Experimental Methods for Networking Research and Practice

Jay Aikat
UNC-Chapel Hill
aikat@cs.unc.edu

Kevin Jeffay
UNC-Chapel Hill
jeffay@cs.unc.edu

F. Donelson Smith
UNC-Chapel Hill
smithfd@cs.unc.edu

ABSTRACT

In this paper, we propose an upper undergraduate / graduate level course titled: "Experimental Methods for Networking Research and Practice". This course is designed as a learning experience with two primary goals – conceptual understanding of the various research methods used in empirical networking research, and hands-on experience by performing two hands-on projects, one in a controlled lab setting and the other in the GENI environment.

Networking researchers have long used experimental networks and distributed systems for designing and evaluating new networking technologies. However, the networking research community lacks a coherent, shared view of best practices for experimental methods for networking research. And hence, the teaching of experimental research methods lags even further behind.

So, even as we make significant methodological advances in empirical networking research, and build large-scale collaborative research testbeds like the GENI infrastructure, these resources have yet to make their way into our classrooms and textbooks.

In this paper, we present a design for teaching experimental methods to networking students, using a simple laboratory testbed as well as the clusters in the GENI infrastructure.

We welcome feedback from this community on this ongoing effort.

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]

General Terms

Measurement, Performance, Experimentation.

Keywords

Research Methods, Network experimentation, Laboratory testbeds, Traffic generation, Experimental evaluations, Performance evaluations.

1. INTRODUCTION AND MOTIVATION

Networking researchers have long used experimental networks and distributed systems for designing and evaluating new networking technologies. Indeed, experimentation, either via software simulation using simulators such as the Network Simulator (NS) [NS], or via hardware emulation using laboratory testbeds, has been the primary means for evaluating existing and newly proposed protocols and algorithms for improving the Internet. Nonetheless, experimental networking remains a challenging research endeavor.

Floyd and Paxson cataloged some of the fundamental issues in experimental networking in their seminal 2001 paper "Difficulties in Simulating the Internet" [FP01]. Since that time, numerous research projects funded by the NSF, including the Global Environment for Network Innovations (GENI) projects and the traffic generation project at UNC – Chapel Hill (UNC-CH) have sought to address some of the difficult methodological issues raised in that paper. More recently, Floyd and Kohler document in their 2008 Internet Draft [FK08], that there has been some effort to formulate evaluation scenarios specific to congestion control experiments.

In [AMF+08], the authors create a case for a common evaluation standard for TCP evaluations. This paper does not present any results of experimentation, but acts as a powerful catalyst for discussions on this topic. There is also a related and ongoing effort by the "Transport Modeling Research Group" [TMRG] to come up with a consensus for a baseline standard for protocol evaluation.

However, despite all these efforts, the networking research community lacks a coherent, shared view of best practices for experimental methods for networking research. And hence, the teaching of experimental research methods lags even further behind. So, even as we make significant methodological advances in empirical networking research, and build large-scale collaborative research testbeds, these resources have yet to make their way into our classrooms and textbooks.

In this paper, we present a design for teaching experimental methods to networking students, using a simple laboratory testbed as well as the clusters in the GENI infrastructure. We propose a course titled: "Experimental Methods for Networking Research and Practice". This course would employ the large, publicly available GENI testbeds, and use these research infrastructures for educational purposes as outlined in this paper.

This course will serve two classes of students: (i) those who will work in industry developing and/or testing new devices or protocols for network deployment, and (ii) those will employ experimental methods for networking research in an academic or industrial laboratory setting. The course will focus on key issues in empirical methods in networking research, including network measurements, data analysis and traffic generation for network experiments.

2. Course Design

This course is designed as a learning experience with two primary goals – conceptual understanding of the various research methods used in empirical networking research, and grappling with issues by performing two hands-on laboratory projects.

2.1 Target audience

This course is proposed as a second networking course. It will be assumed that the students, who could be upper-level undergraduates or first/second year graduate students, would have had an introduction to networking course already. Hence, the students would be quite familiar with the TCP/IP protocol suite, local area networks, and the Internet routing protocols.

2.2 Course Design

This course will be designed as follows:

- Part I: One-third of the course will consist of an overview of key topics such as protocol design, network architectures, congestion control mechanisms (end-to-end and router-based), performance evaluations, and network science.
- Part II: One-third of the course will delve into methodological issues pertaining to measurement, modeling, traffic generation, simulation and emulation, statistical methods, and basics of control theory.
- Part III: One-third of the course will be project based. This will consist of two projects, each running for a half semester. The projects are designed to create a two-step learning experience that will include concrete, hands-on exposure to the following important steps of a network experiment:
    - define the problem
    - design experiments
    - design the experimental testbed
    - identify the methods of traffic generation
    - detail the process of running experiments
    - collect and analyze experimental data
    - write a brief report

2.3 Course Projects

The course projects will follow a two-step process. There will be two course projects constituting the core of this course. The first project will be conducted in the laboratory and the second will be run on a GENI testbed.

2.3.1 Laboratory Project

The first project will be pre-selected and all students will be assigned the same project.

2.3.1.1 Experimental Design and Methodology

All students will all use a laboratory testbed provided to them. This testbed will have been setup and tested, ready to use. This lab project will allow for controlled experiments, giving the students concrete experience in running specific experiments assigned to them. They can thus apply the concepts discussed in class.

The testbed we propose will not involve large commitments of laboratory space and equipment. At UNC – Chapel Hill, we have successfully taught hands-on networking courses by building small laboratory testbeds using commodity PCs and software routers as shown in Figure 1. The traffic generators can be old PCs with 1.5 GHz processors and 500 MB of memory. The software routers and monitors are server class machines. Although we show a DAG monitor, we have also successfully used monitors running *tcpdump* to collect packets on the link.

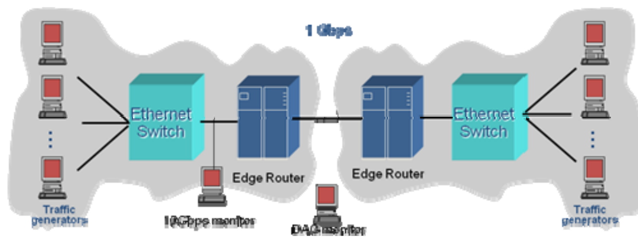

Figure 1: Experimental Network Setup

By working on a real testbed in a controlled environment, the students will gain an understanding of network design and architecture, and how this relates to the design of experiments on that network. They can experience asking and answering questions such as: how do you design experiments given a specific problem, what are the different variables in each experiment, how do these variables relate to each other, which variables can (or should) be controlled, what measurements should be made and at what points in the lab testbed, which metrics of performance should be studied, and how should the results be presented, what are the trade-offs being made, and how do you justify your choices? These are just a sample of the many questions that can be explored in the lab.

2.3.1.2 *Traffic Generation*

In the last decade, significant advances in tools and methods to generate synthetic traffic in simulators and testbeds have been made. These include the Harpoon [SB04], Tmix [WAH+06], and Swing [VV09] suites of tools. These traffic generation platforms enable researchers to perform experiments with traffic that mimics the traffic found on production networks to varying degrees. While such tools exist, we note that these are not widely used even in research, let alone in teaching.

Many networking courses still use only *iperf* [iperf] or similar tools to generate traffic on laboratory testbeds. While *iperf* is a great first step, we need to use more sophisticated tools like Tmix. As shown in Figure 2, traffic generation involves a series of steps from measurement of real traffic on a production link to generating synthetic traffic in the lab that bears all the necessary statistical properties of the original traffic. How and why this is done are questions that can be discussed and evaluated if a realistic traffic generation system is used in these networking projects.

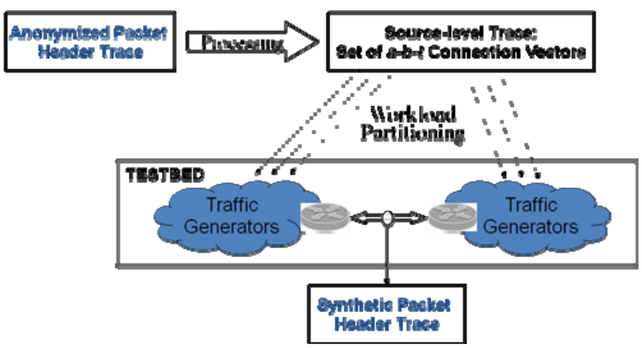

Figure 2: Traffic Generation

We must prepare the students to be asking the right questions about traffic generation in network experiments to prepare them to succeed as researchers or network practitioners. Whether they evaluate a new transport protocol, test a router algorithm, or inject malicious traffic into a network for security evaluation, they will need to generate realistic network traffic in their experiments to obtain reliable results. Hence they must be prepared to answer questions like:

- How (where from) do you obtain realistic Internet traffic data for your experiments?
- How do you model the application workloads that constitute this network traffic?
- What parameters should you include in your workload model?
- How will you generate these workloads in your experimental testbed?
- What network path characteristics must you measure from the original traffic?
- How will you emulate these characteristics faithfully in your testbed?
- How can you generate realistic traffic, and run hundreds of automated experiments using slightly different variations of the same traffic input?

Using the two projects as case studies of experiments for network performance evaluation, we can walk them through the entire process from measurement of production network traffic to modeling workloads and end-to-end paths to generating synthetic traffic for network experiments in the lab to analyzing experimental results using sound statistical methods.

2.3.1.3 *Calibration*

Calibrating a network is a major step in understanding the process of network experimentation in a real testbed. Although calibration itself does not lead to results, it is an important educational step. Hence the students will perform a set of calibration experiments in the controlled laboratory environment. The main motivation for network calibration is to ensure that the network, or any of its individual components, does not present any resource constraints (unless otherwise designed to do so, as in a bandwidth constrained link) when running experiments. The way we verify this is through calibration. Calibration involves first identifying the set of all inputs to the experiment, deciding what the outputs will be, and figuring out the correlations, if any, between these inputs and outputs. The goal of calibration then is to ensure that these correlations are not influenced by an unintended lack of resources in the network.

2.3.1.4 *Process of Experimentation and Analysis of Results*

The actual process of experimentation in the laboratory testbed will have been fairly well defined. Yet, the students will use this step to look at how this ties in to the results of the experiments. Hence, questions that can be explored include: should you run multiple experiments, do you see variance in the results, what statistical methods must you use for data analyses, how do you present the results, why so units matter so much in the presentation, and how does the granularity of plots make such a huge difference in what inferences we make.

2.3.2 *GENI-based Project*

The second project will use one of the GENI testbeds or a specific GENI cluster. While the laboratory testbed allowed for controlled experiments, not all GENI testbeds can be controlled. Hence this will give the students a different experience, and enhance their understanding of the process of network experimentation.

This project will be slightly more open-ended. However, to make it feasible for the students to complete the projects, they will be given a choice of six to ten pre-selected projects that are inspired by leading or classic publications. If the student designed their own project, the scope of that project would be scaled to compare with one of the preselected projects.

Since the projects on the GENI testbeds are not necessarily controlled and, are by design, remote, they will present a significantly difference learning experience for the students. For example, the GENI testbed will not allow for strict calibration as discussed in the controlled environment in the lab. Hence, the student must follow a different process of understanding the network infrastructure and the environment in which they will run their experiments.

Similarly, the methods for measurement and traffic generation will vary, and thus present different challenges to the students in how they run their experiments, collect data, and analyze the results.

The Global Environment for Network Innovations (GENI) is a novel suite of infrastructure being designed to support experimental research in network science and engineering [GENI-Sol3]. The GENI testbeds and resources are designed such that they enable experimenters to use an interconnected set of reserved resources on platforms in diverse locations. The students will be able to remotely reserve, configure, program, and operate distributed systems established across parts of the GENI suite.

The GENI project is ideal for teaching network experimentation because it is fundamentally designed to support such ventures. It supports a wide range of experimental protocols, measurement techniques, running over a range of equipment. While the GENI infrastructure is still in development and testing mode, it is currently beginning its Spiral 3 phase, which encourages use of its resources by the larger community. This is an ideal time to influence some of the development of this infrastructure so that it is made adaptable to teaching as well.

3. Course Development and Maintenance

For successful development and offering of such a course, we, as educators, will need to build and maintain a repository that will consist of traffic datasets, tools for traffic generation, manuals for GENI cluster usage, and tested experimental modules that the community can tap into. With this paper, we hope to start a dialogue within the networking educators community on teaching such a course, and developing and maintaining the resources for its success.

The results of this effort will be documented as part of this community forum and integrated into an experimental networking curriculum. The approach to the development of these materials will be comprehensive and inclusive. It is envisioned that materials ranging from course outlines to software tools will be developed and distributed.

4. Challenges ahead

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. Hence, this only complicates developing and offering courses on experimental methods.

This paper aims to start a discussion and reach some consensus, resulting in some concrete steps on this topic.

5. Conclusion

In this paper, we propose an upper undergraduate / graduate level course titled: "Experimental Methods for Networking Research and Practice". This course is designed as a learning experience with two primary goals – conceptual understanding of the various research methods used in empirical networking research, and hands-on experience by performing two hands-on projects, one in a controlled lab setting and the other in the GENI environment.

The students will be made familiar with the major design trade-offs and decisions used for Internet protocols and systems, and be able to critically assess empirical methods and results reported in the networking research literature. They will be able to successfully perform networking research projects.

We welcome feedback from this community on this effort.

6. REFERENCES

[AMF+08] L. Andrew, C. Marcondes, S. Floyd, L. Dunn, R. Guillier, W. Gang, L. Eggert, S. Ha, and I. Rhee, *Towards a common TCP evaluation suite*, Proceedings of PFLDnet, March 2008.

[NS] The Network Simulator. http://nsnam.isi.edu/nsnam/index.php/Main_Page

[iperf] Iperf – bandwidth measurement tool. *http://dast.nlanr.net/Projects/Iperf/*

[FK08] S. Floyd and E. Kohler. *Tools for the Evaluation of Simulation and Testbed Scenarios*. Internet Draft <http://tools.ietf.org/html/draft-irtf-tmrg-tools-05>

[GENI] The Global Environment for Network Innovations. <http://www.geni.net/>

[FP01] S. Floyd and V. Paxson, *Difficulties in simulating the internet*, IEEE/ACM Transactions on Networking, 9(4):392–403, August 2001.

[SB04] J. Summers and P. Barford, *Self-configuring network traffic generation*, Proceedings of Internet Measurement Conference, 2004.

[Tmix] http://netlab.cs.unc.edu/Tmix

[WAH+06] M. C. Weigle, P. Adurthi, F. Hernandez-Campos, K. Jeffay, and F. D. Smith, *A tool for generating realistic TCP application workloads in ns-2*, ACM Computer Communication Review, 36(3):67–76, July 2006.

[UNCnet] University of North Carolina at Chapel Hill's Networking Laboratory. *http://www.cs.unc.edu/Research/dirt/*

[GENI-Sol3] GENI Solicitation 3. <http://www.geni.net/>

[VV09] K. Vishwanath and A. Vahdat, *Swing: Realistic and responsive network traffic generation*, IEEE/ACM Transactions on Networking, August 2009.