

CS 485/ECE 440/CS 585 Lab 1, Part 1

Due by 11:59pm on Friday, 23 September, as an e-mail to the instructor (jedcrandall@gmail.com). Please send only PDF files.

50 points (out of 200 total for Lab 1)

The purpose of Lab 1 is to understand resource allocation problems in computer networking, especially in terms of routing, flow control, and congestion control.

We'll be exploring problems with a particular aspect of the Tor network, specifically the token bucket filter (filters) that is (are) used to limit Tor's bandwidth usage as an onion router. You will be forming a hypothesis, setting up and performing experiments, and then trying to draw some conclusions about why Tor's token bucket filter mechanism(s) may be causing resource allocation issues. The problem statement we'll work with is this:

There have been unspecified issues reported of resource allocation problems with Tor's token bucket filter mechanism(s) for limiting the bandwidth utilization of an onion router. There are unconfirmed reports of receiving onion routers "ignoring" senders who are trying to send them cells and excessive TCP retransmissions.

For Part 1 of Lab 1 your group will need to form a hypothesis about why these resource allocation problems may be occurring and then design an experimental methodology to test your hypothesis. You won't run any experiments for Part 1 of this lab, Part 1 is a proposal in which you'll propose to me a hypothesis and a methodology for testing your hypothesis. I strongly recommend that you keep in constant touch with me between now and the due date to get lots of feedback on your proposal.

The proposal should be 2 pages in 12-point font with normal margins, submitted to me as a PDF. Please make sure your spelling and grammar are adequate for me to be able to easily read your proposal, spelling and grammar can affect your grade.

Your proposal should clearly state:

- Your hypothesis
- Your proposed experimental methodology
- The resources that you'll need to carry out your experiments, and the setup tasks that will be necessary

About the hypothesis, your hypothesis should be *testable*. It should be a technically detailed sentence

that will either be confirmed or refuted by your experiments. Examples of hypotheses include: “The token bucket filter mechanism generates tokens too infrequently---by generating half as many tokens twice as often performance will improve.”; “Tokens are not allocated fairly per circuit, high-bandwidth circuits exhaust all of the tokens and starve low-bandwidth circuits.”; and, “Sending is bursty, and this bursty usage of the tokens causes unnecessary TCP retransmissions for the receiving of cells.” In general, your hypothesis should lead to a “thumbs up” or “thumbs down” based on your experiments. If you're feeling ambitious and want to try a screening experimental design instead, come talk to me.

About the proposed experimental methodology, try to keep the methodology as simple as possible, but your hypothesis needs to be meaningful and your methodology needs to adequately test that hypothesis, so the methodology shouldn't be *too* simple. As you form your hypotheses, you should be in constant contact with me, and then as you work on your methodology you'll want lots of feedback from me if you want to get a good grade on the proposal for Part 1 and on the lab as a whole. Some things you can Google or ask me about to help you form your experimental methodology: Factors, levels, quantitative vs. qualitative levels, full-factorial design, treatments or runs, randomization, unknown influence, bias or distortion, blocking, interaction, fractional factorial design (for screening experimental designs).

About the resources, here's what we have: 35 lab machines that all have VirtualBox and can run around 5-10 virtual machines at one time, two server machines available for your use (fordyce and shasta) that can run around 50 virtual machines at one time, 20 laptops that each have both wireless and wired interfaces, a few x86 machines, and half a dozen miscellaneous machines (MIPS, Sparc, Alpha) that can run OpenBSD. Keep in touch with me about the resources so I can gauge what the different groups need, what they can share, *etc.* Be specific about the setup for your proposed experimental methodology, *i.e.*, how many onion routers, how many exit nodes, how many clients, *etc.*

To get started on understanding Tor and the token bucket filter mechanism, re-read the papers I assigned the first week (especially the DefenstraTor paper) and then Google “dhungel waiting for anonymity”, “why tor is slow and what we're doing about it”, and “token site:torproject.org”.