

Getting a Ph.D. at the University of Southern California

Yuriy Brun

brun@alum.mit.edu

<http://alum.mit.edu/www/brun>

1 This document and my story

The purpose of this document is to summarize my experience as a computer science doctoral student at the University of Southern California, in hopes that others may benefit from my knowledge and make fewer mistakes than I did. Some of the information here will be specific to my case, some to research in software engineering, some to the Computer Science Department, some to the Viterbi School of Engineering, and some to the University of Southern California. Further, some information will likely become dated, and it is the responsibility of the reader to base his or her decisions only on the proper subset of the information I present here.

To help the reader, let me describe my story of becoming a Doctor of Philosophy. In the summer of 2003, I received a Masters of Engineering in Computer Science from MIT, and in the fall of that year came to USC to do my doctoral work. I started as a School of Engineering (at that time, not yet carrying the Viterbi name) fellow and worked, for roughly three years, with [Prof. Leonard Adleman](#), on theoretical computer science and molecular computation. After that, I worked for two years with [Prof. Nenad Medvidović](#) on applying my ideas to software engineering to design large distributed software systems. I defended my dissertation on March 11, 2008, and officially received my Ph.D. in May 2008. One aspect of my story that is likely to greatly influence the information in this document is that my plan is to remain in academia, work as a post doctoral researcher for roughly two years, and then pursue a position as a professor at a research university.

The rest of this document is split up into three sections and a summary. Section 2 discusses the first few years of the Ph.D. program, up until about two years before the expected graduation date. Section 3 covers the penultimate year of the Ph.D. program, and in particular, taking the qualifying exam. Finally, Section 4 covers the final year of the program, which includes writing and defending the dissertation.

It is not the intent of this document to devote to each topic an amount of space proportional to the importance of that topic. There are numerous resources out on the web, most notably the very complete [Tao Xie's collection of advice pages](#), and the less complete but often much more succinct [Michael Ernst's advice page](#), as well as many others. My goal is to provide some information specific to the USC Computer Science Department, and augment some of this other advice where I deem necessary. Thus, while I hope to give a good overview of what a Ph.D. student needs to do, I will go into depth on some subjects that may perhaps be less important than others I will only mention. Where possible, I will try to identify such topics, and reference other sources of advice.

This document has a number of hyperlinks to websites with further information and the reader should note that these are not visible as hyperlinks when this document is printed in grayscale.

2 The early years

This section covers what a student should do, in my opinion, during the first few years of the Ph.D. program, until roughly two years before graduation. The length and intensity of this period will vary from student to student. The average stay for a Ph.D. student in Computer Science keeps increasing (when I started in 2003, the national average was 6.5 years, while USC's average was a bit shorter), and often, the majority of that time is spent in this period. A student can do quite a bit to shorten his or her stay by becoming more focused early, and avoiding some common pitfalls during these early years.

2.1 Getting to research

While the Ph.D. program has a number of requirements, such as a breadth requirement of taking a number of classes across the different areas of computer science, a depth requirement of taking more classes within some one area, a screening requirement, the qualifying exam, and the writing and defense of a dissertation, the only real requirement is that you do research. If you work hard and do good research, you will see the requirements become satisfied one after another. The main point I want to make is that if your goals are to get a Ph.D., become a researcher, and especially become an academic (though the advice still applies without the latter two goals), at no point in your time as a Ph.D. student should you “focus on classes.” Some people make the mistake of thinking “let me get the classes out of the way over the first one or two years, and then I can get to research.” If you do that, you will find that you have wasted your first two years and will essentially be in a position similar to a fresh incoming student, despite working hard for two years. I use the word “wasted” a bit lightly — you will have learnt plenty about computer science, possibly have had a lot of fun, and met many interesting people while taking classes; however, you are unlikely to be any closer to getting a Ph.D. than you were before.

Your first and most important task as a Ph.D. student is to find an advisor. Some students are admitted to the Ph.D. program with a particular professor bidding to be that student's advisor; others are admitted as free agents. While the department should help, in reality, it is your job to find an advisor, and make sure he or she does work you will be happy doing for the next five or more years. If you are an early Ph.D. student and find yourself without an advisor, all your efforts should go toward finding one. There are numerous ways to do so. One way particular to USC is to attend the first-year seminar (you have to do this anyway), which brings professors who might be interested in new students for weekly talks about their research. If you find yourself in the position of searching for an advisor, you should read the advice compiled on [Tao Xie's website](#), [Michael Ernst's advice](#) on this topic, or any number of books or advice pages from other researchers.

2.2 Passing screening

Within the first two years of your program, you are supposed to take the screening exam. Despite the name, you in no way participate during the actual exam. The entire (or the majority of) CS faculty get together and “discuss” you. In reality, this most likely means they look at your grades, and then your advisor gets up and says something along the lines of “yes, yes, he or she can do research and I will advise him or her” (although the advisor should probably use just a single pronoun or the response might raise some flags). It is your responsibility to: (1) maintain a proper GPA to pass screening (currently it is 3.5), (2) stay current on the requirements for when to file for screening (currently before the end of the 4th semester), (3) fill out and submit the screening application form on time (include a high-resolution photo so that professors who do not recognize

the name may recognize the face), and (4) ensure that your advisor attends the faculty meeting. Of all these responsibilities, (4) is crucial. If the advisor fails to attend, you will fail your screening. You should talk to your advisor ahead of time to ensure he or she is happy with you taking the screening exam and make sure he or she is aware of when the meeting is. It is appropriate to send an email reminder the day before the meeting to make sure it does not slip the advisor's mind.

2.3 Research and publishing

I mentioned before that your primary responsibility is to do research. At first, your advisor will likely have to formulate problems and methodologies for you, or you may join existing projects with other students. Over time, you will pick up the skills necessary to select your own problems and formulate your own methodologies. Eventually, you will formulate a large-enough problem in the area of your work to write a full-bodied dissertation. Make sure you do not wait until that time to start publishing. The only footprint you leave in the academic world (at least until a few years after you become a professor) are your publications. If you are working on projects in your group, help write papers. Look for workshops in your area of research and write down your ideas into position papers or research proposals for these workshops.

One important, and not always visible piece of advice is to aim high with the quality of the conference or journal for your publications. When you first start out, workshops might be appropriate; however, publishing in workshops will do little to help you graduate or get a job (and they may, in some circumstances, hurt your ability to publish in high-quality conferences). Aim for the top several conferences in your area (e.g., ICSE, FSE, FOCS, STOC, etc.) — publications in these will be your currency as a researcher as you near graduation. I should note that attending workshops and publishing periodic position papers can be very helpful to meeting people and advancing your research and career, so take every opportunity to attend these, but be wary of publishing your main work at these less-reputable venues.

There are numerous sources of advice on how, where, when, and why to publish (e.g., [Tao Xie's collection of advice](#) and [the SIGPLAN author guide](#)), but I do want to say one thing myself. **Use the scientific method in your research!**

The scientific method applies differently within different fields and to different problems, but in one view, the scientific method consists of formulating hypotheses and testing those hypotheses through design and execution of experiments and data analyses. Some of computer science, and particularly software engineering, have become less of a science and more of an engineering discipline. While there are positive effects of this transition, such as a close relationship between the academic and industrial worlds, there are also some very detrimental effects.

Forming a hypothesis

It provides little benefit to the research community to publish a paper that says “I thought hard about a problem and came up with this solution, which I think solves this problem.” I estimate that roughly 50% – 60% of the papers I review are of this variety. The problem is that this paper, and the underlying research, neither poses nor evaluates a hypothesis. A more proper research project would state “I examined a clearly defined problem, identified the shortcomings of several existing solutions, and reviewed the literature to be satisfied that no other solution without these shortcomings exists. I then hypothesized that by following a novel approach, I can generate a solution without [some of] these shortcomings. I then evaluated the new solution against the existing ones.” You might be surprised just how many of today's research projects fail the litmus test of whether one can summarize the project within the above framework, and also how easy it

might be to identify shortcomings of your own work by trying to explain it using that framework.

The missing evaluation

Even worse-spread than skipping forming a hypothesis is the problem of the forgotten evaluation. I estimate that roughly 90% – 95% of the papers I review have insufficient evaluation. Again, it is not acceptable to simply say “I thought hard about this problem and my solution.” You must design *experiments* to test your hypothesis, *compare* your solution to other existing solutions *quantitatively* where possible, and if you must be qualitative, you must design a well-defined *methodology* and *framework* for the compared qualities. In the rare cases where there is no existing solution, there almost always is at least a naïve solution awaiting a comparison.

By following the scientific method, developing a hypothesis, and properly evaluating your work, you will discover many more interesting deep questions and greatly improve your research than you do otherwise. Further, the extra work you put in to follow the scientific method will pay off with higher quality research, better respect from your peers, and better acceptance rates on your papers and, eventually, funding grant proposals.

2.4 Being professional

As a Ph.D. student, you will interact with other researchers. Most notably, your advisor, your class professors, and your committee, but also researchers from other labs and universities. It is important to be professional in your interactions with them. Unprofessional communication displays your collaborating abilities in a poor light and will likely lower your chances of success. Professionalism can make you stand out over other students and go a long way in assisting your career. Whenever you send an email to a professor, including your advisor, follow these guidelines.

- Make sure you read the email twice, carefully. I see students misspell words, misuse grammar, and drop entire clauses from sentences rendering those sentences nonsensical. Read the email out loud, or, better yet, have a computer program read it out loud for you. I frequently ask a friend to read my email before I send it to a professor.
- Check the spelling of the names in your email. Most notably, make sure you spell the name of the professor who you are addressing correctly. Find his or her website and copy and paste from there. Many people who email me misspell my name. I try not to hold it against them, but I certainly notice it every time and it tells me just how much thought they put into their email and how important I am to them.
- Use the proper title: if the addressee is a professor, then call him Prof. Addressee; if the addressee holds a Ph.D., M.D., or some other doctorate but is NOT a professor, then Dr. Addressee is appropriate; otherwise you can simply use first and last name, or Mr., Mrs., etc. is the way to go. Avoid strange contraptions some students come up with such as “Prof. First_name,” or the German version “Prof. Dr. Last_name.” In America, the highest title is Prof. Last_name, which trumps Dr. Last_name.
- Make sure you are saying everything you want to say. It looks highly unprofessional to have to send two or three emails within an hour correcting what you said earlier.
- Be cordial when addressing professors. Whether or not you are happy with them right now, they have gone through a grilling and challenging process to become professors primarily because they enjoy educating students like you. While various circumstances can go awry,

resulting in what you may sometimes think are unfair actions on behalf of some professors, their original intentions are always positive and selfless, so they deserve your respect.

2.5 Attending conferences

I will not go into great depth on the subject of presenting your work, as the before-mentioned advice resources do more than an adequate job on that task. I will say that you must make every effort to travel, attend workshops, symposia, and conferences. Talk to people! This is hard. Everyone says it is hard. Everyone says force yourself to do it anyway. Do not just agree with that advice — find a way to force yourself. Go to workshops without your comfort friends. Sit away from your friends. Invite yourself to go to dinner with people you do not know. Do not let yourself hang out with your friends and advisor at conferences. Spend minimal time in your room. Come to a conference with your talk prepared — valuable hours at a conference can be spent practicing a talk, while the same practice could have been accomplished during far less valuable time at home.

Having said that, **practice your talk!** After following the scientific method, this is the second most important advice I can give. Again, I estimate that around 75% of the talks I hear at even the highest-ranked conferences are unacceptable. You, as the author, get 30 (or so) minutes of other people's time to impress them with your work, make them want to read your paper, and make them reference you. As a student, you get at most one or two of these opportunities annually, and often far fewer, making these half hours roughly 17,500 times more important than your average 30 minutes. That, in turn, means it is reasonable to spend 20, 40, or 100 hours preparing for your talk, and still come out with a time-bargain on your hands. As a student, I estimate that you should practice a talk at least ten times before you give it at a conference. That is not an overestimate! Do it three times yourself. Do it for your research group. Follow the group's suggestions and give it a couple more times to yourself. Get a larger audience at your university and give the talk again. Finally, practice it a couple of times the day before you give it at the conference. If you find yourself not knowing what is on the next slide before you flip to it, you are not ready.

I really cannot emphasize enough how important it is to practice your talk and to give a good one. If you give good talks, you will be able to publish more, graduate faster, and get a better job. Of most relevance to this document is the ability to graduate faster. Experience giving good talks will boost your chances taking the qualifying exam and defending your dissertation like you would not believe.

There is a wealth of advice out there on how to prepare a good presentation. As before, [Tao Xie's advice compilation](#), [Michael Ernst's web site](#), and [the SIGPLAN author guide](#) contain a number of good pointers. The majority of people who read this advice, as well as books on giving good presentations, seem to ignore it. Do not ignore it! This includes, but is not limited to:

- Minimize the text on your slides. A good presentation has almost no text on it and plenty of diagrams. This is not an exaggeration! Research (which I do not bother to locate and reference here) has shown that people can pay attention, listen, and learn while eating, fantasizing, doodling, and even sleeping, but they cannot pay attention to you while reading. By leaving text on your slides you are shooting yourself in the foot. Some people mistakenly use the argument that the slides should tell the story on their own — that is not true. The slides are a guide and a helping tool for you to give a good presentation. Without you, the slides need not be able to present your research. Your paper does that, which is why it contains words.
- Speak from the first person. Both in papers and in your presentations, use first person instead of passive voice or third person narration. It has been several decades since the arcane belief that first person is informal was deemed, well, arcane. Say “we built ...” for

multi-author work and “I built . . .” for single-author work. Using first person will make your communication clearer, easier to understand, and make more people acknowledge the great work you do.

- Do not clutter the slides with your school’s name, your department, today’s date, etc. It is appropriate to say some of these things on the title slide, but do not throw this information in the footer and header of every slide. Further, many schools, including USC and the Viterbi School of Engineering ask you to put the appropriate logo on the slides and to use their templates. Again, it is fine to put the logos on the title slide, but the templates are usually terrible and will hurt you rather than help you.
- Keep the slides simple. No complex backgrounds, no flying objects. In general, I strongly discourage animation. The rule of thumb should be: If the animation adds to the content of the talk, use it; otherwise, even if it seems an animation makes something “prettier,” it really just distracts from the substance of the talk.
- Keep to the time limit. You can destroy a great impression by speaking longer than your time slot allows. Practice your talk and if you seem to be going over, your *must cut material and slides*. It is not sufficient to promise that you will talk less or cover a slide or two in less detail. Cut out several slides. Remove some bullets / diagrams from slides. You must prepare for the proper-length talk beforehand. Many people use the one slide = one minute rule of thumb when preparing their slides, but you should remember that is a lower bound. A good slide with a diagram and no text takes around three minutes to describe. In an hour-long talk, it is not uncommon for me to spend ten or so minutes on a single pivotal slide that explains the meat of my approach.
- Avoid the laser pointer. There are very few cases when one actually needs to use a laser pointer. If at all possible, you should avoid using one except in these rare situations. While there is nothing wrong with using a laser pointer, having one in your hand makes it far too easy to perform a number of regretful actions. I have literally walked away from the talk remembering nothing about it except that the person underlined everything he or she said with a pointer, and shined that pointer into several people’s eyes. It is very easy to turn your audience into a cat whose attention is drawn only to the moving red dot on the wall. That, in turn, makes your job of focusing the audience’s attention on your work nearly impossible. It is appropriate to use a pointer to point out pieces of diagrams or figures about which you are speaking, though some clever slide design can remove that necessity by highlighting the proper pieces at the proper times.
- See Section 3.2 for comments on using an Outline slide, and on generic headings and titles. In short, you should avoid headings and slide titles such as “Evaluation” and make them specific to your work.

2.6 Do summer research internships

Doing summer internships can greatly open up your research horizons, introduce you to new people, create employment opportunities for after graduation, and is in general a very good idea. In order to help in getting your Ph.D., you should only consider research labs (e.g., IBM research, Microsoft research, Google research, though the latter has been slow on implementing a summer intern program). Avoid working for the industrial cousins of these companies’ research branches. Your aim, when you go to a summer internship at a company, should be to publish a paper by the end

of the summer. Anything less will likely be a bit of a waste of your time. That is a somewhat extreme view, and you can certainly benefit from an internship that does not result in a paper, but your goal and motivation should be to get enough research done to write one.

You should only do these summer internships during the first few years of your Ph.D. program. By the time you are in your final two years, you will need the summers to work on your dissertation, and doing an internship can slow down your progress. There are, of course, exceptions to this guideline.

2.7 Up-to-date website

Every computer science Ph.D. student should have a website. This website should be simple and informative. Follow the basic outline that most professors in your area use. Feel free to use [my website](#) as a template. Make sure there is a clear and professional photograph of you on the front page so people can identify you when they meet you. While you can be creative in your web page, it is very important to make your page simple and easy to navigate. Following a commonly used scheme will make it easier for people familiar with that scheme to navigate, so do not stray too far from it unless you have a real dynamite idea. While I have the utmost respect for Tao Xie, and the information he provides on [his website](#), I would recommend staying away from the cluttered style of information presentation he has followed. Oftentimes, less is more. Get the important information clearly visible, and avoid adding unimportant data.

Your website should include a .pdf of each of your papers. This has become the format of choice, and the easiest format for many to access. While some people still put .ps files on their websites, it can be an annoyance for many to have to convert in order to read. I have seen people skip .ps papers and read only .pdf papers for that reason! Many publication venues have strict copyright rules about disseminating published work. Some of these rules include a provision for the authors to distribute copies for educational and research use only. Whether legal or not, I have found it to be typical to assume such a provision exists, and post all your papers on your website. It may be appropriate to include a copyright notice on your website regarding these papers ([example](#)).

Make sure the website is *up-to-date* at all times, and include a current curriculum vitae (again, follow your advisor's, other researchers', or [my curriculum vitae](#) as a template). There is no excuse for a student to not have an up-to-date website, and having one will help you reach a broader audience with your research and graduate faster.

2.8 Attending seminars

Go to the CS seminars held almost weekly at USC — these are professors, or top potential professors, coming to talk about their most recent research. Attending these is the best way to stay current with the major directions in computer science, and to see applications of your research you would never discover on your own. Seminars is a horribly underused resource at USC. Very few people attend these talks, which reflects poorly on USC and also prevents our research from reaching the heights it could and should reach via exposure to ideas outside of our small circles. Take advantage of this great opportunity the department gives you by setting aside the one or two hours per week to attend a seminar — you would just waste the majority of that time surfing the Internet for useless amusing articles.

3 Penultimate year: Qualifying exam

The qualifying exam (qual) at the USC Computer Science Department differs from what many other departments, both at USC and elsewhere, refer to by the same name. At the USC CS department, the exam is typically taken roughly one year before graduating (\pm a semester). It consists of writing somewhere between $\frac{1}{2}$ and $\frac{2}{3}$ of your dissertation, submitting it for review by your five-person “Ph.D. Guidance Committee” roughly a month before an oral exam, and a two-hour closed oral exam before that committee. As you might expect, this is a major piece of work. Some professors consider this to be “the real dissertation defense,” as opposed to the one held at the end of your program.

In essence, your qual is like a final defense, but you get to say “I will do so and so” instead of “I have done so and so.” Note that this implies you already know everything you will do, you just may not have done some of it yet. There is no room for “I will explore so and so and decide on how to proceed” except with very minor and noncrucial aspects of your research.

In order to be able to write your qual document, you must have the following pieces of your proposed research:

- A well-defined problem you plan to solve in your dissertation.
- Your hypothesis.
- A complete review of literature related to:
 - the problem,
 - existing solutions,
 - techniques you may use for your solution,
 - techniques you may use to evaluate your solution.
- Your proposed solution.
- A detailed plan of the evaluation, including which experiments you will perform, what you will measure, and how you will know if you are successful.

There are no specific formatting or organizational instructions you need to follow for the qual document, but you will save yourself a lot of time if you lay out the structure of your dissertation and follow the same structure here. The problem statement, related work, and possibly your proposed solution chapters may not even change after the qual. See section 4.3 for information on formatting and typesetting your document.

3.1 Your committee

Your Ph.D. guidance committee consists of five people. This committee contains your advisor and four other professors. In 2007, the guidelines state that at least four of the five professors must have tenure, one must be from an outside department (called the outside member), and at least three must be from your home department. These rules change from time to time, so check for the most recent set.

This committee will evaluate your work, decide whether or not you may proceed with your Ph.D., and, ultimately, allow you to graduate. Thus who you ask to be on your committee is important. You should keep an open conversation with your advisor about your committee from

the time you have a research topic for your dissertation. These are the people who decide your fate, so select them carefully and make sure they know your best side.

The defense committee, who attends and decides the outcome of your defense, is usually a three-person subset of the Ph.D. guidance committee.

3.2 Oral exam

Your oral exam will last two hours. You should aim to present your work in roughly the first hour and field questions for about half an hour. The exam is closed, meaning that only the five members of your Ph.D. guidance committee and you will be present. At the start of the exam, the committee chair (your advisor) will ask you to leave and the committee will briefly discuss you. After about five minutes, you will be invited back in, and be asked to present your work. The committee will interrupt you with questions. They will expect you to be able to answer hard questions, controversial questions, and be able to stand up for the statements you make with certainty and conviction. If you are very good, you will get through your presentation within an hour, and the committee will ask you follow up questions. More likely, the committee will have many questions along the way, will interrupt you constantly, and it will take an hour and forty five minutes before they ask you to skip the rest of your slides and conclude. In some cases, the committee will essentially force you to abandon your talk early on (as early as on the second slide in some cases) and question you about your research directly. Once the committee is satisfied (or more likely tired), they will ask you to leave the room and discuss you some more. When you return, you will either be congratulated for having passed or told that you failed. Each committee member will share his or her concerns for your research and what he or she thinks you can do better. If you fail, your advisor will explain your options, which largely are to take the exam again, most likely next semester or a bit after that. It is common for professors to say that they must run but that they want to talk to you and give you feedback later. In that case, it is your responsibility to follow up with them and make sure you get the feedback. If you do not, they may be irked at you and that is likely to hurt you when you defend.

You should make sure that every member of your committee has your qual document at the exam. Do so by bringing extra copies to hand out. You should also make copies of the print-outs of your slides (six-per-page), one for each member.

Your talk should consist of two parts: the overview, and the details. Start the overview with a brief description of the problem you want to solve and why it is important. Then state your hypothesis about the problem and in a few words describe your approach to solving the problem. Summarize the evaluation you will perform, and show the contributions you will achieve if you are successful. This overview should take no more than ten minutes (no more than five slides). Incidentally, this overview approach is a very good idea for most hour-long talks, including your defense and your job talks.

After the overview, present the audience with an outline of the remainder of you talk (more on this below), and then give the details of why your problem is important, the related work, your approach, the proposed evaluation, and some initial results. This part of the talk is similar to the traditional talk.

I have heard some people complain about the outline slide wasting the audience's time. Others, including myself, think the outline slide is crucial to orienting the audience to what is going on. After my overview, I always present an outline slide, and flash back to it every time I finish a major (and sometimes minor) section. I believe a great compromise between the two schools of thought is to make sure that your outline slide has no generic headings. In other words, an outline such as:

- Proposed problem

- Related work
- My approach
- Evaluation
- Summary

is in fact a waste of people’s time. If, however, your outline is specific to your work, it can convey important information for the audience. Here is a sample outline, modified from my defense:

- Distributing computation on the Internet
- Previous distribution and discreetness attempts
- The tile architectural style
- Proving correctness and building a tile-style system
- Contributions.

You will likely see the advice of not using generic terms such as “Evaluation” and “Related work” as titles of your slides. I fully agree with this recommendation. In general, I strive to make sure that no one could easily reuse a title of one of my slides in an unrelated talk (except “Contributions”). Seeing mundane and general information on the screen puts the audience to sleep. Personalization keeps it awake.

3.3 Qualifying exam timeline

You should speak with your advisor early and often about your plans to take the qual. I knew my target date about a year before I took it. Your advisor and you can lay out an outline of what needs to be done before you take the qual. It is very common to have to push the exam to a later date. It is almost unheard of to say “hey, I think I am ready to take my qual now,” and be able to take the exam less than a year from that first mention, regardless of how much work you have done. Other than getting on the same page with your advisor early, you must do the following things, in roughly the following order.

- Form a Ph.D. guidance committee. You should discuss with your advisor who should be on your committee at least six months before your exam. The six month date is a lower limit, and it is preferable to discuss the committee a year or more in advance.
- Schedule your qualifying exam. You need a two-hour block of time. See my description of scheduling the exam with your committee in Section 4.1.
- Declare your Ph.D. guidance committee. You must fill out a proper form, which each member of the committee must sign, signifying agreement to serve on your committee. The department requires you to file this form roughly a month before the exam. Do this early. Getting the five signatures is hard enough, but after you submit the form it must go around and get approved by the department’s advisor, signed by the department chair, and even the dean. I know of very few cases that went smoothly, and of several, including my own, that imploded with complications and almost ended in tragedy. So do not wait until the last minute to do this; form your committee and submit the form a semester before your exam. Incidentally, you should ask the department’s graduate advisor for the committee formation form. In my experience, this form changes slightly, and for no good reason, every semester, and the administrators are sticklers for making sure you use the latest form.

- Discuss your document outline with your advisor and start writing. These steps should begin a semester before your exam. Try to get the advisor to approve the outline and give you meaningful feedback, to make sure he or she thinks at least some about it. I had sent my outline to my advisor and got it approved without much discussion for both the qual document and my dissertation. The result was that I wrote three chapters, about 100 pages, that I did not end up using in either document, and I could have saved myself a lot of time and work if I had forced my advisor to think in more depth about my outline and document early on.
- You should give a fairly complete draft of your document to your advisor roughly two months to six weeks before the exam. It is a good idea to give him or her parts of your document as you finish them, but you should have a completed draft by this deadline. Most people miss this deadline, and that can cause numerous complications, and often a delay in taking the exam.
- Having gotten comments back from your advisor, you should satisfy his or her concerns and distribute the document to the rest of your committee a month to three weeks before the exam. An ideal way to do this is to email them a digital copy, and ask to set up an individual meeting. At these meetings, you should describe more or less the overview part of your talk, and see if they have any questions or concerns. You should also bring a printed copy of the document with you to the meeting — it will greatly increase the chances that the committee member reads at least some of your document, and make your life easier during the exam. Some people think that if a committee member does not read the document, he or she is unlikely to ask hard questions and the student is better off, but that is a mistake. Try attending a talk on material you are not familiar with and a talk on material with which you have some familiarity. The latter talk will seem infinitely better to you than the former talk. If the committee member does not understand portions of your work, he or she will blame you for your inability to explain them well.
- You need to reserve a room for your exam. Ideally, do this more than a month before the exam so that you can tell your committee the room and time when you email them your document. It is also appropriate to email reminders to your committee a week before the exam, and the day before the exam.
- Also roughly a month before the exam, you should let the department graduate advisor know of the date you plan to take the exam. Make sure he or she has all the documents needed to take the exam!
- Practice your talk. Start making your slides a full month before the defense — do not wait until the last minute. The earlier you start making the slides the better. Set up a session with your research group and a wider-area group and make sure you know the talk by heart, as I described in Section 2.5.
- The day of (or the day before) your exam, your advisor should pick up your documents from the department's graduate advisor. These documents contain that committee formation form, signed by all kinds of folks, and a form that needs to be signed by your committee after the completion of the exam (whether you pass or fail).

3.4 Doctoral symposia

In software engineering, the highest-ranked conferences have doctoral symposia. These are fairly competitive venues that, if accepted, allow you to take a qualifying exam in front of ten or so

professors from outside universities. This can be a very helpful process for three reasons. First, this experience can help you prepare a much stronger dissertation. Second, being part of a doctoral symposium helps you meet well-known professors in your field, exposes them to your work, and shows you how to present your work to the outside world in the best light. Third, the experience allows you to publish a short paper (two to four pages) on your work at a relatively competitive venue. For an example of a short paper, see [my paper](#) from the ICSE 2007 doctoral symposium.

One should note that being a part of a doctoral symposium can greatly help you with your qualifying exam, and taking your qualifying exam can greatly help your doctoral symposium experience. As a result, you should do the one that is less important to you first. If you are worried about passing the qual, try to attend a doctoral symposium to get early feedback and improve your presentation. If you feel confident about your qual, do that first so that you can impress the doctoral symposium committee. I encourage the latter option, as it can greatly increase your visibility with the outside world and form long-lasting academic connections.

Typically, doctoral symposia provide fairly extensive funding for their participants, so there should be few reasons not to attend one. As I mentioned before, these symposia are competitive, so you are likely to have to apply to several in order to get accepted. The application process consists of the short paper on your work and a recommendation letter from your advisor.

4 The final year: The defense and the dissertation

On average, it takes a year from the time you pass the qualifying exam until you are in a position to defend. If you are careful, meet with your committee somewhat regularly, and keep the committee informed of your progress, the final defense should be more or less a formality. That is to say, the committee will not let you defend unless they think you are ready. I have, however, seen multiple people ignore their committees, or chat with the members from time to time but not keep them up-to-date on the research, and find themselves surprised to be failing their defenses. Do not let that happen to you.

Your dissertation committee is between three and five professors. One of them is still the advisor, another is still an outside member. I would form a committee of three people, each of whom were on the Ph.D. guidance committee, but discuss this with your advisor. Again, you must fill out a formation of committee form, and half the university must sign this form, so do it early. The deadline to do this is a month before your defense.

Your defense is not much different from your qual exam, and your dissertation from your qual document, except that you need to have completed your solution and evaluation. You will likely have small developments regarding new related work and such, but if you did a good job on your qual, the structure and much of the text from the early sections should remain the same. As with the qual document, you should make sure that every member of your committee has your dissertation at the defense, and also make copies of the print-outs of your slides (six-per-page), one for each member.

4.1 The final timeline

After passing your qual, consider looking for a “final-year fellowship” to support you. These are offered by several companies (e.g., IBM) and USC also has one called the Dissertation Completion Fellowship. These options change annually, so investigate your opportunities and discuss them with your advisor. In addition to providing some funding for you, and possibly freeing up your advisor’s funding to let you travel to conferences, these can be quite reputable and look good on your curriculum vitae.

For this section, I will assume you wish to graduate in the Spring semester. The deadlines are very early, so you must plan ahead to hit them. I will refer to several deadlines as they existed in my semester, but you should visit [the graduate school website](#), assuming it still exists, for up-to-date deadlines.

At the start of the semester in which you plan to graduate, download “the packet” from the graduate school website. For some unknown reason, this packet is split up into a dozen .pdfs, which you must pick and choose to assemble a valid Ph.D. packet. This packet will contain four forms that you must fill out before your defense.

In order to have a chance of graduating in May, you must submit your dissertation by the thesis filing deadline (April 1st in 2008). That means that you must have defended your dissertation, had the committee give you feedback, and corrected the dissertation and satisfied the committee, all by that date. Therefore you must defend at least several weeks before April 1st. I have several friends who thought they could defend March 31st and submit their dissertations the next day, but were rudely surprised to learn otherwise. My defense date was March 11th, and even that, to some degree, was pushing the envelope.

In order to hit a March 11th defense, you should start scheduling the date and time on January 11th. It can take a long time to figure out when your committee is willing to make time for you. As with the qualifying exam, you should schedule a two-hour block. Get a week calendar, look up the teaching times of every professor on your committee and cross it off. Pick a day or two with large chunks of open space, and email the entire committee asking if they are available on those days. People may be out of town or they may be busy with visiting faculty or potential faculty, so give them some options. Once you find a time everyone agrees on, mail it out to the entire committee and be clear that is your time. Ask them to put it into their calendars.

Around February 1st, you need to inform the department’s graduate advisor of the defense date. This means you should also reserve a room. In my experience, you will have to inform the advisor several times over the next month and a half — every time you are asked.

The process of submitting your dissertation to the school’s editor is unnecessarily complex. Once a semester, the thesis editor holds an hour-long workshop describing the process. If you have not attended one of these previously, do so now. I found this to be the single most useful place to learn how to submit your dissertation. Without it, it is not clear I would have been able to navigate the poorly designed website and collect all the necessary forms.

As described at the thesis editor workshop, you need to create a GradSchl profile. Do this in February. The website is terrible and confusing, the directions are nonexistent, and the regulations are absurd (e.g., you will be asked to create an account and be given a cryptic password you cannot change and once you enter some information, such as you address, you will not be able to alter it, so be very careful filling out the forms). Once you create a profile, you will get almost daily emails telling you that they have not yet received your dissertation. These emails become welcome comic relief after a few days.

On February 11th, email the department’s graduate advisor and ask him or her to send out an announcement of your defense. This step is required by the university and technically, you cannot graduate without it. I had a lot of difficulty making this happen, and in fact it never did, despite my repeated attempts, so do your best. No one, except my advisor, noticed and the university let me graduate. The announcement needs to contain an almost-final abstract, so run that by your advisor. Adhere to the word limit. Your abstract will serve as the face of your dissertation, so make sure it is complete and well-polished. Do not cut it short but do not include so many details as to blur the big picture, either.

This entire time you should have been conversing with your advisor about your dissertation. Around the February 11th mark, you should email your committee a digital copy of your dissertation

and ask to meet with each member individually, just as before the qual. Bring a hard copy of the dissertation to the meeting. It is appropriate to continue working on fixing your document based on the feedback you receive, but avoid mailing new copies of the document to the committee. I sent out a copy with one chapter missing, and specifically instructed the committee that I will supplement that copy with the missing chapter later. Otherwise, you are almost better off with your committee having a *slightly* outdated version than to act unprofessionally and confuse them with several different versions. The bottom line is that if you are timely and professional, your dissertation should be finished around mid-February.

The day before your defense, your advisor needs to again pick up some paperwork from the department's graduate advisor. You should augment this paperwork with the four filled-out forms (depending on the department's graduate advisor, your file may contain empty copies of these forms). I downloaded the .pdf copies and filled them in digitally, since one of these will actually become part of your dissertation and it is nice if it looks professional.

If you pass your defense, each member of the committee has to sign two of the four forms. Make sure they do, I forgot to get signatures on one of the forms and had to track the professors down later.

Before the submission deadline of April 1st, you must go to the graduate school and submit the four forms, as well as the receipt indicating that you paid the publication fee (you need to pay a publication fee, \$110 my year; find out from the department's graduate advisor where to pay this fee). Once they process your forms, you will get an email from the GradSchl profile website telling you that your paperwork is submitted (and of course that they have not yet received your dissertation).

You will now have a couple of weeks to finish the required corrections and upload the dissertation to the GradSchl profile website. Do not wait until the last minute, as the website is quirky. If you have questions, you can email the [thesis editor](#). The editor entity (that is actually several (two or three) assistant editors, who, in my experience, are late-term graduate students from the English department) is pretty good at responding to email but gets very annoyed if you call it and interrupt whatever it is doing, which is most likely reading hundreds of dissertations.

Once you upload your dissertation, the editor will read it and check it for formatting consistency. Your dissertation must be formatted according to a set of rules, some of which I describe in Section 4.3. The editor does not examine content — that is your and your committee's job. This also means that the editor will not check spelling or grammar errors. Your dissertation is a large document and it will reflect on you for a long time. I spent roughly two weeks doing little else but reading my dissertation and correcting mistakes. No matter how much work you put into it, you will still find problems later. Make sure that you check spelling one last time before you submit the final copy.

The editor has hundreds of documents to read, but is amazingly efficient. If you upload within one of the last days before the deadline — the busiest time — you may have to wait as long as a week to hear back. Once you do, the editor will list some corrections you need to make, such as correcting your margins, making sure the numbers in the table of contents are not bold, etc. Once you correct these mistakes, you reupload the dissertation (following the confusing directions the editor gives you), and the editor makes another pass. This process can repeat, and typically, if you work with the editor, there are no problems in finishing in time. If, however, you prove to be a difficult case, you may miss the conferral deadline, in which case your graduation date is pushed to a later semester. I discuss what to do if you miss a deadline in Section 4.2.

Once the editor accepts your dissertation, you must upload it to the USC Library. You get some more instructions about submitting it some places, buying bound copies, purchasing a copyright for the work, and other optional possibilities. Following the library's acknowledgment of receiving your

dissertation, the graduate school notifies degree progress that you are done, and over the course of roughly a week, you officially graduate.

If you are a software engineer, you should also submit your dissertation to the [SIGSOFT software engineering dissertations repository](#). This process takes roughly 30 seconds, but requires you to host the dissertation on your own website (which you should be doing anyway).

If you consider your dissertation to be a top-notch piece of work, you may consider asking your advisor to nominate you for the [ACM Doctoral Dissertation Award](#).

4.2 Missing a deadline

If your defense is pushed to a later date, you do not finish the content corrections in time for the April 1st deadline, or you do not satisfy the thesis editor by the conferral deadline, your graduation date will be pushed forward by a semester. This may not be as terrible as it sounds. You still have until the end of classes of that semester to submit your dissertation without having to register and pay for the next semester. Further, since the Summer semester is “special,” you need not register for it in order to submit your dissertation, so if you miss a Spring deadline, you really have until some time in July to submit without paying a penny more, and your graduation date would be August. If you find yourself in, or near such a situation, you should speak to the thesis editor or someone at the graduate school about your options. There is a right way to navigate this system that will make your experience very smooth, but it is also too easy to make a mistake, forget to submit some extension form, or miss some other deadline, which could cause you quite a headache.

4.3 Formatting the dissertation (and other scientific papers)

There are many different tools available to computer scientists to write scientific papers. Among these are various flavors of Word, Adobe FrameMaker, and \LaTeX . My personal preference is \LaTeX , and I present some of the reasons why I like it below. There are advantages and disadvantages to each tool, and it is up to you to make up your own mind on what you should use, but whatever technology you go with, I urge you to learn it well. Poor knowledge of your editor will result in a final product that looks unprofessional and will diminish the chances your manuscript has of being accepted.

I have several years of experience using \LaTeX , and I have used Microsoft Word for at least half a dozen academic projects. I have never used FrameMaker, but know several researchers who use it and like it quite a bit. Most importantly, in my experience, collaborating with other researchers at other universities often requires the knowledge of \LaTeX . The nature of \LaTeX allows it to be used with CVS and SVN, which are integral to multi-author projects. I have written, or am writing, three papers with collaborators outside of USC and not once has anyone even suggested using anything other than \LaTeX .

I strongly discourage you from using Word. Sadly, you will be forced to use Word from time to time when collaborating with other folks and will quickly learn to hate writing scientific papers in Word as it simply lacks the support (or has malfunctioning support) for many required tasks, e.g., well-controlled figure placement (no functionality), cross-references (terribly implemented and malfunctioning functionality), and citations (EndNote works with Word to provide this functionality, and some people swear by it, though I have used it and find it quite lacking). I should say that Microsoft is, of course, always updating its software, and it may be the case that the latest version of Word is better than the ones I have used; nevertheless, in my experience, Word simply lacks the functionality necessary to write professional-looking papers and to collaborate on projects.

\LaTeX is a combination of a hypertext language, like HTML, and a universal programming

language, like Java. It has extraordinary support for scientific writing that has been developed, open source, over the last several decades. The cost of using L^AT_EX (not a financial cost, as L^AT_EX tools, e.g., MiKTeX and TeXnicCenter, are available completely free) is that it has a much steeper learning curve than any WYSIWYG (what you see is what you get) editor, such as Word or FrameMaker; however, many people will agree that L^AT_EX is worth learning and knowing it will pay off greatly over your career. In my experience, knowing L^AT_EX is a requirement because without it, you would have a very hard time collaborating with people on projects and papers. Every conference and journal to which I have ever submitted has provided L^AT_EX style files to format your manuscript the way they want it, without you having to put any work into the formatting. Further, L^AT_EX allowed me to use text I had written for papers in my dissertation without having to make any changes. L^AT_EX took care of changing the formatting. Further, L^AT_EX allows you to seamlessly move text. You just copy and paste the text, and L^AT_EX takes care of changing the section, figure, and equation numbering, updating cross-references, and properly placing the citations.

Like HTML, L^AT_EX allows you to create fairly simple documents with ease, and add functionality as you learn more of the language, so you can get started without too much trouble and become an expert at your own pace. Almost all of [my scientific papers](#), [my masters thesis](#), [my Ph.D. dissertation](#), [my curriculum vitae](#), and this document have been written using L^AT_EX and I give it my highest recommendation. The bottom line is L^AT_EX makes your documents look professional. E.g., L^AT_EX handles line spacing, hyphens, and ligatures (compare “fi” and “fi”) the way professional editors handle them — just contrast this document to a book of your choosing.

As of recently, I have also begun using L^AT_EX to prepare digital slides for presentations. L^AT_EX has a number of advantages over PowerPoint, although PowerPoint’s WYSIWYG capability is quite important for slides. I do recommend L^AT_EX for slides, but perhaps not for beginners. You should have some experience with L^AT_EX before you force yourself to deal with controlling figure placement. Once you are ready to use L^AT_EX, consider the Beamer package — the best I have encountered.

There are many resources out there for learning L^AT_EX, and this document will not become another one. Simply search the web for a tutorial. If you prefer books, the L^AT_EX Companion is a good one that covers both beginner and advanced topics.

As far as formatting your dissertation, using L^AT_EX can save you a lot of trouble. USC has a set of formatting requirements that you can look up in the dissertation packet (see Section 4.1). These are incomplete! That means that there are a few things the editor requires, but it leaves many of the formatting decisions open to the author. I have updated a previously-existing L^AT_EX style file to conform to the 2008 requirements, and these requirements do not change often, so it is a great place to start. You are welcome to download [my style file and a sample dissertation source](#) and use them for your dissertation. No promises, of course, that the requirements won’t change and that the editor will be satisfied, but these should get you pretty close to the final product.

4.4 What to check before you submit

Over time, I have compiled a list of items I do before submitting a paper. Of course, this list also applies to the dissertation. This is by no means a complete list, but having it physically written down may remind you and me to actually follow the steps.

- Make sure you recently read the abstract and the figure captions. These items are very easy to skip when you are proofreading, but they get more of the readers’ attentions than any other text in your document.
- Spellcheck. Run a spellcheck explicitly. You’ll be amazed how many red-underlined words you missed.

- Search for “e.g.” and “i.e.” and make sure they are always followed by a comma.
- Search for “this”, “that”, “these”, and “those” and make sure they are always followed by a noun. A sentence is rarely unambiguous without these [nouns]. This [grammatical guideline] is not well-known, but should be.
- Search for “Figure”, “Fig.”, “figure”, “Section”, “Equation”, etc. and make sure these are capitalized and abbreviated consistently.
- Search for “its”, “it’s”, “there”, and “their” and make sure you used the correct ones.
- Search for the following contractions: n’t, it’s, ’ve, ’d, ’ll, ’m, and ’re and remove them because they are traditionally too informal for scientific papers and dissertations.
- Search for “which” and “that” and make sure that you used them correctly. Proper uses of “which” without a comma before it, and proper uses of “that” with a comma before it are rare. Look up the proper use of these words if you are not sure how to use them.
- In a multi-author paper, search for “I”, “my”, “me”, “mine” and replace with a plural first-person pronoun. In a single-author paper and in your dissertation, search for “we”, “us”, “our” and replace with a singular first-person pronoun.
- In \LaTeX , make sure you have used the proper dashes. Words are connected with a single dash (- results in “-”), number ranges contain an n-dash (-- results in “--”), and clauses are connected with an m-dash, separated by spaces on both sides (--- results in “—”).
- In \LaTeX , make sure you have examined each of the bounding box warnings you get. Do not ignore these warnings — the paper looks highly unprofessional with an obvious overflow error.
- In \LaTeX , search for “etc.” and “et al.” and make sure if these abbreviations do not end a sentence, that you cancel the next space with a backslash. This ensures that the space following the period is not too large. For example, `dogs, cats, etc.\ all eat food.`
- In \LaTeX , search for the " quote and replace with proper ‘ ‘ and ’ ’ quotes.
- In Word, make sure you were consistent about the number of spaces you used after periods at the end of sentences (I strongly prefer two spaces).

5 Final thoughts

You may have heard older folks say things such as “do not waste your youth — these are the best years of your life.” A similar sentiment applies to academia, though with a fairly specific implication. Your time as a graduate student is your time to do research. After graduate school, you will either enter industry, in which case your time to “have fun and explore via research” will be greatly limited, at best (most likely, it will be nonexistent — not because your boss or your company won’t let you but because you will have a plate full of other tasks, ones for which you are rewarded by your company); or you will become a professor and be more of a manager, guiding other people’s research, rather than doing it yourself. Guidance of your students’ work can be very rewarding, of course, and you get to tackle longer-term and larger-scale and impact goals, but you will rarely get the time to sit down and explore a topic or find an elegant solution to a hard

problem as a professor. As a graduate student, you have little to no worries about applying for grants, getting your peers to worship your work, or dealing with academic service, so do not waste your valuable time. Do research. Do good research. Follow the scientific method, and evaluate your work in a way that will both elevate the quality of your work and make it bulletproof to attacks. Attend conferences, workshops, and symposia to generate new and exciting ideas. Remember that an integral part of doing good research is presentation. It is virtually impossible to put in too much time on improving your ability to communicate via papers and presentations, so practice practice practice. And to reiterate, do good research because that is what will make you have more fun and be more successful.