Predicting Global Failure Regimes in Complex Information Systems

C. Dabrowski, J. Filliben and K. Mills – National Institute of Standards and Technology

Our motivation for participating in the workshop on Networks of Networks: Systemic Risk and Infrastructural Interdependencies is threefold: (1) to learn about the goals and methods of other researchers investing risk in large infrastructural systems, (2) to encourage feedback on our goals and methods, and (3) to seek collaborators from industry, academe and government. Over the past five years, we investigated methods to characterize global behavior in large distributed systems and applied those methods to predict effects from deploying alternate distributed control algorithms (a complete record of this research is available elsewhere [1]). The methods we used assess global behaviors under a wide range of conditions, enable significant understanding of overall system dynamics, and yield insightful comparisons of competing control regimes. On the other hand, such methods do not provide information about potential for rare combinations of events to drive system dynamics into global failure regimes, leading to catastrophic collapse. Our ongoing research aims to address this topic using two complementary thrusts: (1) *design-time methods* that enable system architects to identify and evaluate global failure scenarios that could lead to system collapse and (2) *run-time methods* that alert system operators about incipient transition to global failure regimes, and subsequent collapse. Effective design-time methods will enable architects to devise mechanisms that can prevent high-risk scenarios. Since no design-time methods can identify all possible failure scenarios, effective run-time methods will signal operators when system trajectory trends toward collapse, allowing remedial actions to forestall or mitigate catastrophic failure. In this short contribution, we reprise our previous work on methods to characterize global system dynamics and compare alternate control regimes, and then describe our ongoing work toward design-time and run-time methods for predicting global failure regimes.

Most large distributed systems (e.g., the Internet and computational grids and clouds) rely on decentralized algorithms to achieve scalable control of global behavior. As researchers develop new control algorithms, questions naturally arise as to how widespread deployment will influence global system behavior. We developed a four-step method [2] to answer such questions using system models. Our method mainly aims to reduce the state-space that must be searched, while providing as much information as possible about system behavior, and the effects of alternate control algorithms. Step one requires creating a reduced scale model [3] tailored to address a specific question. Step two entails conducting a sensitivity analysis [4-6] to determine the essential behaviors present in the model and to identify the input parameters that significantly influence those behaviors. Step three applies orthogonal fractional factorial design to compare alternate control algorithms [2] under specific combinations of significant model parameters. Step four uses various multidimensional statistical and visualization analysis techniques [2] to examine how essential model behaviors are influenced by alternate control algorithms. We first applied these methods to study [2, 7] proposed congestion-control mechanisms intended to replace standard Internet congestion-control (i.e., the transmission control protocol, or TCP). We were able to demonstrate that the proposed TCP replacements could provide users with improved throughput only under strictly bounded circumstances that were likely to occur rarely in the Internet, and we found one particular proposed TCP replacement could oscillate with high frequency and large amplitude when traversing areas and time periods of heavy congestion, leading to exacerbated congestion. The methods and related findings were quite well-received when briefed to appropriate industry research groups and standards bodies. Subsequently, we applied [8] the same four-step method to compare 18 possible virtual-machine (VM) placement algorithms for infrastructure-as-a-service (IaaS) clouds. All of this work has also been well-received by a panel of experts from the National Research Council (NRC), which regularly reviews the technical quality of NIST research [9]. These same NRC experts expressed interest in methods that could be used to predict the effects of low-probability events on global system behavior. This interest, along with our own finding that IaaS clouds could experience unexpected and widespread resource leakage [10], stimulated us to develop a related research agenda, covering both design-time and run-time methods to detect global failure scenarios, leading to system collapse.

We have identified two possible design-time methods for detecting failure scenarios in large distributed systems. One method [11-13] constructs a graph (i.e., Markov model) from an instrumented version of a system model, and then applies cut-set analysis to identify where cutting essential edges could lead to system collapse. Subsequently, one can apply perturbation analysis to quantify numeric thresholds marking the onset of such collapse and to predict the collapse trajectory. We were able to verify the predictions of the Markov model by observing simulated target systems. This combination of techniques has proven useful to predict selected potential causes for collapse in the performance of computational grids and clouds, as well as predicting that Internet throughput is generally robust, collapsing only under the loss of all communication paths between corresponding parties. On the other hand, this combination of methods has yet proven unable to make predictions in cases resulting from complex indirect feedback paths in a system model, which Markov chains cannot easily capture. While we continue to work on unconventional Markov representations in an effort to capture complex feedback paths, we have also identified a second method, which we call Anti-Optimization + Genetic Algorithms (AO+GA), which should be able to predict failure scenarios caused by complex feedback paths. AO+GA requires defining a metric that denotes successful system performance (e.g., 100% of users are served) and inverting it to denote system failure (e.g., 100% of users are not served). Then, given a system model (the fitness function), a genetic algorithm (GA) can be used to generate parameter combinations (the chromosomes) that, over time, drive a population (of system models) toward a collection of failed systems. The GA can be stopped using a variety of criteria, for example, when 70% of the simulations yield responses where more than 50% of users are not served. Along the evolutionary path to the final population, the GA will identify combinations of input parameters that result in failed outcomes, and statistical analysis and clustering techniques can be used to derive sets of parameter combinations that lead to failed systems, where each set would suggest a cause of collapse, such as supply-demand mismatches, cascading failures and so on. We have recently begun to develop the infrastructure necessary to apply this technique to predict failure scenarios in IaaS clouds.

We have also begun to consider run-time methods that can signal incipient change in behavior toward a failure regime. A recent Nature article [14] reports that "…certain generic symptoms may occur in a wide class of systems (e.g., physical, biological, geological and financial) as they approach a critical point." The generic symptoms relate to a systemic slowing down, which implies systems become increasingly slow in responding to small perturbations. The Nature article and others [15-17] describe a number of mathematical, statistical and computational techniques (e.g., eigenanalysis, autocorrelation and variance analysis, analysis of skewness and flickering) that might be applied to time series measurements to identify an approaching critical point. The articles also discuss selected patterns (e.g., scale-invariant power-law structures) that may appear spatially as systems approach criticality. The generic nature of these suggested techniques promise wide applicability to both natural and engineered systems[1]. In fact, other researchers have begun investigating the applicability of such techniques to electric power grids [18], and we have applied such techniques to communication networks [19]. We have outlined a multiyear research program [20] to investigate the application of similar techniques to predict the onset of phase transitions in large distributed systems, such as the Internet and computational grids and clouds. While the technical merits and innovativeness of our proposal have been recognized, ultimately we need to enlist collaborators, with large operational systems or laboratories, who can provide access to real-time system data in order to fully evaluate the practicality of any promising techniques that we identify through theoretical studies and prototypes.

Understanding and predicting global behavior in large distributed systems is a problem that is both important and difficult. We have identified a collection of helpful methods to study such global behavior, and we have applied those methods to predict how changes in Internet congestion-control procedures would affect users and influence network behavior. We have applied the same methods to predict how various VM-placement algorithms will affect users and influence system behavior in IaaS clouds. Predicting rare combinations of events that can lead to global failure regimes is also a difficult and important problem. We have made some initial investigations into design-time methods to predict failure scenarios that could lead to widespread collapse in large distributed systems. We have ongoing research to extend those methods and to investigate alternate methods. We have defined a program of research to investigate run-time methods to predict changes in global behavior in large distributed systems. We seek collaborators who are also interested in design-time and run-time methods to predict global failure regimes in infrastructures on which modern society increasingly depends.

References

[1] http://www.nist.gov/itl/antd/emergent_behavior.cfm

[2] K. Mills, J. Filliben, D-Y. Cho and E. Schwartz, "Predicting Macroscopic Dynamics in Large Distributed Systems", *Proceedings of ASME 2011 Conference on Pressure Vessels & Piping*, Baltimore, MD, July 17-22, 2011.

[3] K. Mills, E. Schwartz and J. Yuan, "How to Model a TCP/IP Network using only 20 Parameters", *Proceedings of the 2010 Winter Simulation Conference* (WSC 2010), Dec. 5-8, Baltimore, MD.

[4] K. Mills and J. Filliben, "An Efficient Sensitivity Analysis Method for Network Simulation Models", presented at the *2010 Winter Simulation Conference* (WSC 2010), Dec. 5-8, Baltimore, MD

[1] The techniques we identify [14-17] measure changes in behavior over time; thus, such techniques would be applicable to predict catastrophic phase transitions only for system behavioral collapses that unfold sufficiently slowly with respect to the operating time-scale of the particular detection method used. Note that some phase transitions in engineered systems might occur too quickly to provide sufficient time for such predictive detection.

[5] K. Mills, J. Filliben and C. Dabrowski, "An Efficient Sensitivity Analysis Method for Large Cloud Simulations", *Proceedings of the 4th International Cloud Computing Conference*, IEEE, Washington, D.C., July 5-9, 2011.

[6] K. Mills and J. Filliben, "Comparison of Two Dimension-Reduction Methods for Network Simulation Models", *Journal of Research of the National Institute of Standards and Technology*, 116-5, September-October 2011, pp. 771-783.

[7] K. Mills, J. Filliben, D. Cho, E. Schwartz and D. Genin, Study of Proposed Internet Congestion Control Mechanisms, NIST Special Publication 500-282, May 2010, 534 pages.

[8] K. Mills, J. Filliben and C. Dabrowski, "Comparing VM-Placement Algorithms for On-Demand Clouds", *Proceedings of IEEE CloudCom 2011*, Nov. 29-Dec. 1, Athens, Greece, pp. 91-98.

[9] Panel on Information Technology Laboratory Assessments Board, AN ASSESSMENT OF THE NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY INFORMATION TECHNOLOGY LABORATORY FISCAL YEAR 2011, Division of Engineering and Physical Sciences, National Research Council of the National Academies, The National Academy Press.

[10] C. Dabrowski and K. Mills, "VM Leakage and Orphan Control in Open-Source Clouds", *Proceedings of IEEE CloudCom 2011*, Nov. 29-Dec. 1, Athens, Greece, pp. 554-559.

[11] C. Dabrowksi and F. Hunt, "Using Markov Chain Analysis to Study Dynamic Behavior in Large-Scale Grid Systems", *Proceedings of the 7th Australasian Symposium on Grid Computing and e-Research*, Wellington, New Zealand, Jan. 2009.

[12] F. Hunt, K. Morrison and C. Dabrowski, "Spectral Based Methods that Streamline the Search for Failure Scenarios in Large-Scale Distributed Systems", *Proceedings of the IASTED International Conference in Applied Simulation and Modeling*, June 22-24, 2011.

[13] C. Dabrowski and F. Hunt, "Using Markov Chains and Graph Theory Concepts to Analyze Behavior in Complex Distributed Systems", *Proceedings of the 23rd European Modeling and Simulation Symposium*, September 12-14, 2011. Rome, Italy.

[14] M. Scheffer et al., "Early warning signals for critical transitions", Nature, v. 461, no. 3, Sep.3, 2009, pp. 53-59.

[15] J. Lim and B. Epureanu, "Forecasting a class of bifurcations: Theory and experiment", PHYSICAL REVIEW E 83, 016203, 2011.

[16] C. Kuehn, "A mathematical framework for critical transitions: normal forms, variance and applications" Cornell University Library, arXiv eprint, 1101.2908. arXiv:1101.2908v1 [math.DS]

[17] R. Washington-Allen, D. Briske, H. Shugart, and L. Salo, "Introduction to special feature on catastrophic thresholds, perspectives, definitions, and application", Ecology and Society, 15(3): 38, 2008.

[18] B. Carreras, V. Lynch, I. Dobson and D. Newman, "Critical points and transitions in a electric power transmission model for cascading failure blackouts", *Chaos*, 12(4), (2005)985-994.

[19] J. Yuan and K. Mills, "Monitoring the Macroscopic Effect of DDoS Flooding Attacks", *IEEE Transactions on Dependable and Secure Computing*, 2/4, October-December 2005, pp. 324-335.

[20] K. Mills, J. Filliben, B. Rust, C. Dabrowski and D. Genin, "Predicting the Unpredictable in Complex Information Systems" a proposal to the FY2012 NIST competition for Innovation in Measurement Science, http://www.nist.gov/itl/antd/upload/MillsIMS-Proposal-Predicting-the-Unpredictable-RevG-CostsRemoved-ReviewerQ-A.pdf.