Contribution ID: **6** Type: **not specified**

Evolution of Networking Research into a Science

Tuesday, 11 September 2012 11:52 (12 minutes)

Evolution of Networking Research into a Science

Introduction
Building COmputational Models for BIg NEtworks will consist of developing a complex set of models, both in the vertical and the horizontal space as defined in the call for abstracts. Among the large number of variables that must be considered, a major component is network traffic – how it has evolved over the years, what drives it, and how can we predict traffic patterns, if at all, and thus build future models for network traffic. Understanding the characteristics of network traffic, modeling traffic and studying the effects of different models of traffic workloads on network and application performance has been a major thrust of our research. We have mainly studied traffic models as a means to study better experimental methods for networking research. Our broader and longer-term vision is to help evolve networking research from its current adhoc status into a science.

Background
Experimental methods for networking research have evolved almost ad-hoc over the last decade. That there is no consensus among the research community on best practices in empirical research is only part of the problem. The fundamental problem that calls to question the results of any empirical research is this – as a nascent field, empirical networking has yet to develop into a science. In the physical sciences, experimentation and evaluation of new ideas has been honed so that experiments conducted by one research group can be repeated and tested by another. Due to lack of such coherent standards in experimental methods, this is simply not possible in networking research today. Even within the wider field of computing for example, there are well-known benchmarks for testing new processors against existing ones. In most fields, benchmark standards test new inventions. For example, if Intel develops a new processor, several benchmarks test the new processor to demonstrate that it performs better than an existing one. This is possible because there are agreed-upon scientific processes for experimentation and evaluation in that field.

`For over a decade now, networking researchers have built small and large research testbeds using hardware`

Despite the significant advances made in development and deployment of very large testbeds and the software framework for building different topologies and reserving these resources to run large-scale experiments, there is still little understanding of what constitutes best experimental practices. There are no agreed-upon scientific methods for running experiments, no benchmark workload models, no network emulation standards. The networking research community lacks a coherent, shared view of best practices for experimental methods for networking research. Even today, there are no agreed-upon research methods or standard practices for maintaining traffic datasets, generating traffic, emulating network characteristics, or designing and running experiments. Hence while many researchers propose new and improved protocols for improving our cyber-infrastructure, real progress and deployment of the best protocols is slow.

Research Questions
However, computer networking, as a nascent field with explosive growth, woefully lacks such benchmarks and standards. Establishing such benchmarks remains a challenging research endeavor, and it forms the central motivation for our research. We need to develop models for several components of experimentation which include workload modeling, network path emulations, network topologies, measurement methodologies, and determining which performance metrics best describe the outcome of experiments. Before we develop any models, however, there has to be some agreement in the scientific community of certain standards.

Several research projects have chipped away at this problem of lack of scientific methods for networking research for some years now, but the field is fragmented at best. It is time to take this research to the next level by integrating the research of experimental methods in networking in a fundamental way to create a holistic view.

Some of the research questions that we're grappling with lately have been the following:
– Can we develop a set of necessary and sufficient conditions for running a successful experiment for empirical networking research? For example, topology of testbed, duration of experiment, number of runs for each experiment, source of input traffic, mix of traffic, workload model, network emulation, and measurement methodology. How can we build models if we do not have comprehensive understanding of how each component works and how they all interact to form this complex system?
– Could we classify experiments into sets, based on the input, the output, or the goal?

– Can we develop scientific methodology for calibration of a network for experimentation? We already do this in the lab, but can we generalize calibration for any network?
– Run the same experiments using different sources of traffic to determine if there are some invariants in traffic generation and experimentation and what are the variants?
– How can we reproduce and validate experiments conducted by a different research team?
– How can we develop "experimental methods in networking research" into a science?
Given our research interests and the work we have done in this area, we completely appreciate the dire need to get out of the reductionist stance that we, as a community, have taken in the so-far pseudo-scientific investigation of networking as a field. It brings to mind the blind men and the elephant, and an evolving elephant at that. We have decided that the problem is so complex that we shall simply take one slice of it and analyze that slice independently rather than deal with the very difficult problem of creating holistic models to solve problems in this space. Yet, we do remain skeptical about building such models because networks seem to be evolving so rapidly and they are such complex systems. However, it is all man-made and we are driving its evolution. So surely, if we employ large enough resources to this problem and go at it systematically and globally, we should be able to create reliable and evolving models for networks?
My interests and expertise thus form a slice of the holistic modeling that this workshop seeks to attempt to motivate. My own motivation in attending this workshop would be to lend my expertise and collaborate with this group so we might be able to work toward constructing such large-scale models for networks that would encompass all components emulating the vertical and horizontal spaces while being able to evolve over time.

References

– J. Aikat, S. Hasan, K. Jeffay, and F. D. Smith, Towards Traffic Benchmarks for Empirical Networking Research: The Role of Connection Structure in Traffic Workload Modeling, IEEE MASCOTS (Modeling, Analysis, and Simulation of Computer and Telecommunication Systems), Washington DC, August 2012.
– J. Aikat, S. Hasan, K. Jeffay, and F. D. Smith, Discrete-Approximation of Measured Round Trip Time Distributions: A Model for Network Emulation, GENI Research and Education Experiment Workshop 2012 (GREE12), Los Angeles, CA, March 2012.
– J. Aikat, K. Jeffay, and F. D. Smith, Experimental Methods for Networking Research and Practice, SIGCOMM 2011 Education Workshop, Toronto, Canada, August 2011.

Primary author: Dr AIKAT, Jay (UNC-Chapel Hill)

Co-author: Dr JEFFAY, Kevin (UNC-Chapel Hill)

Presenter: Dr AIKAT, Jay (UNC-Chapel Hill)