National Institute of Standards and Technology

NIST National Institute of Standards and Technology

Predicting Global Failure Regimes in Complex Information Networks

Chris Dabrowski, Jim Filliben and **Kevin Mills**

July 12, 2012 *SFI Workshop on Measurement of Complex Networks*

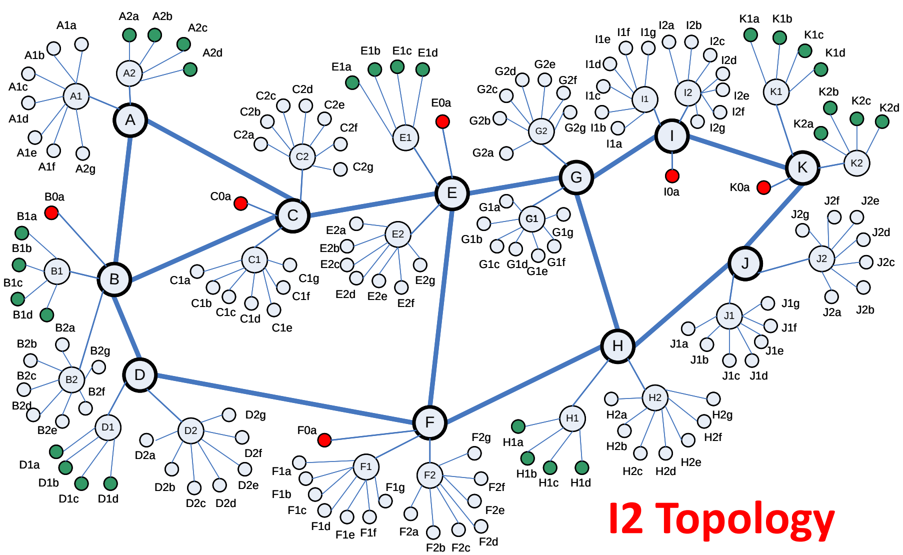

I2 Topology

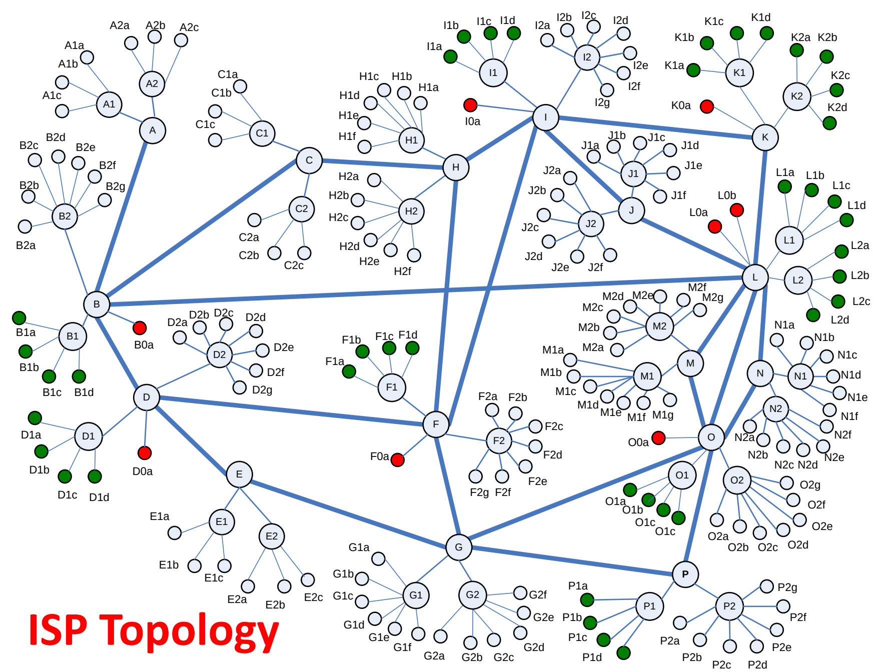

ISP Topology

Today's Topics

- Overview of our Past & Ongoing Research – with application to complex information networks, e.g., Internet, Clouds, Grids
- What is the problem and why is it hard?
- Four Approaches we are investigating:
  1. Sensitivity Analysis + Correlation Analysis & Clustering
  2. Combine Markov Models, Graph Analysis & Perturbation Analysis
  3. Anti-Optimization + Genetic Algorithm
  4. Measuring Key System Properties such as Critical Slowing Down
- Example of Sensitivity Analysis + Correlation Analysis & Clustering applied to a TCP/IP Network Model – closely related to the theme of this topical event: Measurement of Complex Networks

Past Research

Our Past Research: 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).
- 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).

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

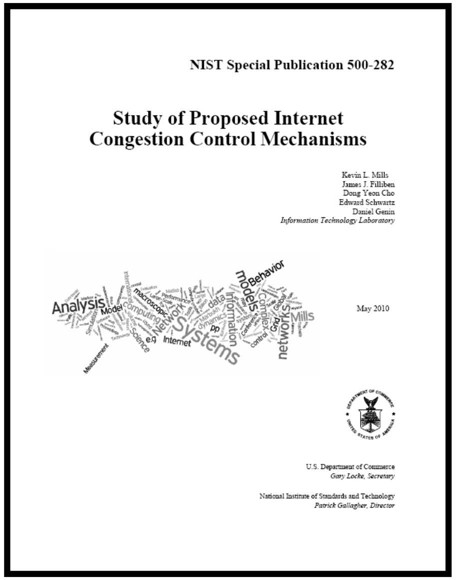

http://www.nist.gov/itl/antd/Congestion_Control_Study.cfm

National Institute of Standards and Technology

Ongoing Research

- **Our Ongoing & Planned Research**: How can we help to 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.

- **Ongoing**: investigating
  a. **Sensitivity Analysis + Correlation Analysis & Clustering**
  b. **Markov Chain Modeling + Cut-Set Analysis + Perturbation Analysis (MCM+CSA+PA)** (e.g., Dabrowski, Hunt and Morrison, "Improving the Efficiency of Markov Chain Analysis of Complex Distributed Systems", **NIST IR 7744**, 2010).
  c. **Anti-Optimization + Genetic Algorithm (AO+GA)**

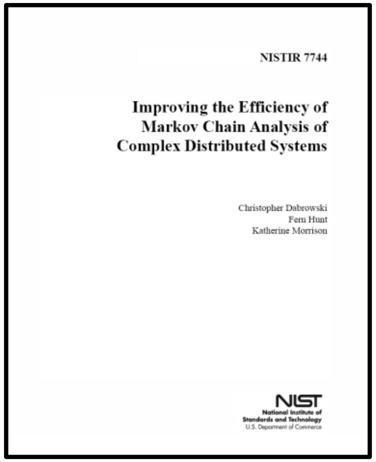
NISTIR 7744

Improving the Efficiency of
Markov Chain Analysis of
Complex Distributed Systems

Christopher Dabrowski
Fern Hunt
Katherine Morrison

NIST
National Institute of Standards and Technology
U.S. Department of Commerce

http://www.nist.gov/itl/antd/upload/NISTIR7744.pdf

- **Planned:** investigate run-time methods based on approaches that 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).

What is the Problem?

National Institute of Standards and Technology

➢ **Problem**: Given a complex information network (represented using a simulation model), how can one identify conditions that could cause global system behavior to degenerate, leading to costly system outages?

Koala Cloud Simulator

Why is it Hard? – Reason 1

National Institute of Standards and Technology

Determining causality is difficult – in a complex system, global behavior is not easily predictable, even if behavior of the components is understood completely

For example, unexpected collapse in the mitigation probability density function of job completion times in a computing grid was unexplainable without more detailed data and analysis.

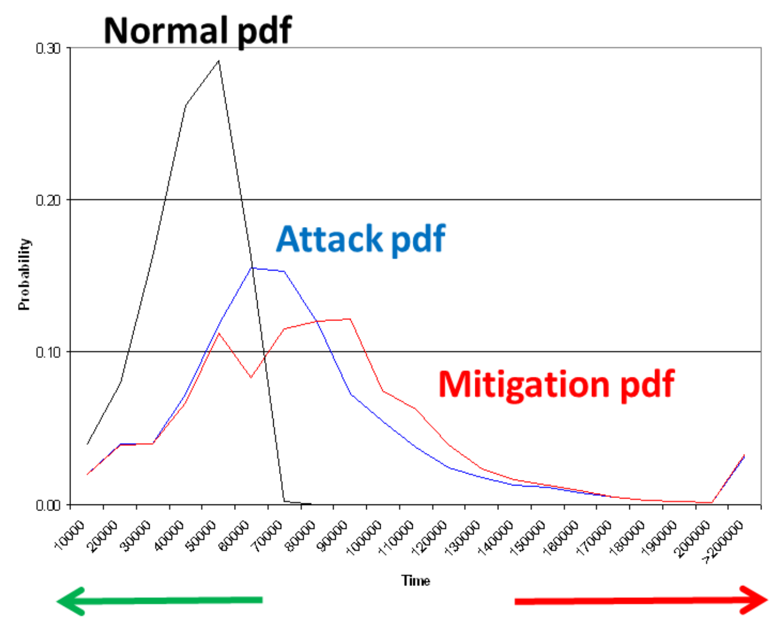

See: K. Mills and C. Dabrowski, "Investigating Global Behavior in Computing Grids", Self-Organizing Systems, Lecture Notes in Computer Science, Volume 4124 ISBN 978-3-540-37658-3, pp. 120-136.

Why is it Hard? – Reason 2

Size of the search space!!

$$y_1, \ldots, y_m = f(x_{1|[1,\ldots,k]}, \ldots, x_{n|[1,\ldots,k]})$$

Model Response Space **Model Parameter Space**

For example, the NIST ***Koala*** simulator of IaaS Clouds has about $n = 125$ parameters with average $k = 6.6$ values each, which leads to a model **parameter space** of ~**10^{100}** (note that the visible universe has ~10^{80} atoms) and the ***Koala*** response space ranges from $m = 8$ to $m = 200$, depending on the specific responses chosen for analysis (typically $m \approx 45$).

Sensitivity Analysis + CAC

- **Sensitivity Analysis**: Determine which parameters most significantly influence model behavior. Reduces parameter search space and identifies conditions under which alternate control algorithms should be compared.
- **Correlation Analysis & Clustering:** Determine response dimension of a model.

Use 2-level, orthogonal fractional factorial (OFF) experiment design and main effects & interaction analyses to identify significant model parameters

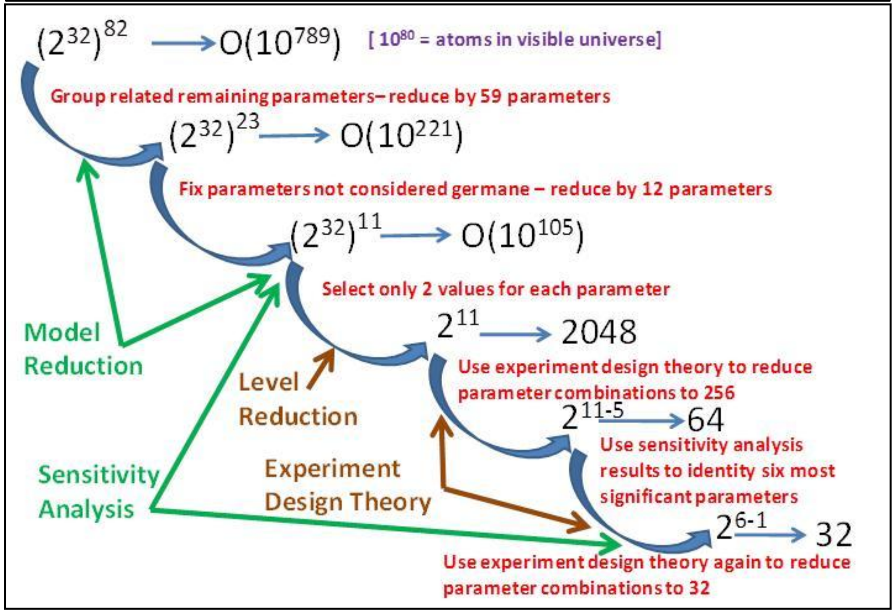

Use correlation analysis and clustering to identify unique behavior dimensions of your model

1. Compute correlation coefficient (r) for all response pairs
2. Examine frequency distribution for all $|r|$ to determine threshold for correlation pairs to retain; $|r| > 0.65$, here
3. Create clusters of mutually correlated pairs; each cluster represents one dimension
4. Select one response from each cluster to represent the dimension; we selected response with largest mean correlation that was not in another cluster*

Response Dimension	SA1-small (9 dimensions)	SA1-large (8 dimensions)	SA2-small (10 dimensions)	SA2-large (9 dimensions)
Cloud-wide Demand/Supply Ratio	y1, y2, **y3**, y5, y6, y8, y9, y10, y13, y23, y24, y25, y29, y30, y32, y34, y36, y38	y1, y2, **y3**, y5, y6, y7, y8, y9, y10, y13, y23, y34, y25, y29, y30, y32, y33, y34, y36, y38	y1, **y2**, y3, y5, y6, y8, y9, y10, y11, y13, y14, y15, y23, y24, y25, y38	y1, y2, y3, y5, y6, y8, y9, **y23**, y24, y25, y38
Cloud-wide Resource Usage	y10, y11, y12, y13, y14, **y15**	y10, y11, y12, y13, y14, **y15**	**y10**, y11, y12, y13, y14, y15	**y10**, y11, y12, y13, y14, y15
Variance in Cluster Load	y16, y17, y18, y19,y20, y21, **y26**, y27	y16, y17, y18, y19,y20, y21, **y26**, y27	y16, y18, y19, y20, y21, y26, **y27** **y17** (Mem. Util)	y16, y17, y18, **y19**,y20, y21, y26, y27
Mix of VM Types	y34, **y35** (WS) **y31** (MS)	**y31** (MS)	y12, y14, y15, y30, y31, y33, y34, y35, **y36**	y14, y15, y30, **y31**, y33, y34, y35 y15, **y36** (DS)
Number of VMs	y29, **y37**	**y37**	y29, **y37**	**y29**
User Arrival Rate	**y4**	**y4**	**y4**	**y4**, y37
Reallocation Rate	**y7**, y22	y7, **y22**	**y7** (cluster) **y22** (node)	y7, **y22**
Variance in Choice of Cluster	**y28**	**y28**	**y28**	**y28**

See: Mills, Filliben and Dabrowski, "An Efficient Sensitivity Analysis Method for Large Cloud Simulations", *Proceedings of the 4th International Cloud Computing Conference*, IEEE (2011) **and** Mills and Filliben, "Comparison of Two Dimension-Reduction Methods for Network Simulation Models", *Journal of NIST Research* **116-5**, 771-783 (2011).

Cut-Set + Perturbation Analysis

Using simulated failure scenarios in a Markov chain model to predict failures in a Cloud

Example: Markov simulation and perturbation of *__a minimal s-t cut set__* of a Markov chain graph:

- **Corresponds to software failure scenario involving multiple faults/attacks.**
- **Simulation identifies threshold beyond which increased failure incidence causes drastic performance collapse**

→ ***__Verified in target system being modeled (i.e., Koala, a large-scale simulation of a Cloud)__***

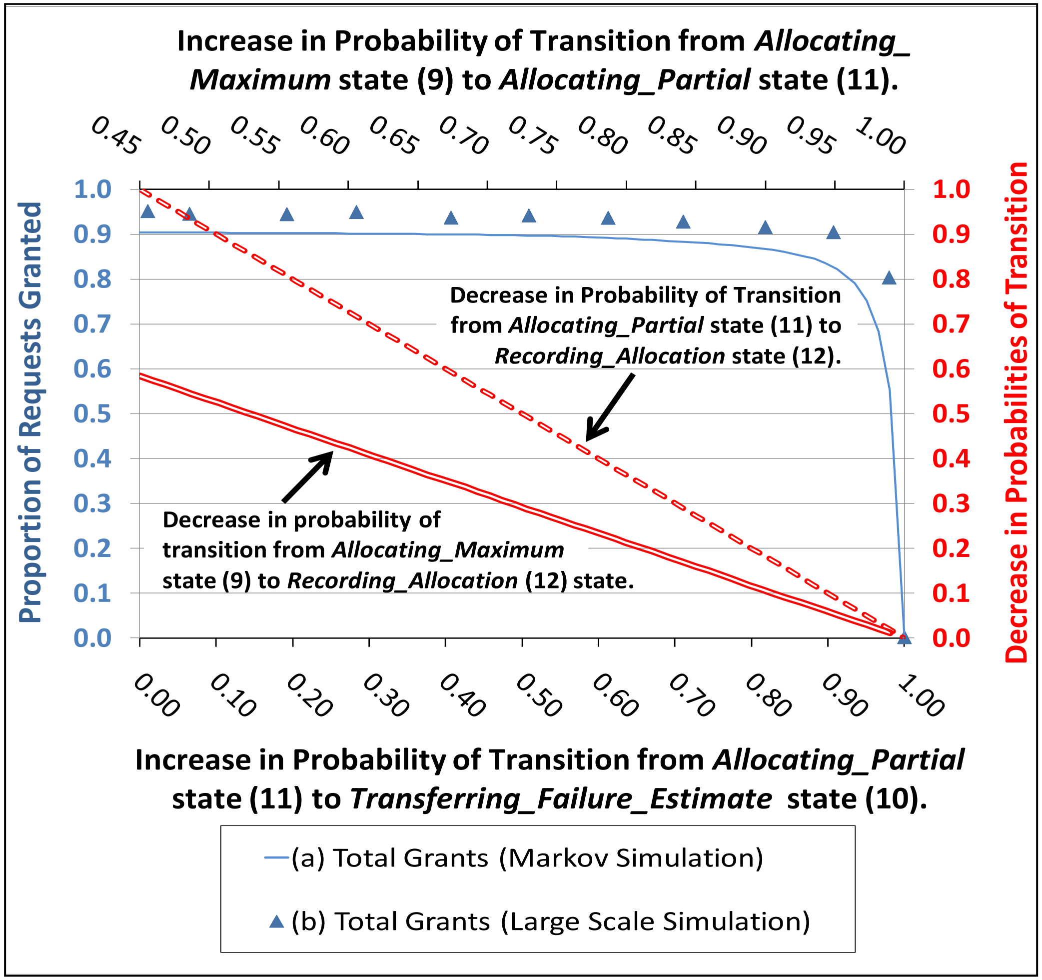

Anti-Opt. + Genetic Algorithm

MULTIDIMENSIONAL ANALYSIS TECHNIQUES

Principal Components Analysis, Clustering, …

Growing Collection of Tuples:

{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}
…
{Generation, Individual, Fitness, Parameter 1 value,….Parameter N value}

GENETIC ALGORITHM

Recombination* & *Mutation

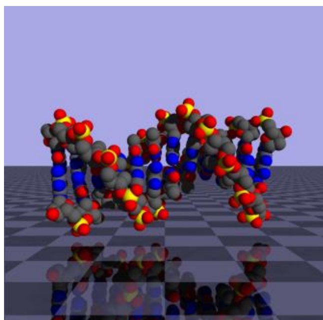

***Selection* based on Anti-Fitness**

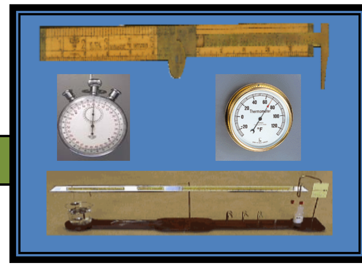

Anti-Fitness Reports

MODEL SIMULATORS

Parallel Execution of Model Simulators

List of parameters and for each parameter a MIN, MAX and precision.

Model Parameter Specifications

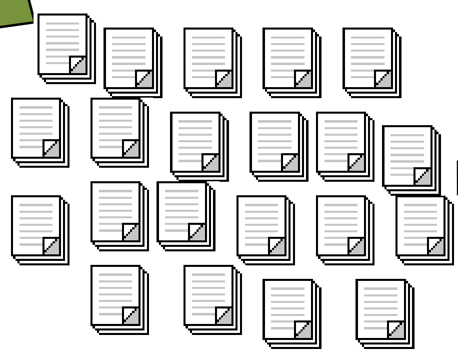

Population of Model Parameterizations

Critical Slowing Down

A simple univariate example predicting power grid blackout in a human engineered system*

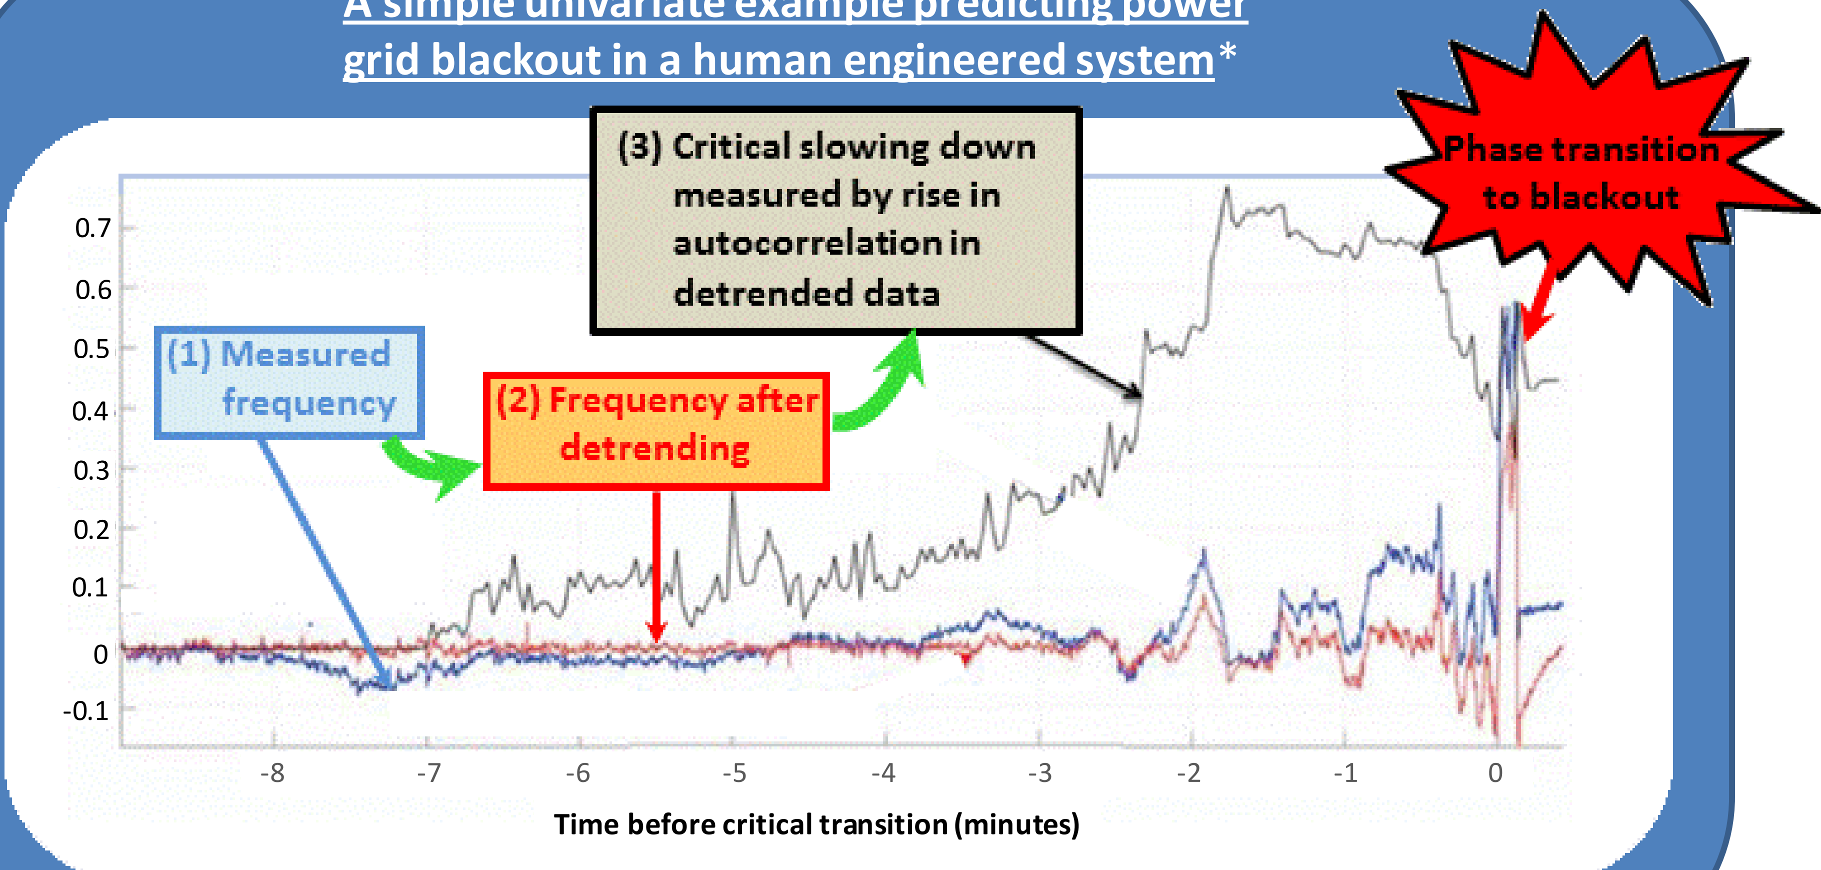

*From P. Hines, E. Cotilla-Sanchez, and S. Blumsack. Topological Models and Critical Slowing Down: Two Approaches to Power System Risk Analysis. Proceedings of the 44th Hawaii Conference on System Sciences. IEEE Computer Society, Washington, DC, USA, pp. 1-10.

SA + CAC Example

Sensitivity Analysis + Correlation Analysis & Clustering applied to a TCP/IP Network Model

(using an 11-parameter subset of a 20-parameter model*)

Questions: (1) What responses characterize system behavior?
(2) What factors drive system behavior?

$$y_1, \ldots, y_m \quad = \quad f(\, x_{1|[1,\ldots,k]}, \ldots, x_{n|[1,\ldots,k]} \,)$$

Model Response Space — **Model Parameter Space**

In the example that follows, m = 22, n = 11 and k = 2

The approach is general: as we have demonstrated on a TCP/IP model with m = 45, n = 20 and k = 2 and on a Cloud Computing model with m = 38, n = 11 and k = 2 and with m = 45, n = 20 and k = 2

*For a discussion of the full 20-parameter TCP/IP model see: Mills, Schwartz and Yuan, "How to Model a TCP/IP Network using only 20 Parameters", *Proceedings of the Winter Simulation Conference* (2010).

Simulated Topology

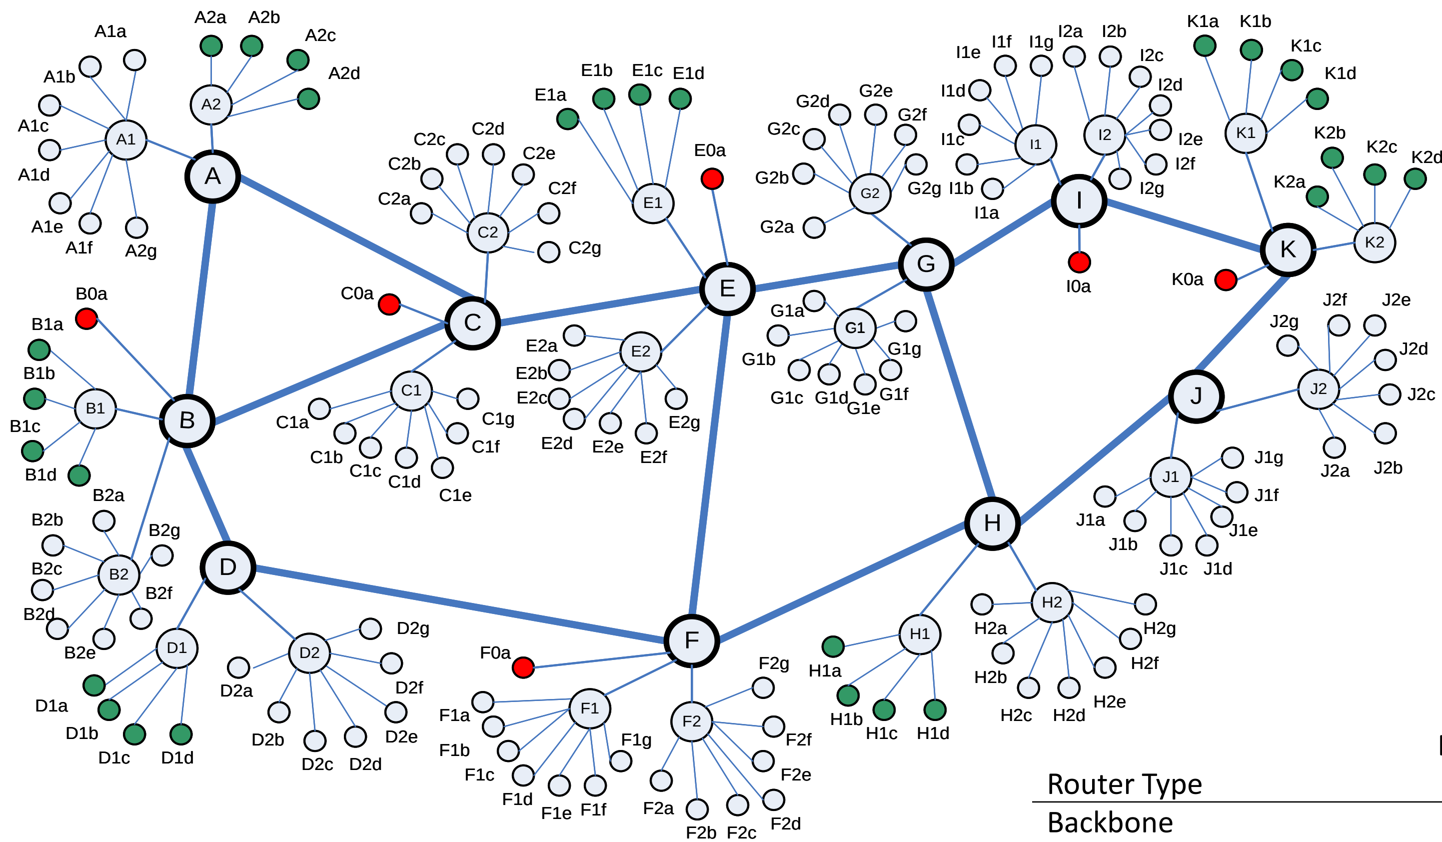

Backbone of Topology based on Internet 2 (including propagation delays & routing paths)

Router Type	Base Speed p/ms
Backbone	400
POP	100
Typical Access	10
Fast Access	20
Directly Connected Access	100

11-Parameter Subset

	Factor	Name	Plus (+1) Setting	Minus (-1) Setting
Network Factors	x1	**Propagation Delay Multiplier**	2	1
	x2 [1]	**Network Speed Multiplier**	1	2
	x3	**Buffer Sizing Algorithm**	*RTTxC*	*RTTxC*/SQRT(*n*)
User Factors	x4	**Average File Size**	100 packets	50 packets
	x5	**Average Think Time**	5000 ms	2000 ms
	x6 [1]	**Probability User Downloads 10x File**	0.01	0.02
Source & Receiver Factors	x7 [1]	**Probability of a Fast Host Connection**	0.2	0.4
	x8	**Multiplier for Number of Sources & Receivers per Access Router**	3	2
	x9	**Distribution Pattern of Sources**	P2P	WEB
	x10	**Distribution Pattern of Receivers**	P2P	WEB
Protocol Factors	x11	**Initial TCP Slow-Start Threshold**	1.07×10^9 packets	43 packets

[1] Unfortunately, we coded these settings backwards from the usual convention of higher value for the Plus setting, so care must be taken when interpreting the results for these factors – mainly the network speed factor. Sorry.

22-Response Subset

16 Responses Characterizing Macroscopic Network Behavior

Response	Name	Definition
y1	# Sending Flows	Active Flows – flows attempting to transfer data
y2	% Sources Sending	Proportion of potential flows that were active: Active Flows/All Sources
y3	# Packets Entering	Data packets entering the network per measurement interval
y4	# Packets Exiting	Data packets leaving the network per measurement interval
y5	Loss Rate	Loss Rate: y4/(y3+y4)
y6	# Flow Completions	Flows Completed per measurement interval
y7	Flow Completion Rate	Flow-Completion Rate: y6/(y6+y1)
y8	# Connection Failures	Connection Failures per measurement interval
y9	Connection Failure Rate	Connection-Failure Rate: y8/(y8+y1)
y10	Retransmission Rate	Retransmission Rate
y11	Average Congestion Window	Congestion Window per Flow
y12	# Window Increases	Window Increases per Flow per measurement interval
y13	# NAKs	Negative Acknowledgments per Flow per measurement interval
y14	# Timeouts	Timeouts per Flow per measurement interval
y15	Average Round-Trip Time	Smoothed Round-Trip Time
y16	Queuing Delay	Relative queuing delay: y15/(x1x41)

Global Behavior

+ 6 Responses Characterizing Instantaneous Throughput for Active Flows by Class

Response	Definition
y17	Average Throughput for Active DD Flows
y18	Average Throughput for Active DF Flows
y19	Average Throughput for Active DN Flows
y20	Average Throughput for Active FF Flows
y21	Average Throughput for Active FN Flows
y22	Average Throughput for Active NN Flows

User Experience

OFF Experiment Design

X1	X2	X3	X4	X5	X6	X7	X8	X9	X10	X11
-1	-1	-1	-1	-1	-1	-1	+1	-1	+1	+1
+1	-1	-1	-1	-1	-1	-1	-1	+1	+1	-1
-1	+1	-1	-1	-1	-1	-1	-1	+1	-1	+1
+1	+1	-1	-1	-1	-1	-1	+1	-1	-1	-1
-1	-1	+1	-1	-1	-1	+1	-1	-1	+1	+1
+1	-1	+1	-1	-1	-1	+1	+1	+1	+1	-1
-1	+1	+1	-1	-1	-1	+1	+1	+1	-1	+1
+1	+1	+1	-1	-1	-1	+1	-1	-1	-1	-1
-1	-1	-1	+1	-1	-1	+1	-1	-1	-1	-1
+1	-1	-1	+1	-1	-1	+1	+1	+1	-1	+1
-1	+1	-1	+1	-1	-1	+1	+1	+1	+1	-1
+1	+1	-1	+1	-1	-1	+1	-1	-1	+1	+1
-1	-1	+1	+1	-1	-1	-1	+1	-1	-1	-1
+1	-1	+1	+1	-1	-1	-1	-1	+1	-1	+1
-1	+1	+1	+1	-1	-1	-1	-1	+1	+1	-1
+1	+1	+1	+1	-1	-1	-1	+1	-1	+1	+1
-1	-1	-1	-1	+1	-1	+1	+1	-1	-1	-1
+1	-1	-1	-1	+1	-1	+1	-1	+1	-1	+1
-1	+1	-1	-1	+1	-1	+1	-1	+1	+1	-1
+1	+1	-1	-1	+1	-1	+1	+1	-1	+1	+1
-1	-1	+1	-1	+1	-1	-1	-1	-1	-1	-1
+1	-1	+1	-1	+1	-1	-1	+1	+1	-1	+1
-1	+1	+1	-1	+1	-1	-1	+1	+1	+1	-1
+1	+1	+1	-1	+1	-1	-1	-1	-1	+1	+1
-1	-1	-1	+1	+1	-1	-1	-1	-1	+1	+1
+1	-1	-1	+1	+1	-1	-1	+1	+1	+1	-1
-1	+1	-1	+1	+1	-1	-1	+1	+1	-1	+1
+1	+1	-1	+1	+1	-1	-1	-1	-1	-1	-1
-1	-1	+1	+1	+1	-1	+1	+1	-1	+1	+1
+1	-1	+1	+1	+1	-1	+1	-1	+1	+1	-1
-1	+1	+1	+1	+1	-1	+1	-1	+1	-1	+1
+1	+1	+1	+1	+1	-1	+1	+1	-1	-1	-1
-1	-1	-1	-1	-1	+1	-1	+1	+1	-1	-1
+1	-1	-1	-1	-1	+1	-1	-1	-1	-1	+1
-1	+1	-1	-1	-1	+1	-1	-1	-1	+1	-1
+1	+1	-1	-1	-1	+1	-1	+1	+1	+1	+1
-1	-1	+1	-1	-1	+1	+1	-1	+1	-1	-1
+1	-1	+1	-1	-1	+1	+1	+1	-1	-1	+1
-1	+1	+1	-1	-1	+1	+1	+1	-1	+1	-1
+1	+1	+1	-1	-1	+1	+1	-1	+1	+1	+1
-1	-1	-1	+1	-1	+1	+1	-1	+1	+1	+1
+1	-1	-1	+1	-1	+1	+1	+1	-1	+1	-1
-1	+1	-1	+1	-1	+1	+1	+1	-1	-1	+1
+1	+1	-1	+1	-1	+1	+1	-1	+1	-1	-1
-1	-1	+1	+1	-1	+1	-1	+1	+1	+1	+1
+1	-1	+1	+1	-1	+1	-1	-1	-1	+1	-1
-1	+1	+1	+1	-1	+1	-1	-1	-1	-1	+1
+1	+1	+1	+1	-1	+1	-1	+1	+1	-1	-1
-1	-1	-1	-1	+1	+1	+1	+1	+1	+1	+1
+1	-1	-1	-1	+1	+1	+1	-1	-1	+1	-1
-1	+1	-1	-1	+1	+1	+1	-1	-1	-1	+1
+1	+1	-1	-1	+1	+1	+1	+1	+1	-1	-1
-1	-1	+1	-1	+1	+1	-1	-1	+1	+1	+1
+1	-1	+1	-1	+1	+1	-1	+1	-1	+1	-1
-1	+1	+1	-1	+1	+1	-1	+1	-1	-1	+1
+1	+1	+1	-1	+1	+1	-1	-1	+1	-1	-1
-1	-1	-1	+1	+1	+1	-1	-1	+1	-1	-1
+1	-1	-1	+1	+1	+1	-1	+1	-1	-1	+1
-1	+1	-1	+1	+1	+1	-1	+1	-1	+1	-1
+1	+1	-1	+1	+1	+1	-1	-1	+1	+1	+1
-1	-1	+1	+1	+1	+1	+1	+1	+1	-1	-1

Template for a 2-Level 2^{11-5} Orthogonal Fractional Factorial (OFF) experiment design specifying the combination of parameter level settings for 64 simulation runs

Balance: 32 (-) — 32 (+) along X_i

Orthogonality: X_j (-, +) by X_i (-, +): 8 8 / 8 8

<u>OFF Benefit #1</u>: Superior Coverage & Robustness as compared with 1-Factor-at-a-Time Designs

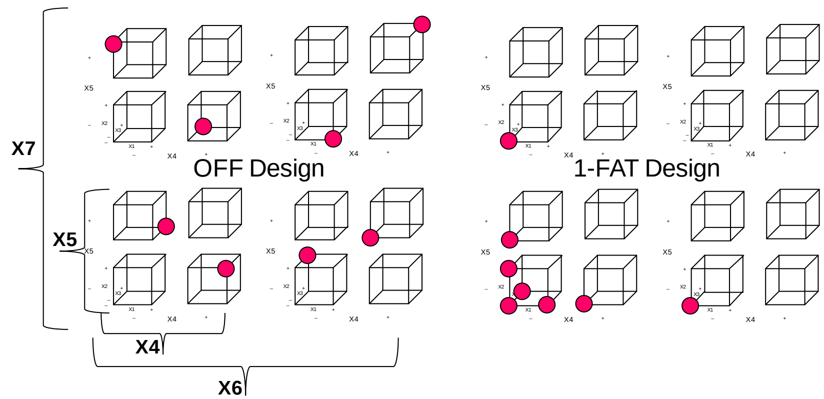

<u>OFF Design Benefit #2</u>: Minimizes Variation in Effect Estimates

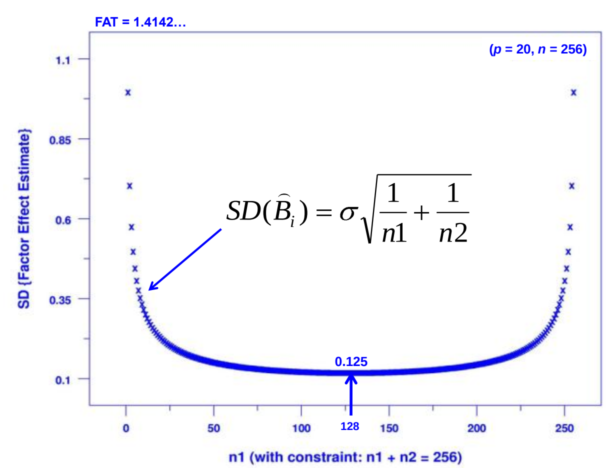

Pairwise Corr. Matrix

Each response observed under the same 64 combinations of parameter settings

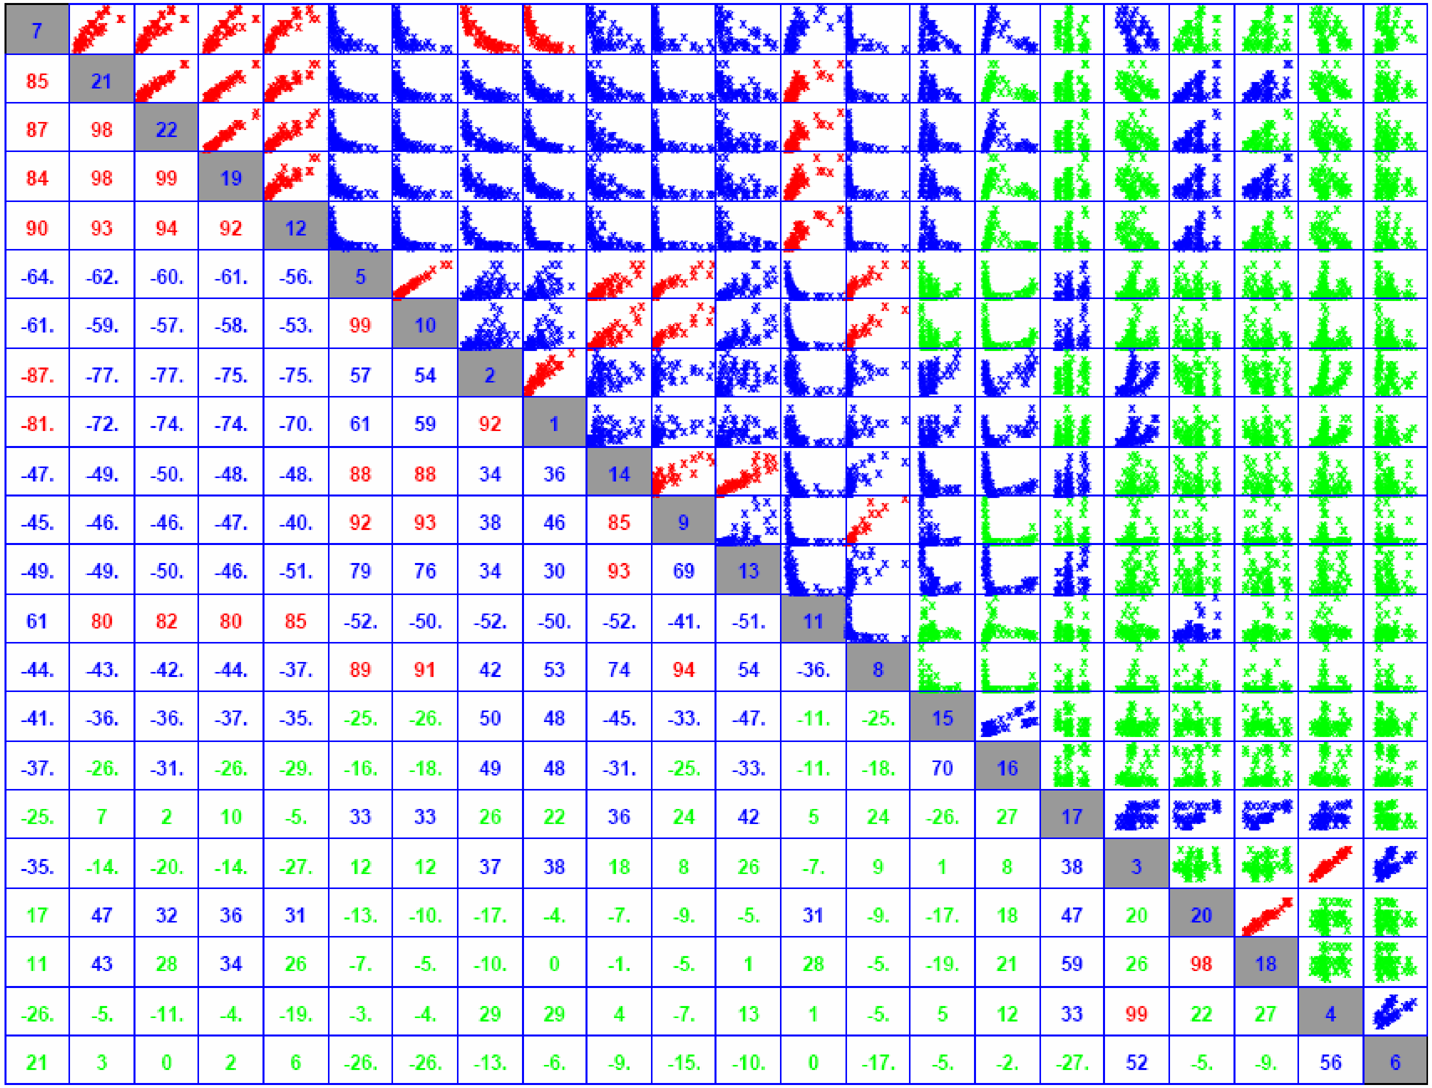

Red 80 ≥ |r|x100 ≤ 100 **Blue** 30 ≥ |r|x100 < 80 **Green** |r|x100 < 30

Histogram of |*r*|

Select a threshold for |*r*| such that correlations above that threshold will be further considered

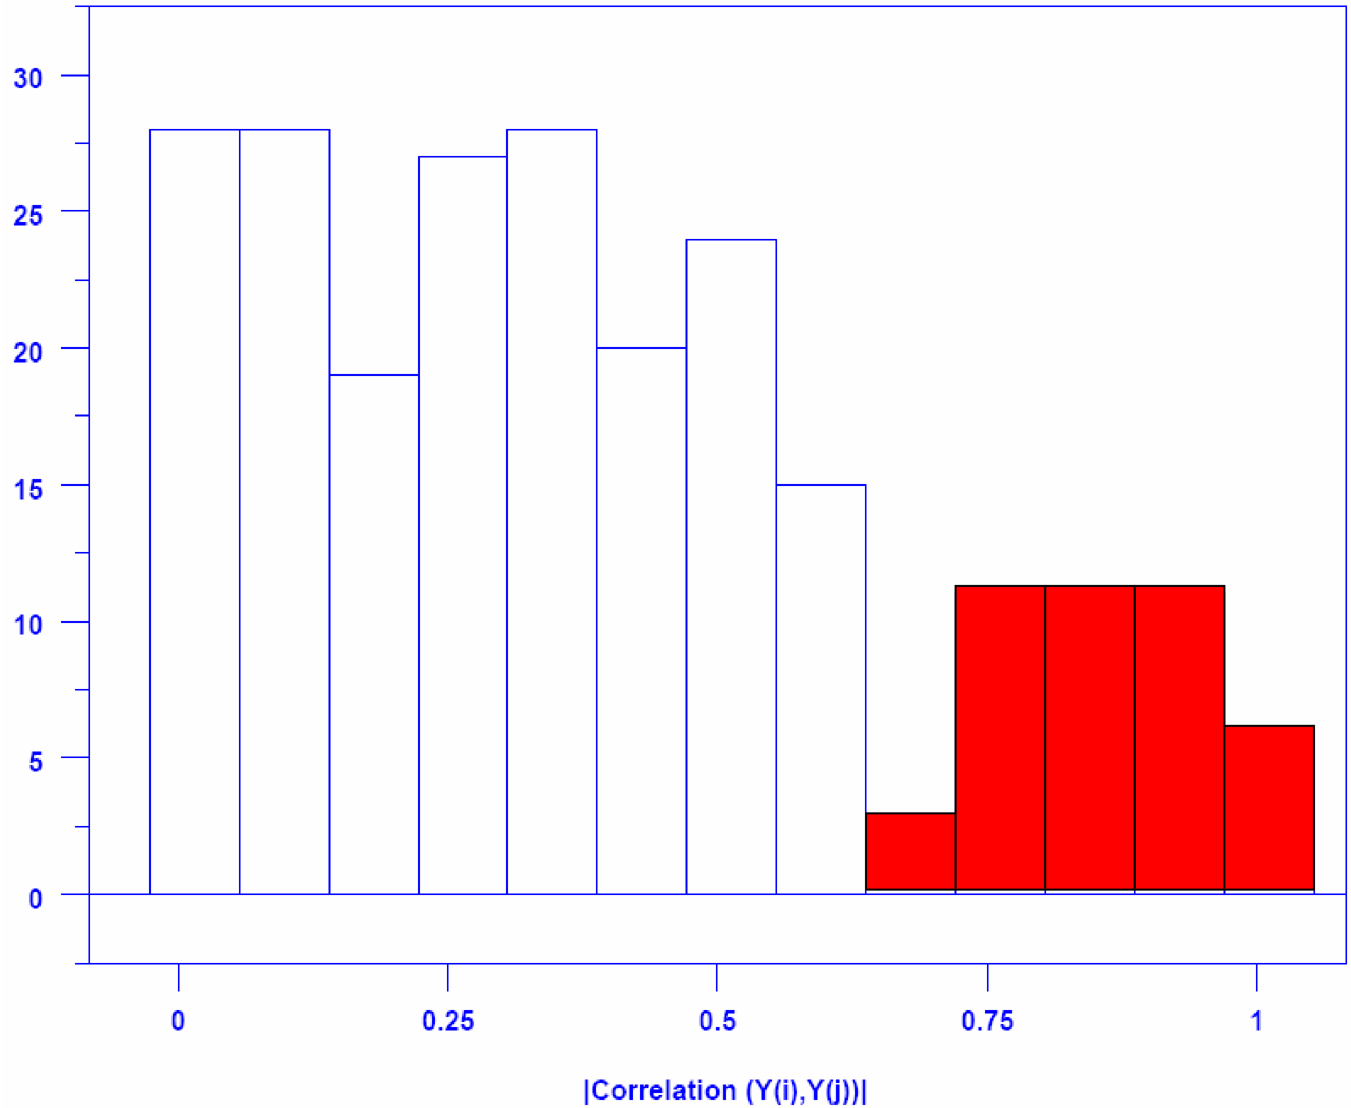

Response Index-Index Plot

where $|r_{i,j}| > 0.65$ are clustered into mutual correlations

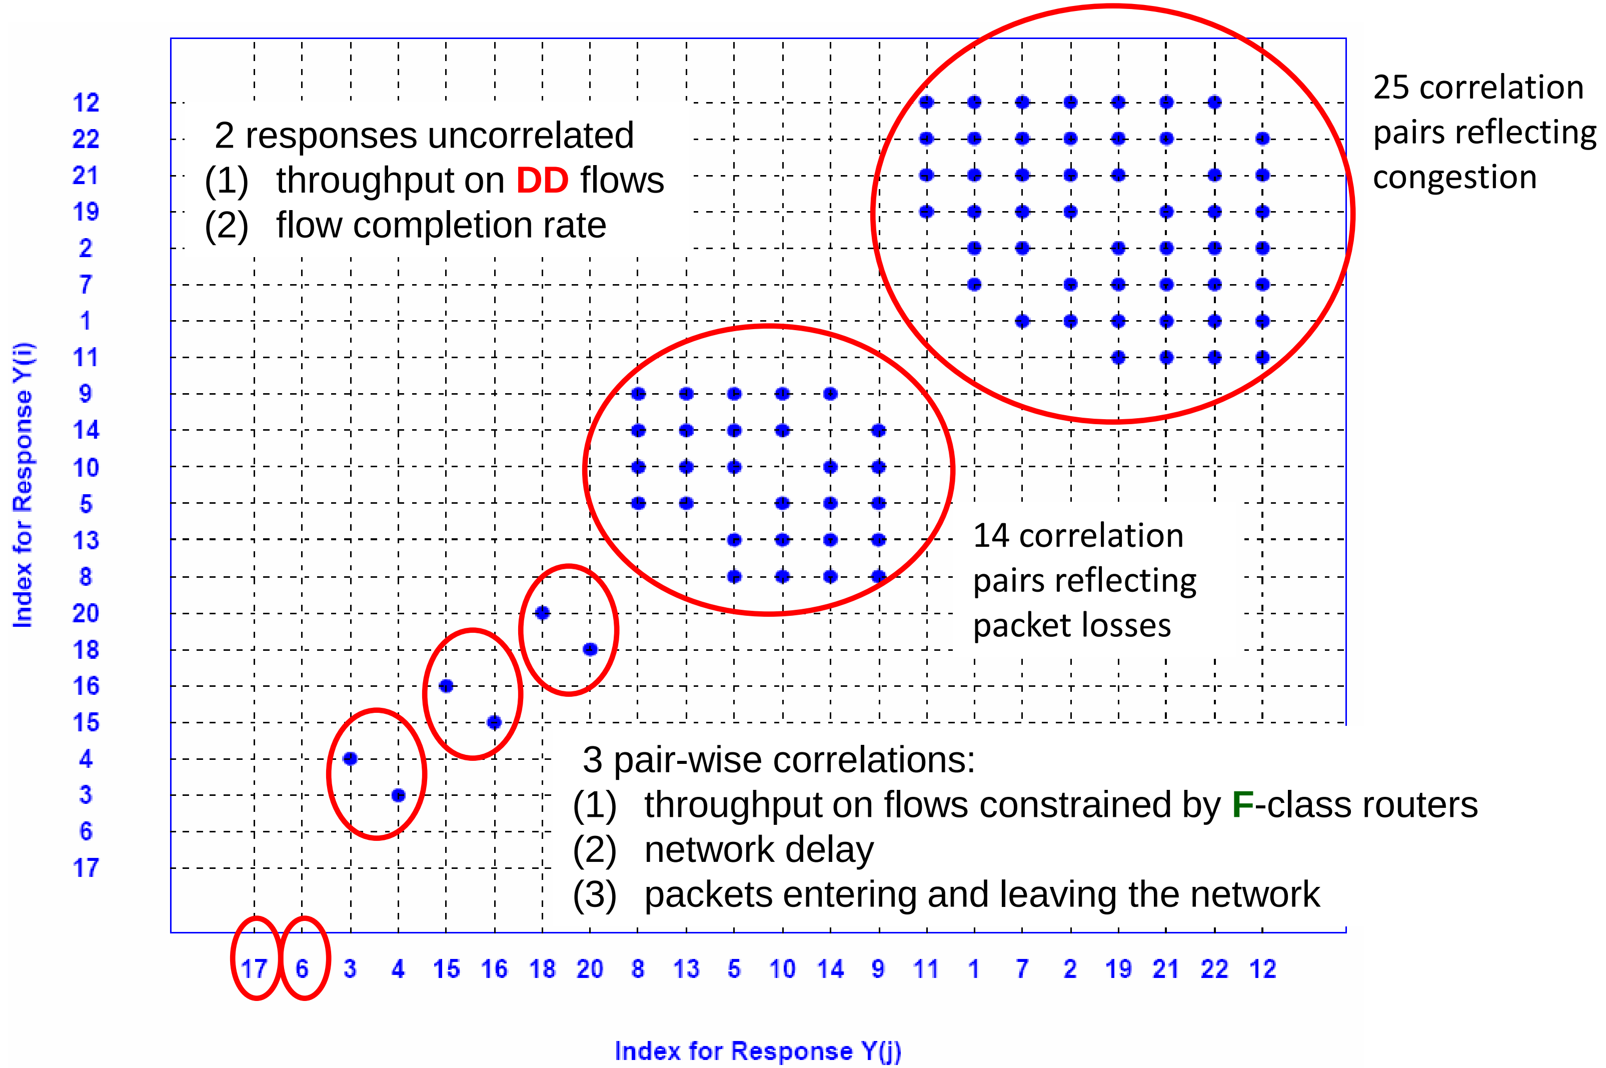

7-D Model Behavior

Question 1: What responses characterize system behavior?

Answer: As shown below, measuring only 7 of the 22 responses suffices to characterize behavior

Representative Response	*Dimension* and a Characterizing Measurement
y4	*D1 - network throughput in packets/sec* measured by average number of packets output per measurement interval
y6	*D2 - network throughput in flows/sec* measured by average number of flows completed per measurement interval
y10	*D3 - packet loss* measured by average retransmission rate
y15	*D4 - network delay* measured by average smoothed round-trip time
y17	*D5 - throughput in packets/sec for the most advantaged users* measured by average instantaneous throughput for **DD** flows
y20	*D6 - throughput in packets/sec for 2^{nd} most advantaged users* measured by average instantaneous throughput for **FF** flows
y22	*D7 - network congestion* measured by average instantaneous throughput for **NN** flows

Main Effects Analyses

For each response, compare mean at 32 Plus settings with mean at 32 Minus settings and conduct *t*-test to determine statistical significance (response here is TP on **NN** flows)

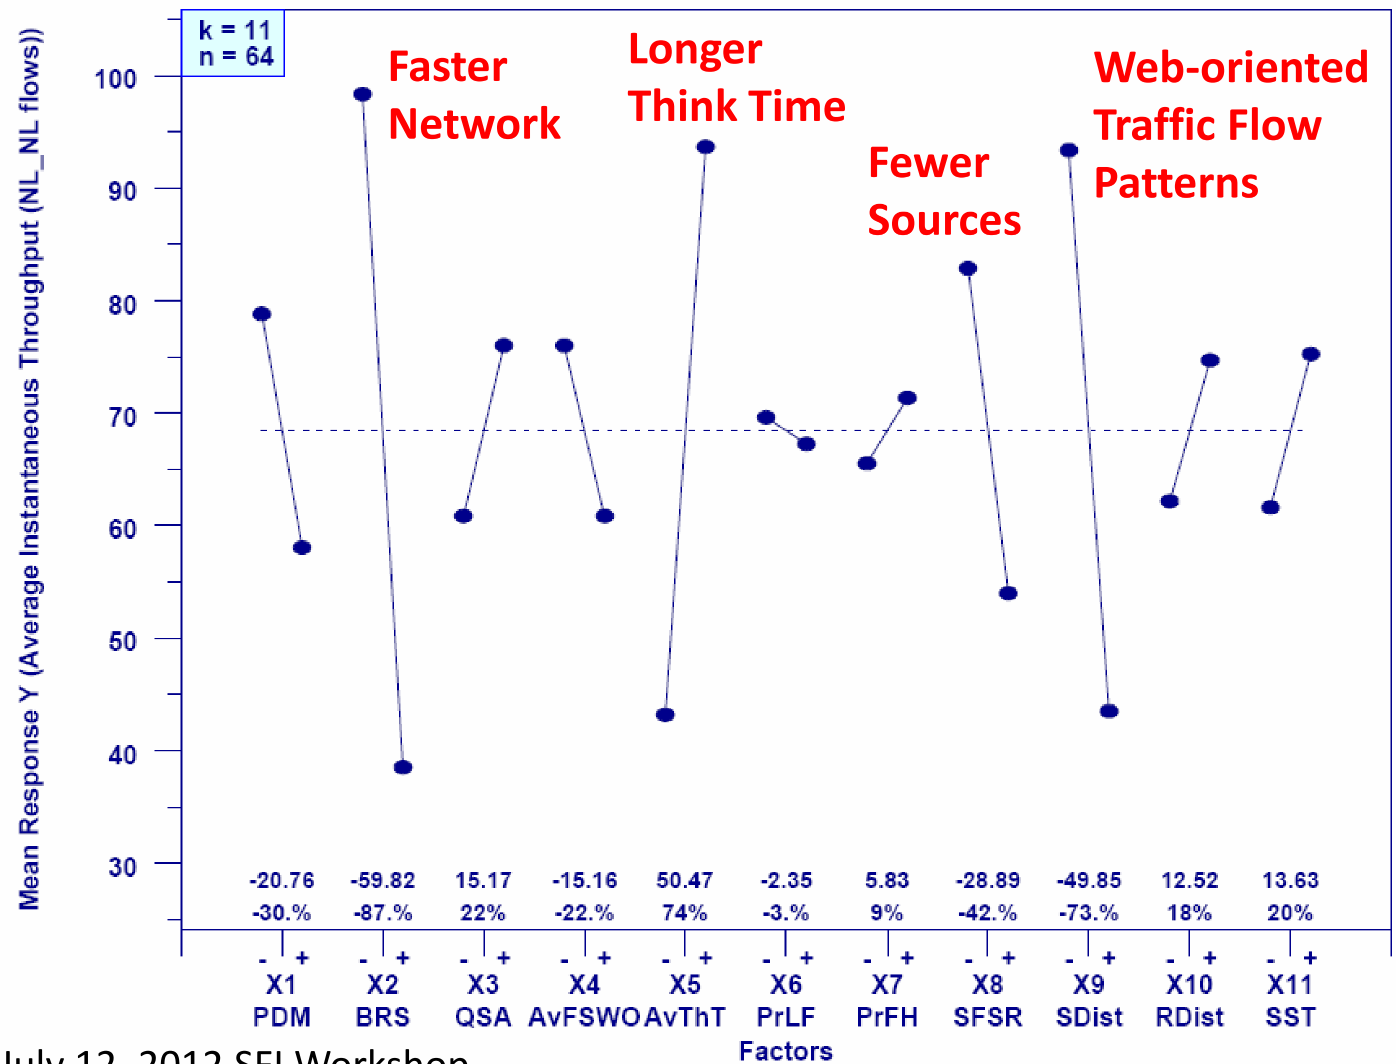

Four factors lead to higher throughput for typical users; inverse settings lead to lower throughput

The other seven factors have little effect on this response

Summary of Main Effects

		x1	x2	x3	x4	x5	x6	x7	x8	x9	x10	x11
y4	*Packet TP*		-**		+**	-**			+*	+*		
y6	*Flow TP*		-**		-**	-**			+*	+*		
y10	*Packet Loss*		+**	-**	+*	-*			+*	+*		
y15	*Delay*	+**	+*	+**	+*	-*			+*	+*		
y17	***DD*** *TP*	-**			+**							
y20	***FF*** *TP*	-**	-**		+*	+*			-*	+**		
y22	***NN*** *TP*	-*	-**			+*			-*	-**		
Ψ	*Net Effect*	50	79	29	64	57	0	0	43	50	0	0

- * $p < 0.05$ and ** $p < 0.01$
- \- means minus value caused response increase
- \+ means plus value cause response increase
- please remember that network speed x2 was miscoded, so – means higher speed and + lower

This chart reveals much about model behavior

$$\Psi = 100\ (|\{y \mid p < 0.01\}| + \tfrac{1}{2}\ |\{y \mid p < 0.05\}|)\ /\ |\{y\}|$$

6 Factors Drive Behavior

Question 2: What factors drive system behavior?

Answer: Network speed mainly followed by file size, then by user duty cycle, propagation delay and source distribution, and finally by number of sources.

Relative Influence	Factors
1	**Network Speed**
2	File Size
3	User Duty Cycle Propagation Delay Distribution of Sources
4	Number of Sources
5	Buffer Sizing
6	Proportion of Hosts with Fast Attachments Probability of 10x File Sizes Receiver Distribution Initial Slow-Start Threshold

The information generated here can be used in two ways:

1. Compare model behavior to experiences of operational networks, as a validation step.

2. Select parameter combinations to explore when asking what if questions, such as what if TCP were replaced by any of 7 competing congestion control algorithms?

For more information see: Mills, Filliben, Cho and Schwartz, "Predicting Macroscopic Dynamics in Large Distributed Systems", *Proceedings of ASME* (2011)

Brief Recap

- Reviewed our Past & Ongoing Research – with application to complex information networks, e.g., Internet, Clouds, Grids

- Defined the problem underlying our ongoing research and identified two reasons why the problem is difficult

- Described 4 approaches to address the problem:
  1. Sensitivity Analysis + Correlation Analysis & Clustering
  2. Combine Markov Models, Graph Analysis & Perturbation Analysis
  3. Anti-Optimization + Genetic Algorithm
  4. Measuring Key System Properties such as Critical Slowing Down

- Discussed an example of Sensitivity Analysis + Correlation Analysis & Clustering applied to a TCP/IP network model

Questions?

Suggestions?

Ideas?

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

Contact information about Information Visualization:
sressler@nist.gov

For more information see: http://www.nist.gov/itl/antd/emergent_behavior.cfm
and/or http://www.nist.gov/itl/cloud/index.cfm