Response to Comments on IMS proposal: Early warning of network catastrophes

The review comments suggest that the reviewers focused mainly on the first part (Phase I) of the proposed research: to test and validate phase-transition theories in the context of realistic network models. While this part of the research is necessary, a significant advance in capability requires completing Phases II and III. Of course, the initial phase of the research agenda is essential because existing phase-transition theories consider networks that are far abstracted from reality. Only if such theories prove valid for realistic network models, can they form a basis for measurement processes in real networks, so that network operators can be alerted to take mitigating actions. Injecting effective measurement processes into real networks is the ultimate aim of the proposed project. In other words, *we aim to advance theory into practice, which would bring substantial innovation to the field of network monitoring and management*.

The reviewers appear to give insufficient emphasis to Phases II and III of the proposed research: to develop measurement methods, and supporting software, for evaluation in network test beds, and in real networks. This part of the research, if successful, will provide a significant advance in industry's ability to monitor complex networks and distributed information systems, on which the nation critically depends. So, ultimately, the aim of the project is not simply to generate data in order to validate a theory, but rather to develop measurement methods (Phase II) that can be deployed in real networks to provide monitoring and early warning of network catastrophes. Phase III focuses on establishing working partnerships with network operators, who can evaluate and refine the NIST-developed measurement methods and software within the nation's operational networks.

The reviewers raise two specific questions about proposal content, which we address in turn. First, the reviewers ask what insight can be gained from the results? We will determine conditions under which the phase-transition behaviors demonstrated by theory, and related studies of abstract network models, also appear (or not) in realistic network models, and in laboratory networks. We will determine if the onset of phase transitions can be predicted based on observable changes in only a subset of nodes, and, if so, how such subsets relate to network topology. At the same time, we will characterize relationships between false positives and the selection of measurement intervals and alerting thresholds. We will evaluate which precursor signals appear most effective in predicting onset of phase transitions, and under what constraints. If phase-transition behaviors do not appear in realistic network models, then we will provide explanations, which can serve to refocus ongoing academic research. Such results may also inform ongoing research that aims at predicting phase transitions in other types of networks, such as electrical grids and cyber-physical systems. These findings are the main insights that can be gained from the proposed research.

In a second question, the reviewers ask: can weak spots in a realistic network be identified and bolstered? This is our main research question, which we aim to answer in Phases II and III. The existing theory of network phase transitions suggests that precursor signals can be identified, and sufficiently timely warnings issued, so that network operators can take mitigating actions. If (in Phase I) we find that theory holds in realistic networks, then (in Phase II) we will develop measurement methods aimed at identifying weaknesses in operational networks. Then (in Phase III), working together with network operators, we will determine if such measurement methods yield practical value. The theory is promising, but unproven. The aim of our project is to evaluate the theory under realistic assumptions, and then, if the theory holds, to transform the theory into practical measurement methods and software that can identify weak spots in operational networks, and allow network operators to bolster their networks to remedy weaknesses.

The reviewers question whether the proposed research has sufficient innovation. NIST has a long tradition of providing empirical measurements to evaluate the applicability of theory. For example, Dr. Ernest Ambler, a former NIST director, was one of a team of scientists who conducted empirical physics experiments demonstrating that the conservation of parity does not hold. The NIST experiments verified a theory developed by researchers at Princeton's Institute for Advanced Study, for which they subsequently won the Nobel Prize in physics. The research proposed in our project follows in this NIST tradition by evaluating empirically whether an existing theory of network phase transitions is applicable in real networks. The theory of network phase transitions has not been evaluated in scenarios that adequately represent the complexity of modern communication networks. Our proposed research would be the first attempt to advance the theory into practice. If the theory holds, then our proposed research will also develop and evaluate measurement methods based on the theory. Industry can apply those measurement methods to achieve a fundamental advance in operational monitoring for early warning and mitigation of network failures. The end aim of our proposed research is to produce a measurement method that provides a distinct leap forward for network monitoring and management (as stated in three support letters attached to the proposal).

The reviewers wonder whether NIST is the best organization to perform this work. Aside from the long NIST tradition of applying empirical measurement to evaluate theories, as described above, there are other reasons, both general and specific, why NIST is the best organization to perform the proposed research. In general, NIST has an illustrious history of converting advances from academe into practical methods that industry can use to improve commercial products and processes. For example, NIST, working with other government agencies and companies specializing in computer security, leveraged a cryptographic algorithm from two Belgian academics into an Advanced Encryption Standard that is used worldwide to protect information, even up to the level of top secret. As another example, NIST worked together with the University of New Hampshire Interoperability Laboratory to create a testing program for the Internet Protocol Version 6 standard. The results from that testing program are now used throughout the nation to provide evidence that commercial products conform to the standard.

More specifically, in the case of our proposed research, the ITL team has significant advantages over other organizations. First, the ITL team has a decade of experience in developing computationally tractable, realistic network models (see K. Mills, et al., "How to Model a TCP/IP Network using only 20 Parameters", *Proceedings of the 2010 Winter Simulation Conference*), and applying such models to study Internet congestion (see K. Mills, et al., Study of Proposed Internet Congestion Control Algorithms, NIST Special Publication 500-282, May 2010). In discussing these prior results with industrial and academic stakeholders, academic researchers often observe that they do not have sufficient incentives and resources to conduct studies of the same breadth and depth as those of the NIST research team. Industrial researchers often observe that NIST models and experiments seem much closer to reality than is the case for academic research. Thus, the ITL research team is well positioned to bridge a gap between academic theory and industrial practice. Second, the ITL team has access to an in-house network emulation facility that can be used for initial validation of simulation findings. While the facility is not unique, it reflects the state-of-the-art in network emulation, as represented by Emulab technology from the University of Utah, on which the NIST facility is based. Third, through various other programs, NIST researchers have ongoing collaborations with network operators and equipment vendors. These ongoing collaborations can be leveraged to identify potential commercial partners for Phase III of the research plan, and to provide general access to commercial audiences who could benefit from measurement methods that are developed by the project. Thus, NIST has a unique combination of: (1) knowledge and experience developing realistic network models, (2) access to state-of-the-art network emulation facilities, and (3) broad and deep connections (and well-known reputation) with network operators and equipment manufacturers. This unique combination of characteristics positions NIST as the best organization to conduct the proposed research.

Finally, the reviewers wonder if the technical plan includes any “stretch goals”. The technical plan is designed as a series of three phases, where each phase defines an objective, and the objectives of later phases depend on successful results from preceding phases. Phase I focuses on validating the theory of network phase transitions, within the context of realistic network models, or at least characterizing the limits of the theory’s applicability. Phase II aims to develop and verify deployable measurement software that can monitor operational networks and raise alerts of impending network catastrophes. Phase III represents one, initial, stretch goal for the proposed research: convincing selected network operators that NIST-developed measurement methods and monitoring software can be deployed safely in operational networks and can be effective in alerting about incipient phase transitions. This is a stretch goal because even if we can successfully validate the theory of network phase transitions and can develop effective measurement methods and related software, adoption by commercial partners is required for the project to achieve maximal effect. Another stretch goal is to convince a wide array of system operators of the safety and utility of the proposed measurement methods and software, as well as extending the approaches to measure behaviors beyond network congestion. Achieving ultimate success on these latter stretch goals is likely to require continued research and engagement with commercial partners long after the other project goals are achieved.