Process and Requirements for Computer Science Systems Ph.D. Students[1]

The purpose of this document is to outline the process and requirements for earning a Northwestern Computer Science (CS) Ph.D. in Systems in the Computer Engineering and Systems (CES) division of the Electrical Engineering and Computer Science (EECS) Department, and to demystify elements of graduate student life within the group and the department. Please consult your advisor if you don't understand anything here.

Overall Process and Model

After being admitted to the program, you may visit Northwestern to get a sense of the possibilities here. It is very important that you find at least one faculty member with whom you would like to work.

The first thing that you should do after you arrive is to start looking for an advisor. ***Your advisor is of critical importance***. He or she will guide you, help you, fund you, and defend you. We believe in the ***apprenticeship model*** of Ph.D. education. How you learn to do research and your style and taste in problems will be formed in large measure by your advisor. Choose wisely. Although it is always possible to change your advisor later, you can lose valuable time in doing so. Advisors have different styles, but it is common to meet once a week with your advisor to talk about research.

In choosing an advisor, you should talk to any faculty member who interests you. The Department typically runs a "Meet the Faculty" series that is helpful. However, you should take the initiative right away. Generally, you should find an advisor by the end of the fall quarter (your first quarter). The remainder of this document assumes you have chosen a faculty member in the Systems Group[2] as your advisor.

By the beginning of your second quarter, you should be starting to engage in research. In your first year or two, you will also be taking classes, but ***doing research is the critical to your success as a graduate student***. Throughout your graduate student years, at least 50% of your time should be spent on research. ***The whole point of a Ph.D. in CS is to become a good, independent researcher***. The only way to learn how to do research is to do research under the guidance of your advisor and other faculty members. You also want to determine very quickly whether research is for you. Remember, you are not in graduate school to take classes.

At the end of your first or second year, you will take the systems qualifying exam, which is described in detail below. The next step after the qualifying exam is to find a thesis topic. This can take some time and it is easy to get lost during the process. This makes it all the more important to work with your advisor. One you have a good topic, you will embark on the thesis process as described below.

---

[1] This is a living document and is subject to change. The latest version is available on http://nsrg.cs.northwestern.edu.

[2] Broadly defined, this is the group of faculty who advise students for a Ph.D. according to this document. This currently includes Bustamante, Chen, Dinda, Kuzmanovic, and, for some students, Horswill and Dennis.

It is important to note that the Department and the Graduate School also have requirements that need to be abided by. Here is a summary of how certain milestones map:

- The qualifying exam maps to “Admission to Candidacy”. Either an "unqualified pass" or a "qualified pass" on the exam is immediately reported to the Graduate school so that the student is immediately admitted into Ph.D. candidacy.
- The thesis proposal defense maps to “Approval of Prospectus”. A successful thesis proposal defense is immediately reported to the Graduate School as an approved prospectus.
- The thesis defense maps one-to-one. A successful defense is reported as such, and the student is allowed to file his dissertation with the Graduate School (and as a departmental technical report).

The target duration for the Ph.D. process is five years.

Importance of Self-motivation and Initiative

A Ph.D. student is expected to be strongly self-motivated. Unlike undergraduate school or a masters program, Ph.D. level study involves long periods where the primary driver is the student himself. The search for the thesis topic is the most critical of these periods.

A Ph.D. student is also expected to increasingly take the initiative in research as he or she progresses in the program. By the thesis proposal, and ideally well before, the student should feel comfortable suggesting research directions, disagreeing with literature, and taking on side projects.

Lab Expectations

Systems research necessarily involves computers and networks, often many of them. This research infrastructure does not manage or configure itself, nor does the systems support group support all aspects of research computing. Systems students are expected to help in configuring, updating, and maintaining the infrastructure for the Systems group’s overall benefit.

Understanding Funding

Students in the Northwestern EECS Department are funded during the academic year through university fellowships, external fellowships, teaching assistantships (TA), and research assistantships (RA). ***Funding depends on adequate progress toward the Ph.D. and available funding sources. It is not guaranteed.***

University fellowships, such as Murphy and Cabell, typically apply only to first year students. These funds are generally provided, in a department-level competition, on the basis of the perceived quality of the incoming students and the policies of the Graduate School. After the thesis proposal, the Dissertation Year Fellowship and other fellowships may apply. External fellowships, such as the NSF, NASA, and DOD Graduate Fellowships and others, are awarded directly to students and provide the maximum flexibility. ***We highly recommend that students take the initiative in seeking***

external funding. It provides maximum flexibility to the student and will also be rewarded by the systems group.

Teaching assistantships can fund students at any stage in the career. TAships are distributed according to a department-level competition and generally require that the student teach. The time involved in TAing a course should not exceed 20 hours a week on average.

Note on teaching requirements and time: Independent of funding, all EECS Ph.D. students are expected to be involved in teaching to some extent. In addition to TAing a course, a student can also be a Teaching Trainee (TT). A student must be a TA for 3 quarters, a TA for 2 quarters and a TT for 1 quarter, or a TT for two quarters. Teaching can become a time sink, but it should not be. ***If you find that you are spending more than 15 hours per week on average TAing or TTing a course, immediately inform your advisor so that he or she can help you fix the situation.***

Research assistantships are funding that is provided as part of a research grant, generally your advisor's grant, and generally a grant from the federal government. If you are funded from an RAship, the expectation is that ***you will do, in part, research and development related to the grant,*** as determined by your advisor. There is generally a very workable situation as you hopefully share at least some of your advisor's interests and those interests are partially reflected in the grant and its work. Many advisors are extremely happy when students take the initiative in suggesting work to be done while funded on an RAship.

New faculty members generally have some degree of student support as a part of their startup packages. This takes the form of some combination of RAships, Fellowships, and TAships.

There are no fellowships (other than perhaps external fellowships) or TAships during the summer months. ***Summer funding derives almost entirely from RAship funding and is not guaranteed. The expectation is that students funded during the summer will work full time on the research of the underlying grant.***

Students are encouraged to seek out summer funding of their own in the form of internships at quality research laboratories. Students who are interested in doing a summer internship must take the initiative in finding appropriate opportunities. Generally, this must be done in January.

Acquiring Breadth In Computer Science Beyond Systems

Good systems researchers understand the big picture of computer science and related fields such as electrical and computer engineering. You should not embarrass yourself or Northwestern by lack of this basic knowledge. Before taking qualifiers, you should have taken at least one course in each of the following areas. We list here courses from the EECS department. We also strongly suggest that you examine the department's undergraduate computer science curriculum document for a deeper explanation of what

we mean by breadth and depth in computer science. With the consent of your advisor, you may substitute other courses, including 495s and 499s.

- Theory: EECS 310, 328, 336, 356, 357, 457, 459, 495(Current Topics), 495 (Bioinformatics), 495 (e-Commerce); MATH 374. We strongly recommend that students become familiar with algorithms at least to the level of EECS 336.
- Artificial Intelligence: EECS 325, 337, 344, 348, 349, 360, 495 (AI for interactive entertainment,) 495 (Knowledge Representation). We strongly recommend that students become familiar with core AI and machine learning topics as described in 348 and 349.
- Interfaces: EECS 330, 332, 351, 352, 370, 395 (Intermediate computer graphics), 395 (Advanced computer graphics), 495 (Computer animation), 495 (Graphics and perception), 495 (Image-based modeling and rendering), 495 (Human-centered product design)

A student may already have satisfactory background in these areas, either through general knowledge or having taken similar courses at other universities. If the student feels he has satisfied any of these areas, he is encouraged to approach the relevant course coordinator for an assessment, or his advisor if the coordinator is unable to provide an assessment.

Acquiring Breadth in Computer Systems

A systems researcher in some specific area should be familiar with work in other areas of computer systems. The expectation for students is that they have deeper knowledge of systems in general than of computer science as a whole. Each area is listed with appropriate corresponding Northwestern introductory and advanced courses. All EECS systems courses have online syllabi and the advanced courses have online reading lists. You should familiarize yourself with their contents. You need not have taken these specific courses, but you should be familiar with their concepts and content.

We expect that you will take at least six courses in the following areas. You must take at least one course in each of Operating Systems, Networking, and Compilers, unless, for some reason, appropriate courses are not offered. You may already have satisfactory background in these areas, either through general knowledge or having taken similar courses at other universities. If you feel you have satisfied any of these areas, you are encouraged to approach the relevant course coordinator for an assessment, or your advisor if the coordinator is unable to provide an assessment.

Architecture: EECS 361 (452, 453)
Operating Systems: EECS 343 (441, 443)
Distributed Systems: EECS 345 ()
Parallel Systems: EECS 358 ()
Real-time Systems: EECS 397 ()
Compilers: EECS 322 ()
Languages: (no current courses)

Networking: EECS 340 (440)
Performance Analysis: (EECS 410, 442, 486)
Databases: EECS 339 (464)
Security: EECS 350 (450)

Note that there are typically also several additional 395 and 495 courses that may be appropriate for systems depth. Additionally, 499 courses may be used with advisor approval.

Depth in Systems

How to acquire depth in your area will be determined by your advisor. Generally, it takes the form of taking additional graduate level courses and doing guided research and reading. By the end of your second year, we expect that you will have made research contributions.

Qualifiers

The purpose of the systems qualifying exam is to determine whether you have the essential prerequisites of being a doctoral level researcher, namely:

- Have you acquired a breadth of knowledge in computer science and computer systems
- Do you have a depth of knowledge in your research area?
- Can you do research?
- Can you present your research well, both in written form and orally?
- Can you defend your research?
- Can you think and discuss research extemporaneously?
  In other words, can you think on your feet?

If you do not meet these prerequisites, you will not pass the exam. In some cases, such as if you fail due to insufficient breadth or depth, a student may be able to retake the exam. The exam can be retaken only once.

What you should expect

You should ask your advisor if you are prepared to take the systems qualifying exam. If he or she agrees, you should form a committee consisting of your advisor and at least two other systems faculty members. Non-systems faculty are also appropriate in some situations: you should ask your advisor. It is your responsibility to schedule the exam and reserve a conference room for it. Exams have no set length, but past exams have taken from 2 hours to 6 hours. Exams are private: only your committee and you are in the room.

The exam will begin with your presentation of a significant piece of research that you have done. *One week before the exam, you must supply the committee with a paper about the work.* A conference or workshop talk/paper is ideal. The committee will ask you tough questions about the content of the presentation and the work. The purpose of this part of the exam is to determine whether you are capable of doing research, presenting it, and defending it well.

In the next stage of the exam, each of your committee members will have the opportunity to ask you questions. Any technical question related to computer science is fair, however the focus will be on systems. Many faculty members prefer to start with a question designed to test your breadth or depth of knowledge in computer science. The committee may follow up on such questions, probing to find out what you know and what you don't know. The committee is particularly interested in how you respond to questions in areas you don't know or that you don't know the answer to. This is a common situation in doing research and the committee wants to know how you respond to it. It is appropriate and encouraged to ask questions of the committee. The committee also wants to see how you respond in an intellectual dialog.

After the exam, the committee will deliberate and write you a formal letter. Four outcomes are possible:

- Pass. You did great.
- Conditional Pass. You did OK. The letter will explain what you need to do to improve and the process by which you and your advisor will make it happen.
- Fail With Possibility Of Retake. You failed, but the committee thinks there is hope for you. The letter will outline what you need to do before you retake the exam.
- Fail Without Possibility of Retake. You failed and the committee does not believe you will ever pass.

All members of the committee will receive a copy of the letter.

If the outcome is Pass or Conditional Pass, we will immediately tell the graduate school that you should be "admitted into candidacy".

What you should know

Breadth of Knowledge in Computer Science. This is described above.

Breadth of Knowledge in Computer Systems. This is described above.

Depth of Knowledge in Systems. This is described above.

Programming. Good systems researchers build systems, they don't just talk about or simulate them. You must know at least one low-level systems programming language such as C or C++. You must know at least one high-level application programming language such as Java, Perl, Python, Scheme, Lisp, ML, or Matlab. If you haven't written a 1000+ line program in the language, you don't know it. If you haven't programmed on a multi-person project, you haven't programmed. You should look at the web sites of the various labs that comprise the systems group to get a sense of the level of programming you should be up to.

Thesis Process

The point of the thesis process is to demonstrate that you can independently come up with a significant new research question, do the research necessary to answer it, write

compellingly about the question, your research, and the answer, and defend it all. Successfully completing the thesis process earns you the Ph.D. and hence establishes you as a person who can successfully conduct independent research.

The thesis process generally takes from one to two years to complete.

Committee

The thesis is judged by a committee that is chosen by the student in consultation with the student's advisor. The committee commits to reading and commenting on the thesis proposal, attending the thesis proposal defense, providing guidance and advice as the thesis work progresses, reading and commenting on the dissertation, and attending the thesis defense.

The committee must consist of at least three faculty members in the EECS Department that are also faculty in the Graduate School and at least one external committee member. The committee must include the student's advisor, who is generally the chair of the committee. In most cases, the faculty member should be drawn from the systems group, although exceptions can be made. The external committee member should be from outside Northwestern and should hold a Ph.D. Exceptions can be made in consultation with the student's advisor, but a member external to the EECS Department is required.

It is the responsibility of the student to form the committee and to schedule it for the proposal and dissertation defenses.

Proposal

The thesis proposal is a document written by the student that describes the proposed thesis. It must contain:

- Thesis statement. What is the specific research problem being addressed and what is the proposed solution?
- Related work. What have other people done in this area and why is the proposed solution new?
- Prior work. What work has the student done already that suggests that he is capable of addressing the problem?
- Work plan. What the student proposes to do. Of course, research often takes one in unplanned directions. The point of the work plan (and schedule) is to describe what path is currently expected.
- Expected contributions. What artifacts and results are expected?
- Schedule. When will the major elements of the work plan be completed? Notice that writing the dissertation is an important task.

A thesis proposal is generally 10-15 pages long and prepared in consultation with the advisor

The proposal must be given to the members of the committee and posted in written form in a public place in the department at least one week before the proposal defense. It is not necessary to make the proposal available online.

Proposal Defense

The proposal defense is a open, advertised, public talk, given in front of the committee and any members of the EECS department who care to attend. The open segment of the proposal defense is followed by a closed segment with only the committee and the student.

The student must schedule the defense, making sure all his committee members are there physically or via phone conference. The student must assure that the proposal defense is advertised to the EECS department at least one week before it occurs. It will specifically be posted as a thesis proposal talk.

The talk is a summary of the thesis proposal and a defense of its ideas. It's the final sanity check before the thesis work begins and is very important.

Generally, a proposal talk lasts about 50 minutes, although there is no set time. Only clarification questions are permitted during the talk. After the talk, each member of the committee, in an order determined by the chair, will ask in-depth questions. Once the committee is finished with public questions, further questions will be solicited from the audience.

After public questions have been exhausted, the audience will leave and the committee may ask further private questions, or raise other private concerns.

The student will then leave the room and the committee will determine whether the student as passed or failed the proposal defense. The student will be informed whether he has passed or failed on the day of the proposal defense. In either case, the chair of the committee will write a formal letter to the student describing the results and what additional work, if any, is to be done. All members of the committee will be given a copy of the letter.

If the student passes the thesis proposal defense, we will immediately inform the Graduate School that the student's "thesis prospectus has been approved".

All But Dissertation

After a successful proposal, the student will carry out the work described in the proposal, modifying his research plan in consultation with the committee, and, most importantly, his advisor.

Dissertation

A dissertation is a book describing the work carried out during the thesis process and its questions and results. It must be well written and stand on its own.

The dissertation document must be complete, in draft form, before the dissertation defense can take place. It must be provided to the members of the committee at least one week before the defense is to take place. Generally, the student will have his advisor read and comment on the draft well before then.

A summary of the dissertation (generally 10-15 pages) must be posted in a public place in the department at least one week before the defense is to take place.

Dissertation Defense

The procedures for the dissertation defense are similar to those of the proposal defense. The defense is an open, advertised, public talk, given in front of the committee and any members of the EECS department who care to attend. The open segment of the defense is followed by a closed segment with only the committee and the student.

The student must schedule the defense, making sure all his committee members are there physically or via phone conference. The student must assure that the defense is advertised to the EECS department at least one week before it occurs. It will specifically be posted as a thesis defense talk.

The talk is a summary of the thesis work and a defense of its ideas and results.

Generally, a defense talk lasts about 50 minutes, although there is no set time. Only clarification questions are permitted during the talk. After the talk, each member of the committee, in an order determined by the chair, will ask in-depth questions. Once the committee is finished with public questions, further questions will be solicited from the audience.

After public questions have been exhausted, the audience will leave and the committee may ask further private questions, or raise other private concerns.

The student will then leave the room and the committee will determine whether the student as passed or failed the dissertation defense. In either case, the chair of the committee will write a formal letter to the student describing the results and what additional work, if any, is to be done. All members of the committee will be given a copy of the letter.

If the student passes the thesis defense, we will immediately report this to the Graduate School. At this point, the student needs only to deliver the final version of his dissertation in order to graduate.

Wrap up

After a successful defense, the committee will, within 2 weeks, send comments on the dissertation draft to the student. The student will then complete any additional work and make the necessary changes to his dissertation. The student must deliver his final dissertation in two ways. First, he must turn it in to the library. Second, he must publish it as a EECS department technical report.

The purpose of publishing the dissertation as a technical report is to make it widely available to the public.