

DOE/MICS/ECPI Project Summary

Project Title: A Theory of Stability for Communication Networks
PI: David Starobinski
PI Institution: Boston University

DOE Collaborator: Dr. Nageswara Rao, ORNL
Post-doctoral fellow: Dr. Reuven Cohen
Graduate Student: Niloofar Fazlollahi
Undergraduate students: Andres Guedez, Edy Tan
Project Website: <http://nislalab.bu.edu/nislalab/projects/stability/index.html>

Abstract:

This research consists of two major thrusts. The first one is to develop formal methods and algorithms to ensure the stability of control-plane protocols in UltraScienceNet and future DOE production networks. The second one is to devise high performance bandwidth allocation mechanisms for UltraScienceNet that make most efficient use of available network resources. In addition to performing detailed mathematical analysis, we are developing advanced simulation and visualization tools that allow the evaluation of the proposed mechanisms under various topological settings and applications workload.

Project Description:

The US Department of Energy has recently commissioned the design of an ultra high capacity testbed, called UltraScienceNet, that will facilitate the development of new networks with unprecedented capabilities to support large-scale DOE applications. UltraScienceNet substantially differs from traditional communication networks (e.g., the Internet), because of its unique capability to set up and tear down dedicated channels on demand.

Research Thrust 1: The first goal of this project is to guarantee that the control-plane protocols (e.g., routing protocols, traffic engineering protocols, etc.) of UltraScienceNet and future DOE production networks will be resilient to various manifestations of instability, such as increased delay, loss of connectivity, and routing instabilities that may result from switch/router overload, nodes and link failures, maintenance operations, or malicious attacks. As part of our research, we have found that most of the above instability issues can be avoided if the underlying network topology does not contain cycles. Thus, we have been developing new graph-theoretic methodologies, based on the concept of “turn-prohibition,” that prevent *a-priori* the creation of cycles in the network. In particular, we have shown that our methodologies are applicable to layer-2 networks, e.g., Gigabit Ethernet backbones, and can significantly improve their performance and scalability [1]. We are currently investigating the application of our techniques to enhance the robustness of layer-3 routing protocols, with a particular emphasis on policy-based routing protocols such as BGP.

Research Thrust 2: Our second goal is to optimize the process of bandwidth allocation in UltraScienceNet. Our main performance metrics are (i) *delay*, where the goal is to minimize the time between a file transfer request and the completion of the file transfer; and (ii) *throughput*, where the goal is to maximize the number of requests per unit of time that the network can accommodate.

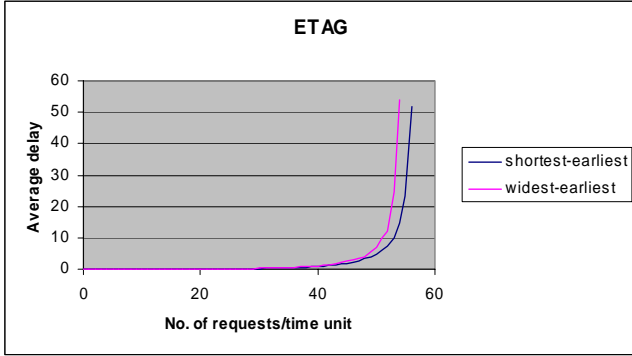


Figure 1: Comparison of bandwidth allocation algorithms for Ultra Net

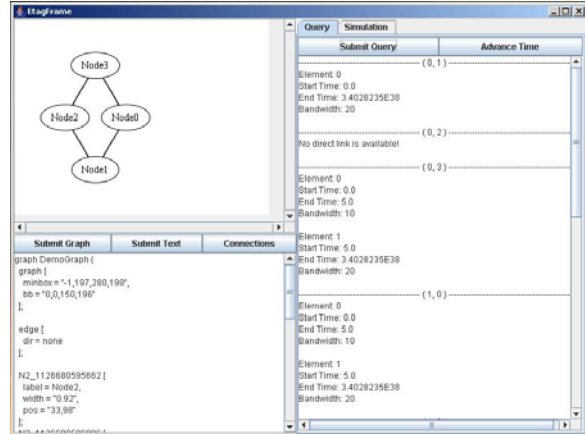


Figure 2: Bandwidth allocation GUI.

Currently, UltraScienceNet uses an algorithm, called ALL-SLOTS, that lists all available time-slots for a path of bandwidth B and duration of at least t between a certain source s and destination d [2]. As part of this research, we are proposing various modifications and improvements to this algorithm that show promise for significant performance gain.

For instance, the current ALL-SLOTS algorithm does not have a provision to select a path with certain desired properties (e.g., the shortest path or the widest path) when multiple paths are available between a source and destination at a certain time slot. To address this issue, we have developed a different algorithm that can compute such an optimized path. This algorithm has polynomial complexity in the number of nodes in the graph and in the number of pending connections that have already been scheduled.

We are also exploring a new concept called “path switching,” whereby a connection between a source s and destination d can switch over time between different paths. We have devised an algorithm called ALL-SLOTS/PATH-SWITCHING that list all available time-slots where it is possible to establish a connection (possibly switching between different paths) with bandwidth B and duration t between a source s and destination d .

To evaluate the performance of our algorithms, we have developed a discrete-event simulator written in C. The simulator allows configuring various parameters such as the network topology, the arrival rate of requests, the average connection length, and the source-destination pair for each request. The simulator returns throughput and delay performance of the algorithm under study. Fig. 1 shows an example where the simulator compares between the performance of a “shortest-path” algorithm and a “widest-path” algorithm. Our simulator has been useful in showing that the “shortest-path” algorithm usually performs better than other heuristics for path selection. In parallel, we have developed a simple GUI interface, written in Java, that allows better visualization and understanding of the operations of various bandwidth allocation algorithms (see Fig. 2).

References

- [1] F. De Pellegrini, D. Starobinski, M. Karpovsky and L. Levitin, “Scalable, Distributed Cycle-Breaking Algorithms for Gigabit Ethernet Backbones,” submitted for journal publication (preliminary version appeared in *Proc. IEEE INFOCOM 2004*).
- [2] N.S.V. Rao, W.R. Wing, S.M. Carter, and Q. Wu, “UltraScience Net: Network Testbed for Large-Scale Science Applications,” *IEEE Communications Magazine*, 2005.