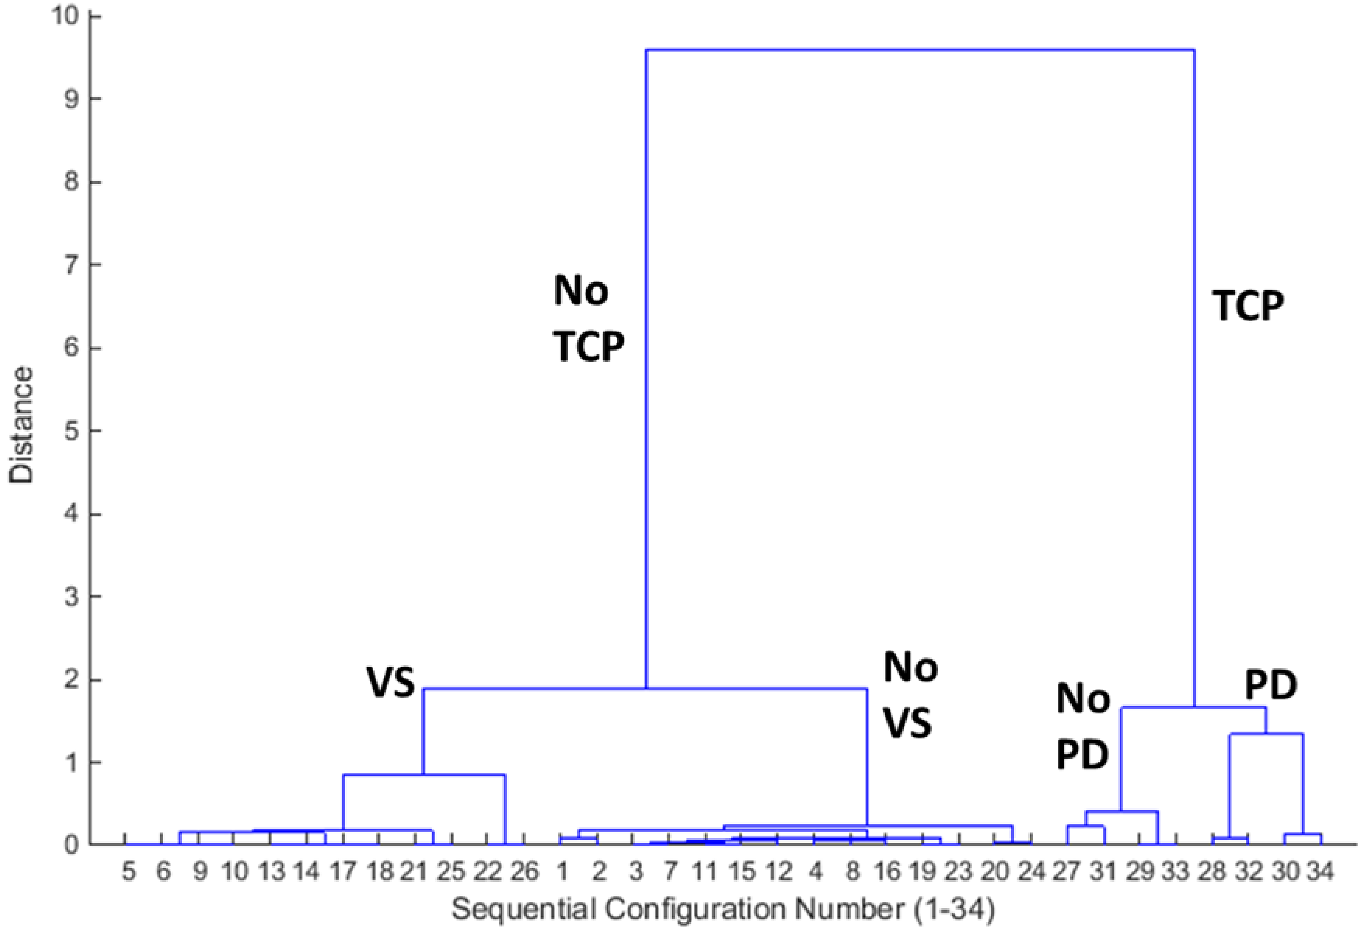

Information Technology Laboratory NIST

Results V – Scaled Packet Latency δ All Combinations

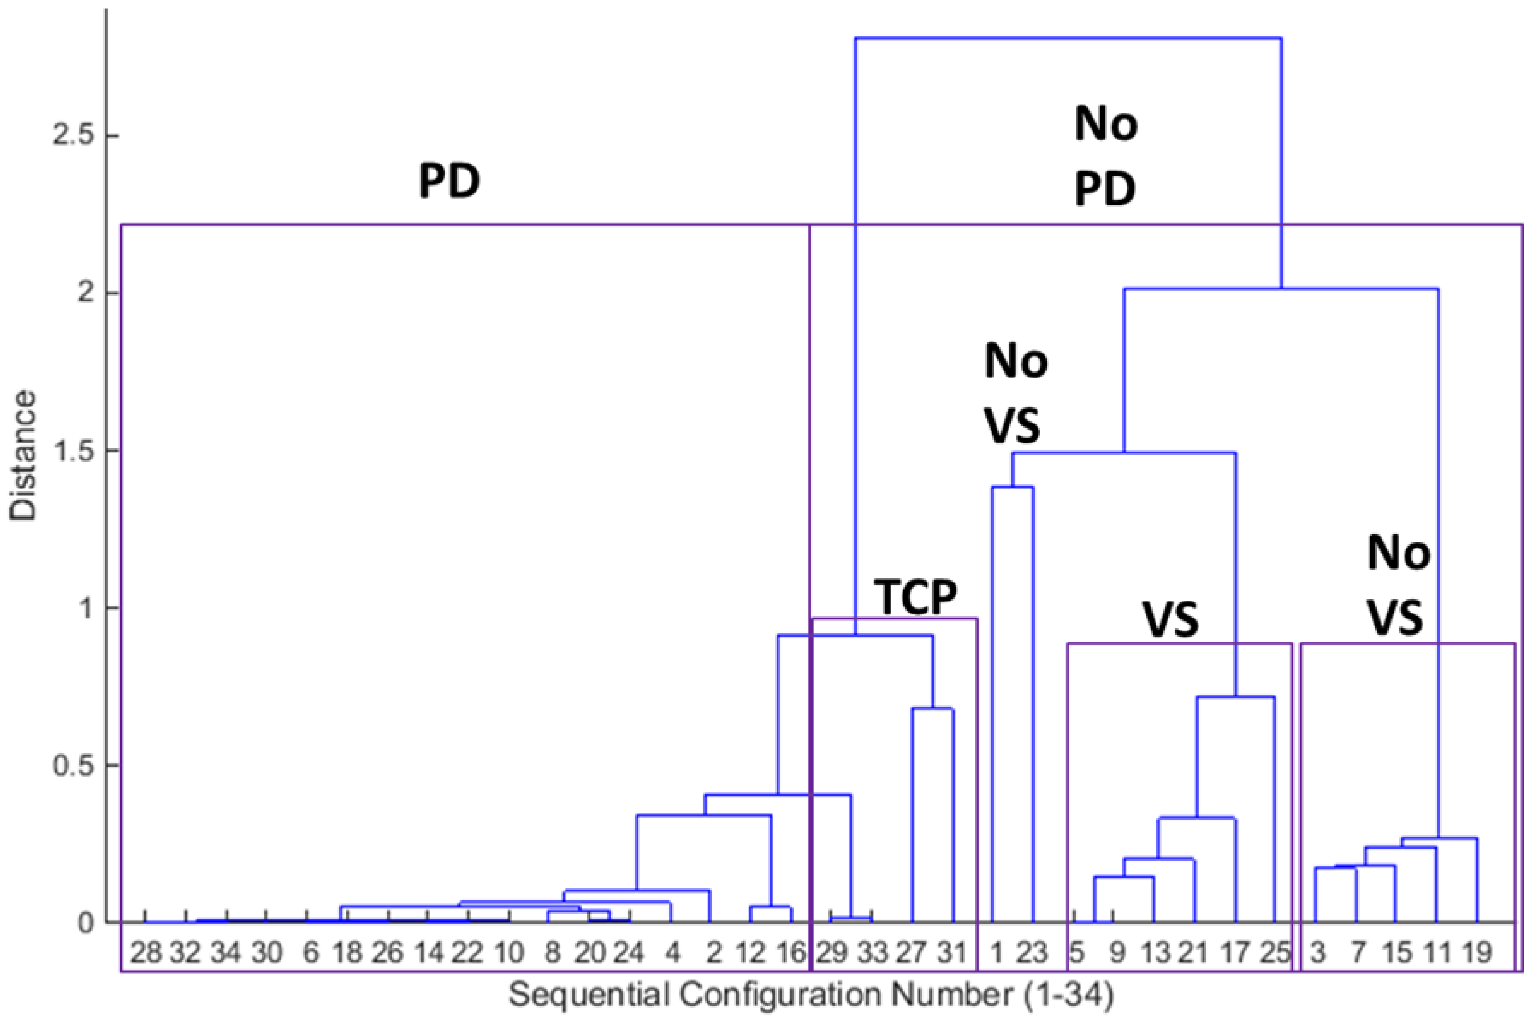

Findings

- Congestion spreads differently in abstract and realistic models
- Hierarchical Router Speeds and TCP very important to model
- Packet dropping important to model for accurate packet latencies
- Propagation delay not important to model in a continental US network, but would be important to model in topologies where propagation delays exceed queuing delays
- Congestion spread, connectivity breakdown and the effectiveness and efficiency of packet delivery can be measured using only two concepts: graphs and packets

NIST Information Technology Laboratory

For more of our research see:

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

NIST | NIST Time | NIST Home | About NIST | Contact Us | A-Z Site Index | Search

Information Technology Laboratory

About ITL ▼ | Publications | Topic/Subject Areas ▼ | Products/Services ▼ | News/Multimedia | Programs/Projects

NIST Home > ITL > Advanced Network Technologies Division > Measurement Science for Complex Information Systems

Select Language
Powered by Google Translate | SHARE

Measurement Science for Complex Information Systems

Summary:

This project aims to develop and evaluate a coherent set of methods to understand behavior in complex information systems, such as the Internet, computational grids and computing clouds. Such large distributed systems exhibit global behavior arising from independent decisions made by many simultaneous actors, which adapt their behavior based on local measurements of system state. Actor adaptations shift the global system state, influencing subsequent measurements, leading to further adaptations. This continuous cycle of measurement and adaptation drives a time-varying global behavior. For this reason, proposed changes in actor decision algorithms must be examined at large spatiotemporal scale in order to predict system behavior. This presents a challenging problem.

Description:

What are complex systems? Large collections of interconnected components whose interactions lead to macroscopic behaviors in:

- Biological systems (e.g., slime molds, ant colonies, embryos)
- Physical systems (e.g., earthquakes, avalanches, forest fires)
- Social systems (e.g., transportation networks, cities, economies)
- Information systems (e.g., Internet and compute clouds)

What is the problem? No one understands how to measure, predict or control macroscopic behavior in complex information systems: (1) threatening our nation's security

Internet Autonomous System Graph Circa 2001 - Image by Sandy Ressler

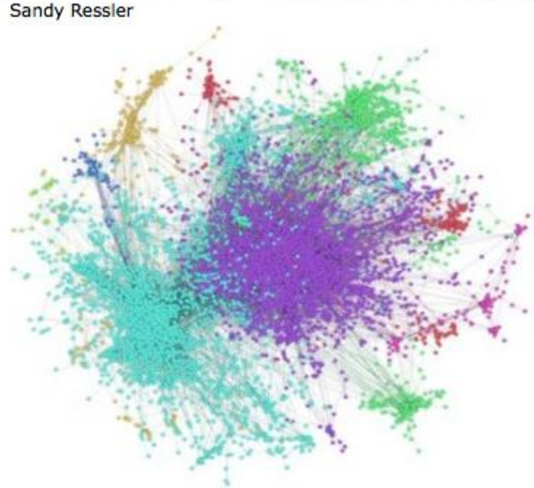

Start Date:

October 2, 2006

End Date:

ongoing