Predicting Global Failure Regimes in Complex Information Systems

Kevin Mills, Jim Filliben & Chris Dabrowski

DoE COMBINE Workshop September 19, 2012

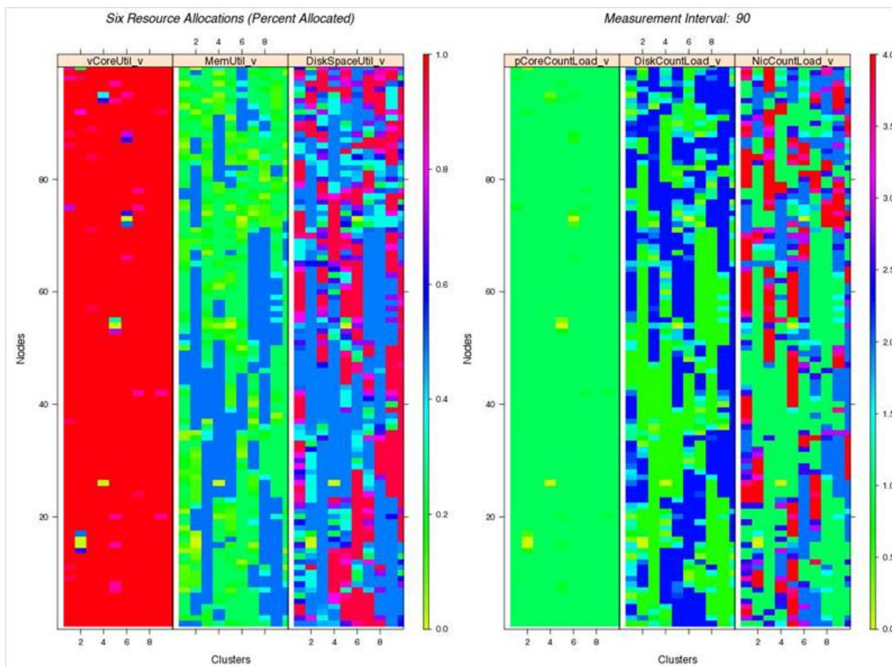

y10 y11 y12 y13 y14 y15

Koala Information Visualizations by Sandy Ressler

(see http://math.nist.gov/~SRessler/cloudviz.html for animations and more)

Today's Topics in 12 Mins.

NIST National Institute of Standards and Technology

- Overview of Our Past & Ongoing Research – with application to complex information systems, e.g., Internet, Clouds, Grids

- What is the problem?

- Why is it hard?

- Four Approaches we are investigating:
  1. Combine Markov Models, Graph Analysis & Perturbation Analysis
  2. Sensitivity Analysis + Correlation Analysis & Clustering
  3. Anti-Optimization + Genetic Algorithm
  4. Measuring Key System Properties Such as Critical Slowing Down

Past Research

NIST National Institute of Standards and Technology

Past ITL 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

Using a single process, simulated > 10^5 simultaneous TCP flows in nationwide topologies

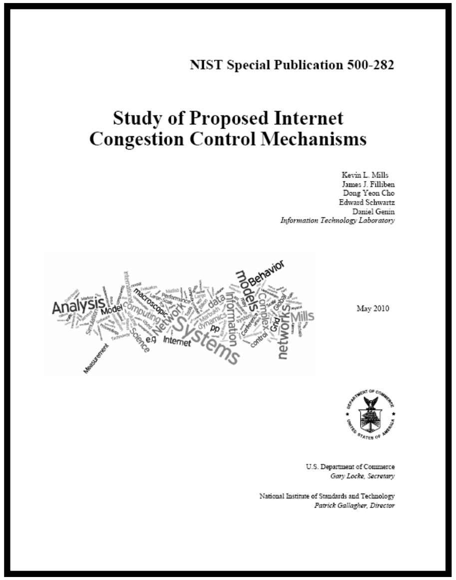
NIST Special Publication 500-282

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

National Institute of Standards and Technology

Ongoing Research

- **Ongoing & Planned ITL 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. **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).
  b. **Sensitivity Analysis + Correlation Analysis & Clustering**
  c. **Anti-Optimization + Genetic Algorithm (AO+GA)**

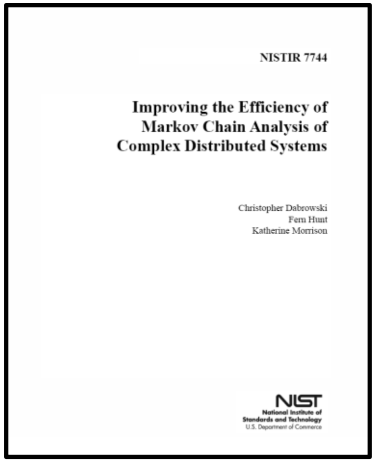
NISTIR 7744

Improving the Efficiency of
Markov Chain Analysis of
Complex Distributed Systems

Christopher Dabrowski
Fern Hunt
Katherine Morrison

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?

- **Problem**: Given a complex information system (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

Size of the search space!!

$$\underbrace{y_1, \ldots, y_m}_{\textbf{Model Response Space}} = \underbrace{f(\, x_{1|[1,\ldots,k]}, \ldots, x_{n|[1,\ldots,k]} \,)}_{\textbf{Model Parameter Space}}$$

For example, the NIST ***Koala*** simulator of IaaS Clouds has about $n = 132$ parameters with average $k = 6$ values each, which leads to a model **parameter space** of **~10^{102}** (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 $m \approx 45$).

Why is it Hard? – Reason 2

Determining causality is hard when only global system behavior is observable.
(in a complex system, global behavior cannot always be understood, even if behavior of components is completely understood)

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.

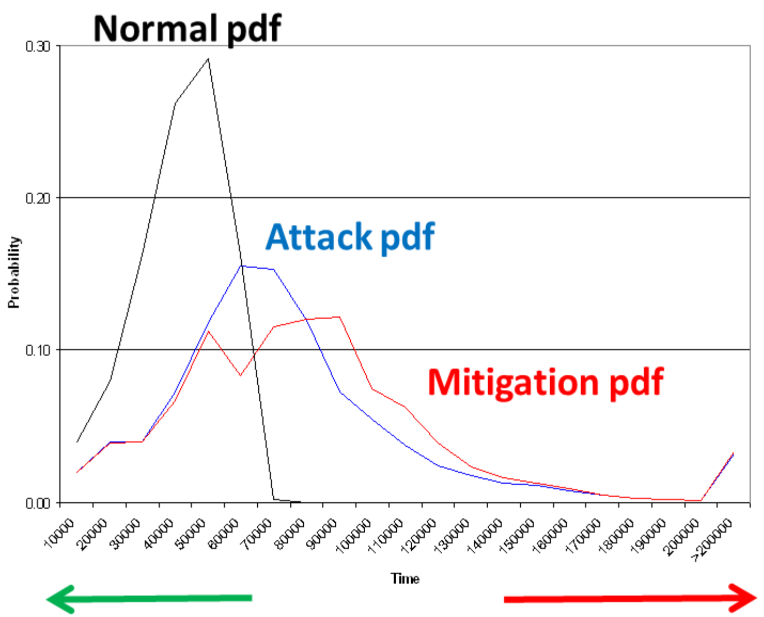

Cut-Set + Perturbation Analysis

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)__***

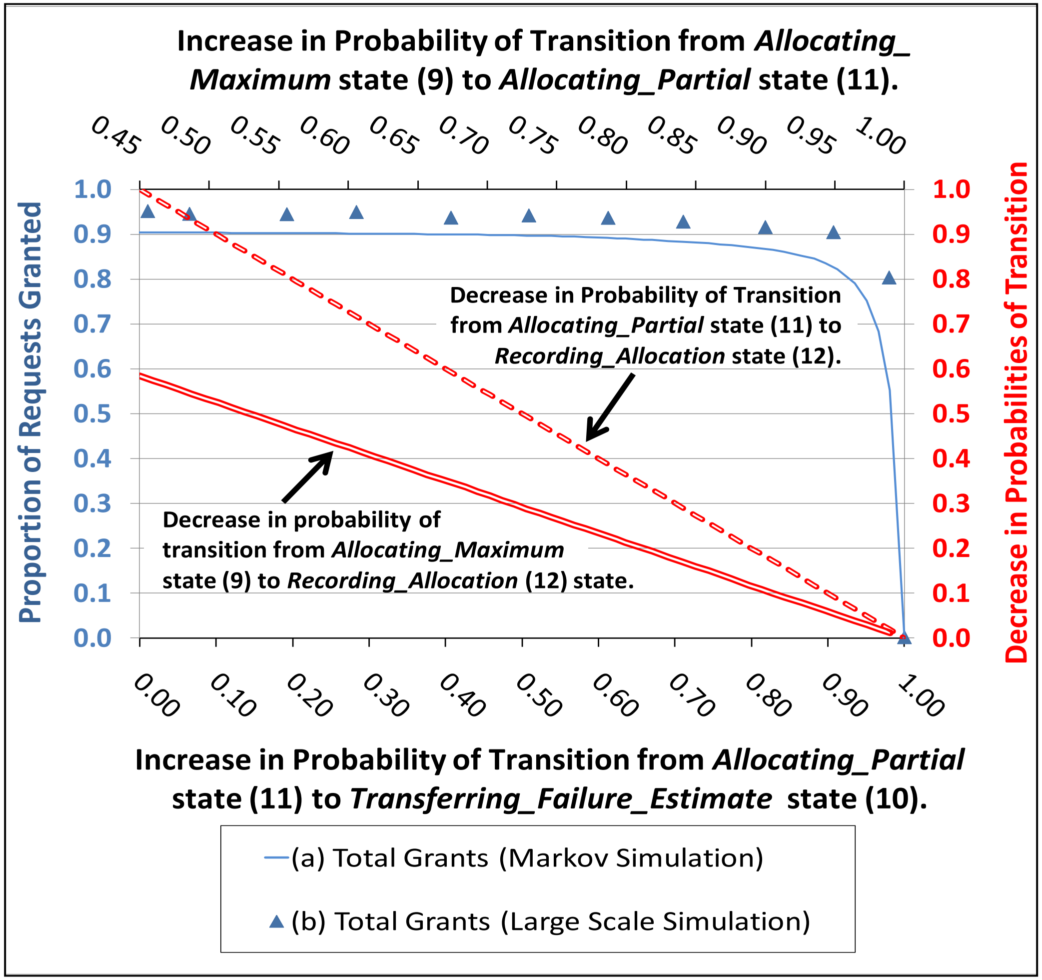

Sensitivity Analysis + CAC

- **Sensitivity Analysis**: Determine which parameters most significantly influence model behavior and what response dimensions the model exhibits. Allows reduction parameter search space and identifies model responses that must be analyzed.
- **Correlation Analysis & Cluster:** Determine response dimensions of a model

Use 2-level, orthogonal fractional factorial (OFF) experiment design to identify the most significant parameters of your model

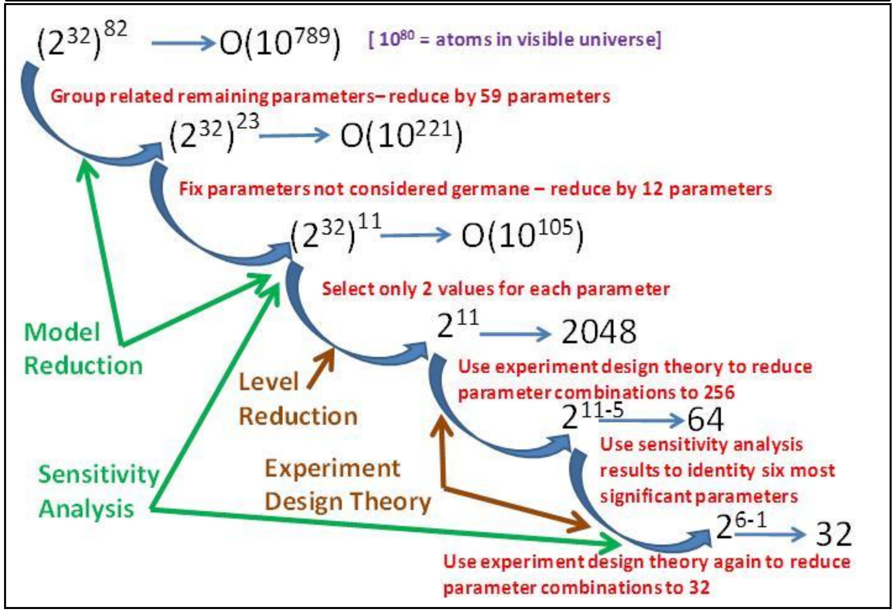

Use correlation analysis and clustering to identify unique behavior dimensions of your model

Compute correlation coefficient (r) for all response pairs

Examine frequency distribution for all $|r|$ to determine threshold for correlation pairs to retain; $|r| > 0.65$, here

Create clusters of mutually correlated pairs; each cluster represents one dimension

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).**

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

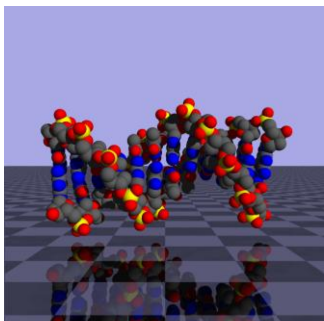

***Selection* based on Anti-Fitness**

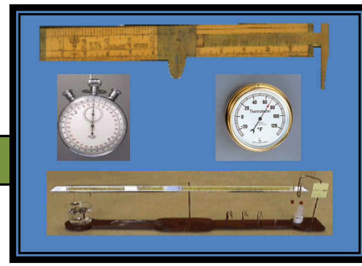

Anti-Fitness Reports

MODEL SIMULATORS

Parallel Execution of Model Simulators

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

Model Parameter Specifications

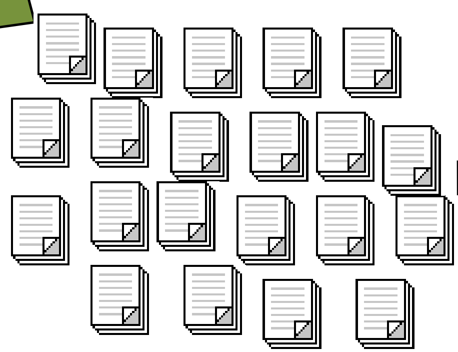

Population of Model Parameterizations

Critical Slowing Down

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

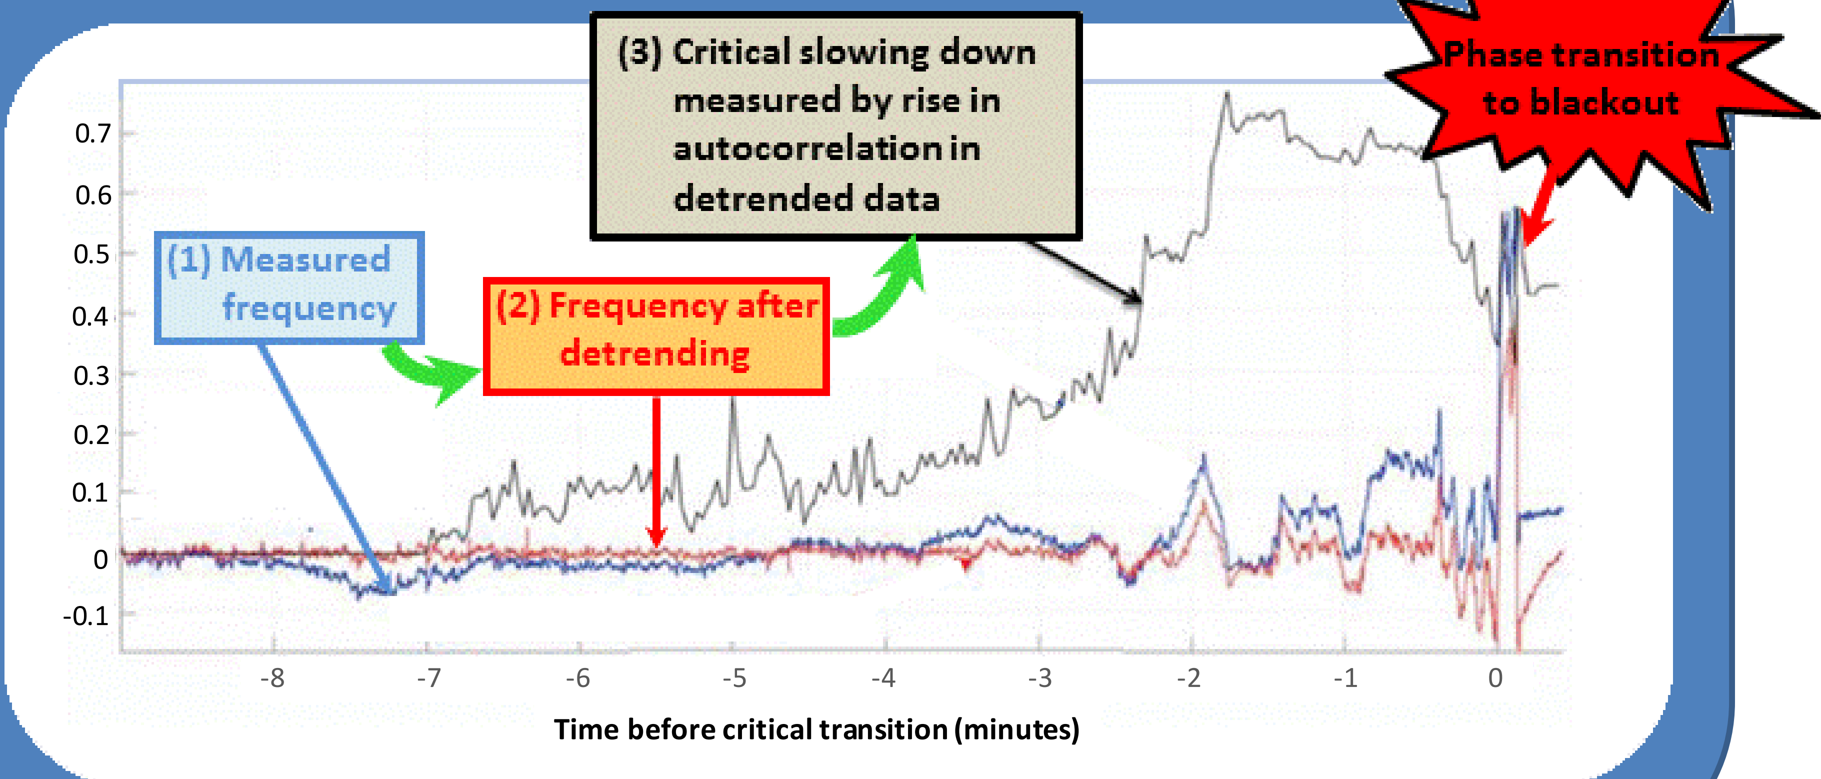

*From P. Hines, E. Cotilla-Sanchez, and S. Blumsack. Topological Models and Critical Slowing Down: Two Approaches to Power System Risk Analysis. Proceedings of the 44th Hawaii Conference on System Sciences. IEEE Computer Society, Washington, DC, USA, pp. 1-10.

Thanks for Listening

Questions?

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