

On the Thesis-Proposal Style General Exam: Experiences and Recommendations for Graduate Students

Neil Spring

July 6, 2003

Abstract

There is surprisingly little communal wisdom among graduate students on the recently added thesis-proposal style of generals exam. This is a short list of do's, don'ts, and don't worry abouts from my experience with this style of exam.

1 Introduction

Generals is personal. My experience with generals is unique, just like everyone else's. I was at the end of my fifth year. Yes, I'm a systems student, at least I hadn't already started as faculty.¹ I had two-thirds of the work of a thesis finished. I already knew the area fairly well, but using the thesis proposal style of generals exam to complete my related work section seemed attractive. However, what works for me may not work for you, so don't assume what you read here will apply to your generals experience. This caveat should be attached to all generals advice. So remember, if you can write off any of the advice here to "he's a systems student," "he's a kth year," or "his advisor is Australian" then you may ignore it.

Generals is fun. Plan your generals six months or a year in advance. Think "generals *charge*" in the context of a sporting event or epic battle, not a crime drama or epic lawsuit. Think about summer, when it's nice out, reading some papers, sunbathing as you consider their meaning. Think about the time of year when you want least to be indoors tied to your machine.

Read the other materials collected for you from the grad seminar. Talk with your post-generals office-mates and those in your research group about what's required. They will tell you how to succeed, which will likely be some combination of:

Synthesis or *framework* or simply "a big table." This is not simply about learning the tabular environment in L^AT_EX, but about classifying the methods and technology that you can use to solve problems. It's about finding common fundamental strengths and weaknesses. Comparing approaches along axes seems to work too.

¹The name "Stefan Savage" is often invoked as an example of the generals process gone wrong.

Personal insight. You undoubtedly have a different way of viewing the world than anyone else, perhaps more theoretical or practical or empirical or operational. Maybe you think more like a user or more like a software engineer. Maybe you spent a year in industry or a year in the peace corps. Maybe your undergraduate minor shapes your world view. However this leads, it's what you bring to the process that makes it interesting for everyone else.

Focus on techniques. Focus on the methods and how they can be applied to solve a problem. You can make an exception if conflicting or changing results motivate further analysis. Often the inputs (workload, applications, processor speeds, network speeds) will change, and so the results (performance, comparisons) and conclusions will change with them.

Focus on future work. The ability to find promising future work is the point. Consider a one-paragraph proposal for research for each of several topics. Each paragraph would consist of one sentence for each of description, motivation, methodology, and perhaps expected results and implications. Promising future work is often an empty cell in the big table, generally *not* a tweak to existing approaches, a head-to-head comparison between them, or a hybrid approach that combines the best of both worlds. Promising future work is more likely a composition of the techniques used in the charge papers to do something new. Finally, future work should fit within the framework – if it doesn't, it suggests that the framework was poorly constructed to understand the area.

The rest of this note is structured as follows. First, I describe my strategy, or how I chose to spend my time. Second, I describe what went badly, including some things I didn't prepare for. Then, what went well and hints for making your generals go as smoothly. The last few "sections" cover dealing with your advisor before you start, what advice about "traditional" generals you might want to ignore, and a few technical tools I found valuable.

2 My strategy

The exam part of generals is getting four people in a room to agree that you'll be able to defend successfully sometime in the near future. For me, this took the pressure off, and it became more about not screwing up than proving myself from scratch.

I looked at the General Examination requirements and saw four pages of stuff, then three lines on what a thesis-proposal generals meant. Those three lines said, "it's not a contract." I interpreted this to mean *do generals first, a thesis proposal second*. Or more cynically, you can screw up the thesis proposal because you're not committing yourself to it, but you can't screw up the generals.

While most advice said "do nothing else" during the exam period, I disobeyed to write a paper with Mira and Maya. Tom Anderson gave me great advice – don't work on the research paper until a week before the submission deadline, and only if you already have an outline of your generals paper. This worked *very* well – not only did it give me an incentive to start writing early, it also limited the time I could spend on the research paper. Don't assume this will always be true for you; we already had a lot of the work done before that week, and we weren't aiming for the best conference in the field.

Most advice also said "use the whole clock." You have two months from charge to exam, and have to schedule four (or more) faculty and a conference room, so schedule the exam more than just the required 3 weeks early. I arranged to present ten days before my clock ran out. If you're studying a new area, this may be more of an issue, but as a thesis proposal exam, you may be able to get away with using less time.

I've also been to a lot of generals practice talks. Make sure that you do because it will help you shape your thinking toward synthesis early on, well before you start preparing the talk. Also, make sure that you invite younger students to your practice talks to keep the community strong. There's much more learning about how to give talks in the practices. (Obviously, how to answer vague faculty questions is something to learn at real exams.)

3 My failures

I'm embarrassed that I didn't produce a good proposed thesis. My excuse is mostly that other graduate students around me have very little experience with thesis proposals. I'll try to describe some of the mistakes so that you might avoid them or help those around you. And if you read my generals paper as an example, you won't mistakenly believe that it made a good thesis proposal.

I was not prepared to answer the questions:

- how do you know when you're done?
- how do you define X in your thesis statement? (X may be "useful," "accurate," "efficient.")

That is, I had a crummy thesis statement. It wasn't broken down, it wasn't defined, I hadn't even thought about how to evaluate success. It's possible that I'm harder on myself than is warranted; just be sure to outline for yourself what tasks you'll need to accomplish, what results you'll show, and how you'll *prove* that what you've shown is true.

Obviously there are other questions, but by now you should know enough to answer them –

- why is this important? who cares about the result?
- why is this the right way? (as opposed to spending money or buying a faster machine)
- how do you know you're right? (validation of measurements)
- how does it relate to X's work on Y?
- why haven't people already done this?
- what are the burning questions in the field and how does this help?
- why is this science?
- is there a bigger, more ambitious goal beyond what you've proposed?

Finally, I didn't sleep well before my exam. I woke up at 5:30 in the morning, which is something I never do. I was sleepy and doing the best I could to caffeinate myself.² I should have been more intent on finding an afternoon slot for the exam. Afternoon slots are hard to come by, as these are faculty you're trying to schedule, and afternoons are often booked. But, scheduling for the afternoon could help if you're not a morning person or full of life and energy at all times of the day.

4 My successes

While I botched the thesis statement, apparently I did well in the ordinary generals requirements. I focused on the techniques, not the results. That is, not what we knew about networks, but about how network measurement tools solved the basic problems of network measurement – heterogeneity, scale, and confidentiality. I also think I erred on the right side with the thesis proposal – too risky and ambitious, (therefore exciting) but not clear that it will work out as planned (therefore worrisome).

Start writing early. Have an outline complete by the end of the first month. For giving the paper to your committee, while the official standard says "two weeks" before the exam, apparent convention is only one.

²There was no coffee in the lounge that day either!

Write densely. Not hard-to-read densely. Just don't waste space; it gives you an excuse or the benefit of doubt for glossing over things when you can't present detail.

Practice your talk early. A week before. It forces you to get going, to get something. And it gives you a chance for a second practice.

Don't worry about the agreement between talk and paper. You can explain things more visually in a talk, which means you can explain different things.

Use Bibtex. The `annotate` keyword allows you to put your own private text describing the research contributions of a paper and what previous work it builds upon. This is somewhat important early on, because you'll forget. But don't spend so much time writing summaries for your `.bib` file and forget what you're supposed to be working on.

Get a *three-ring binder* to hold the papers you're reading. Mine was a 2½" D ring style, though you might prefer something even larger if you read more or longer papers than I did. Carry it with you wherever you go. When you're bored, pull one at random, skim it, and make sure you understand how it fits. After reading later papers (sometimes with exaggerated contributions), the earlier papers may uncover the threads and techniques that were already there or that you were expected to know about already. When you have a classification of sorts, use dividers to organize the papers.

Find relevant work your advisors haven't heard of. In networking, we have one selective conference, SIGCOMM, whose program everyone knows, and one five-track monstrosity, INFOCOM, where treasures are often lost, along with a series of other grab-bags. Such unknown work can be discovered when you have a plan for future work or are searching for work that fills a cell in the table.

Remind the audience of the dual role in your talk. I subtitled my talk "a thesis-proposal flavored generals exam" because I wanted to make clear that it was a generals exam in substance and a thesis proposal in seasoning. It is probably a good idea, at least until such exams are common.

5 Questions for your advisor

Generals is about meeting your advisors' expectations without their supervision. Grok that sentence and realize that it's up to you to learn what those expectations are before (or soon after) you get your charge. What makes a good thesis statement? What should be the balance between understanding context and past work (traditional generals) and motivating future work (thesis proposal)? Keep in mind that you only have 20 pages and 40 minutes and you won't do a complete job of either. They may forget what guidance they gave you, which may make you seem resourceful and clever, so take whatever advice they are willing to give.

6 Diffs to communal wisdom

I want to briefly describe some conflicts with the common wisdom on generals that are specific to the thesis proposal style.

Forget the summaries of papers. I don't think there was a single paper I spent more than a paragraph on, and most only about a sentence. And then, it was all about how the techniques in the paper contribute to the greater understanding sought in the context of my proposed thesis.

Some suggest that you work on something your committee knows little about, perhaps to avoid baggage in how they use terms, or to more easily cover up an incomplete literature search. In a thesis-proposal generals, you will not be able to branch too far from your advisor's area of expertise. Just try not to branch too far from your own.

7 Technical hints

`psrip` and `pdfimages` are tools to lift figures from postscript and pdf papers. This is useful if you want to show figures from papers you've read. I used such tools to extract some crummy looking network maps to motivate my work.

`sidewaystable` is a latex environment that renders a table rotated sideways on a page. This is very helpful if you need a large table to capture a fine-grained classification scheme. It might be in the `rotating` package.

I also used emacs outline mode to help syntax-highlight my thoughts. And, I kept downloaded papers in a subdirectory, which made it easy to recognize which papers I'd already "read."

8 Summary

The thesis proposal style of exam can make a late generals meaningful. I can now talk with just about anyone about the context and goals of my work, something I really couldn't do well before. This means that I'm becoming more ready to interview and more assertive in the quality of my own work. It also gave me a chance to write something on my own, without coauthors rewriting things, so my work alone was being evaluated. No matter what you do, assuming you pass, you can expect to have the same increased confidence and sense of accomplishment.

My talk and paper are at <http://www.cs.washington.edu/homes/nspring/talks/generals.pdf> and <http://www.cs.washington.edu/homes/nspring/papers/generals.pdf>. Feel free to send questions or comments to <mailto:nspring@cs.washington.edu>.

Take it seriously, do your best, be creative, be light, don't stress.