

## Final Project

I have a control problem.

I normally assign carefully designed programming projects that I have explored thoroughly myself. This approach lets me reassure myself that I am not asking students to do unreasonable things (or that if I am I can at least help you out). Rumor has it, however, that some of you enjoy the thrill of undertaking projects of your own design based on your particular interests within the topic of the course. As a result, I am trying my hardest to let go and end this semester with an open-ended final project.

Sadly, I can't quite bring myself to do it! I think, however, that I have found an appropriate compromise.

For your final project, I am going to offer you three choices or, to be more accurate, two categories of projects with a fallback for those of you who are chicken like me.

Regardless of which option you end up choosing, you are encouraged strongly to work in groups on this project (i.e., you will need my permission to do otherwise). The working partnerships do not have to be the same ones you have used for the earlier homeworks or programming projects. They also don't have to be pairs. There is already one group of three working on homeworks. It might be ideal if we could reorganize the 9 of you into 3 groups of 3. I don't want to dictate how you form yourselves into working groups, but please share any thoughts you might have in case some intervention on my part can help the process of forming good groups.

During next week (11/2-11/6), you will be busy. Somewhere during the week you need to take the midterm. You also probably need to be working on the Mininet router assignment. There will, however, be no problems to work in preparation for our weekly meeting. Instead, I ask you to work on choosing the final project you want to work on and confirming who you will work with. We will meet next week so you can tell me your plans either at our usual times or at times when you can meet together with your partners if we can arrange a schedule. I have two goals for these meetings. You will need to make a convincing case that you know enough about the project you plan to complete that I can believe you can, in fact, complete it. We will need to agree on a timeline for your work. In all cases the timeline will include completion in time to demonstrate the work in our last lab slot and then present it in our last weekly meetings.

Now, for your options.

The fallback is to complete the project I assigned the last time I taught this course. A copy of the (unmodified) handout for this project will be linked to this semester's course web site. Since it is unmodified, the due dates in the handout are inappropriate. The one "open-ended" aspect of this option will be that any group who chooses to do this assignment will have to work with me to establish a schedule they will follow

This is the safe option because I know it is doable in the time allowed.

If you look at the project I used in the last offering of this course, you will discover that the underlying idea was to use a simulated environment to reproduce the results of an earlier published experiment. The closed-ended part was that I picked the result they had to reproduce. As proof that this was a good underlying idea, I have since discovered that, Nick McKeown, the Stanford faculty member behind Mininet, had the same idea but approached in a bit more broadly. In a graduate networking course he had students find a networking research paper they found interesting, try to reproduce its experiments using Mininet and report on the results. I want to give you the opportunity to do the same thing. However, since you are not graduate students and I am not nearly as knowledgeable about Mininet as McKeown, I want to make your undertaking a bit less risky. With this in mind, your second option is to take one of the reports by the Stanford students describing the attempt to reproduce some paper's results (they are all online at <https://reproducingnetworkresearch.wordpress.com/about/>) and

reproduce the reproduction. If the students did a good job describing their efforts, this should be easy! (Well, I actually suspect there may be some hurdles to overcome, but easier than option 3.) To assure me that these hurdles will be manageable, by our meeting next week, you should make sure you can describe the work in the original paper you will be working on and the basics of how the Stanford students used Mininet in their reproduction. Did they make kernel modifications? Is source code included in their report? etc.

Finally, Yes, there is Option 3. I am going to actually give you an open-ended option. If you prefer to pretend you already are a Stanford graduate student you can propose to reproduce the result from some networks research paper that the Stanford students haven't gotten to yet. In this case, you need to tell me the paper at least a day before we meet so I can read it and come to the meeting with a concrete plan that will reassure me you have a very good chance of completing the work in the limited time available.