

Advice for Beginning Your PhD Career

by Idris Hsi

Contributions from: Paul Camp, Zack Kurmasz, Vernard Martin, Jochen Rick, Lex Spoon

Last Edited: January 18, 2004

There are a number of things that I wished I'd known to do in my very first year as a graduate student. While I'm not out yet, I thought I would write up a short list of things that PhD students should start doing when they begin in the program. Some of this is covered in Robert Peters' "Getting What You Came For" which was recommended reading when I took CS 7001. You'll probably find similar graduate student pages and graduate advice pages that have lots of useful advice for how to get out here:

<http://www.cc.gatech.edu/student.services/phd/phd-advice/>

What I've listed in this document are activities that you should be starting from Day 1 of your PhD career to save yourself lots of work later on. These are activities that no one will require of you explicitly like improving your technical writing skills, public speaking ability, or general knowledge in an area. In fact, you can probably graduate with a PhD without having done any of them but you will have done it the hard way and will be making things more difficult for yourself in the future. The beginning of the list contains 'mechanical' activities - things that fall under the aegis of 'information maintenance'. The end of the list contains activities that are more personal (how you do it depends on what you're like as an individual) and abstract but important nonetheless.

I've also written this document with an emphasis for those PhD students interested in doing research after graduation but there are some useful tips for those of you walking the hard path to eventually find a teaching position at the college/university level. Whether your objective is research or teaching, you will still need to jump the dissertation hurdle and that means you have to define a thesis topic, do some research that contributes to the fabric of knowledge, and produce a dissertation.

1 Keep a Research Diary:

I keep a file on my main work computer where I write up meetings with my advisor, good ideas that I've had about my research, bad ideas that I've had about my research, problems that I've encountered, talks that I've had with other professors or colleagues, and even emails about research that I copy into the diary in a different font. I started doing this recently when I realized that over the years I'd filled a box with notes that I had no way of searching intelligently. I wish I had started earlier. Every now and then a wacky idea that I ignored at the time I wrote it turns out to be informative and useful. This diary is also a good place to make notes of your anxieties about the program or your progress just to clarify the problems that you think you're having. Hopefully after you've graduated, you can look back on your recorded angst and laugh.

In your early years, you should be generating ideas about potential thesis topics. Even if you're brainstorming, you should write them down. Most of them will not go anywhere or will not be relevant or accomplishable in your time frame. Eventually one will lead you down the right track.

Some examples:

- March 2, 2003
An idea about bipartite graphs: number of portals connected to concept is another measure of importance - reflects activation cost requirements on part of designer to concept.
Yahoo Messenger is freaking insane!!! Maybe I need to pick a different application - like Microsoft Word 2.0.
- April 14, 2003
Wrote up 'Surveyor_Description' for Melody describing Surveyor's basic functionality and relationship to Architect.
Basically getting back in gear after being annoyingly sick with a cold for the last 3 days.

* It's also recommended that you keep a *Daily* entry. This will get you into the habit of making regular entries and to also remind you that you need to be making regular progress towards your dissertation. What do you enter on a daily basis? Record any little bit of progress that you made. If you're coding, write down the modules you coded or debugged. If you did some reading, note the papers and what you learned. If you did nothing, note that you did nothing (but keep those kind of entries to a minimum). When the PhD review comes around and you're feeling evil when they ask you what you did in the last year, send them all the entries you made for that year. [Tip from Vernard Martin]

2 Maintain an Electronic Bibliography:

Whether you use BibTeX or EndNote or design your own database, you should keep an electronic record of the papers and books that you read whether it's in your area or not. This may not seem like an issue for you now but wait til you get into the 300 paper range and are trying to remember the name of the authors who wrote about applying design patterns to multithreaded real-time applications in 1994 so you can cite them in a paper that you're writing. I know that EndNote and BibTeX have add-on tools that work with word processors to do things like automatically generate bibliographies and citations after inserting fields in your document. Very useful when you're writing your proposal or dissertation or research paper.

I also use EndNote as my database for storing my notes of a research paper's key points. It saves time rereading the actual paper and it generates more key terms for searching on this database.

As a corollary to this, go to an Office Depot type store and spend the \$8 or so to get a big box of file folders to store your papers. The best organization scheme that I've seen seems to be alphabetical by first author. Another alternative is to maintain libraries of .pdf files on your computer as a way of storing research.

3 Know Your Search Engines:

There are three very useful search engines that you should be aware of in doing literature searches. They have great collective coverage over the computer science field and often allow you access to the .pdf or .ps files that you need to print them out for yourself.

Georgia Tech has a campus subscription to the IEEE Digital Library that only requires your standard GTEL login and password. It's located here:

<http://gtel.gatech.edu>

Find the databases tab and find IEEE then log in.

The other major repository of literature is ACM's Digital Library. This requires a membership in the ACM but it's something you should think about anyhow if you intend a long career in computer science. It's located here but you need an ID and a password to get to the cool areas:

<http://portal.acm.org/>

The last is a research site maintained by NEC and keeps a very nice database of publications, sometimes with links to the actual files. It also will show you a graph showing the time distribution of your searches.

<http://citeseer.nj.nec.com/cs/>

This page has a list of conversion tools:

<http://www.ecst.csuchico.edu/~jacobsd/bib/tools/converters.html>

4 Develop an Archiving System for Useful Info

Emails, Papers, Revisions of Papers, Classwork, Lecture Notes, Meeting Notes, and so on will eventually swamp you in a Sargasso Sea of Information where nothing useful escapes. It's very useful to design a good filing and retrieval system very early on, both virtually and physically. The more work you do now, the more work you'll save yourself later on when you need to find something.

For example, Business Cards - If you go to enough conferences, job fairs, and corporate socials, you will acquire a collection of business cards that you need to organize. You need to be able to recall: who gave it to you, when, why, what you talked about. One trick is to scan and OCR them, but they will usually have something handwritten on them. Perhaps you just need to tape them chronologically in a journal with a note on each one. [tip from Paul Camp]

Lastly, you absolutely must have a system for **BACKING UP YOUR IMPORTANT RESEARCH REGULARLY!!** I recommend periodic CD burning as well as CNS backups.

5 Maintain a Professional Home Page

Your home page is the first thing that researchers and other colleagues will see. It should be clearly presented and the front page should be professional (not necessarily professionally designed but not sloppy). Your web content should have the following information

- A picture of yourself
- Personal and contact info.
- Your research interests - what I'm working on and why.
- Current and previous project affiliations.
- Your academic history - (previous degrees)
- Electronic versions of papers that you've published.
- A current Curriculum Vitae - The Academic Resume

Here are two very good examples of what else you should have and how it could look:

<http://www.csc.ncsu.edu/faculty/anton/>
<http://www.cc.gatech.edu/fac/Gregory.Abowd/>

A spare example but extremely functional is here:

<http://www.cc.gatech.edu/~spencer/>

When you go to conferences or start communicating with people either in or outside your college these days, the first thing that they will do if they are interested in you or your work is go to your home page to find out more about your work or about your history. I've seen faculty candidates get derided in review meetings for having an inferior web page. In the 21st century, your home page is your virtual "public face" - the first thing people will see when they're checking you out. Don't let it be "Under Construction". Keep it maintained with your latest activities. You should update your home page no less than once per semester.

6 Maintain an Updated CV:

If you volunteer for a committee, a conference, take an internship, publish, or do any of a number of things that should go on a CV, update it immediately. A CV looks like a resume but includes:

- Publications Record
- Research Experience
- Service Record: Community and Professional Activities
- Teaching Experience
- Membership in Professional Societies
- And so on.

Along those lines, you should look at other CVs and see what they're filled with - it's not just about your publications although they are the most important component of your academic resume. Which leads me to my next point.

7 Get Involved Early On

Volunteering for conferences, serving on the various college committees, helping at college activities, attending faculty candidate talks, giving research talks, giving and attending research demos, going to the occasional corporate mixer, and so on are all things that you should be doing as early as possible. Being a TA to a professor in an area of interest (and doing a good job) is also a very good way to meet faculty. This public involvement serves two purposes. One, you're getting exposure to the facets of academic life that are not related to research but can be very important in the intangible ways. Two, you're making contacts (networking) and hopefully making a good impression on the people that you meet. It's important to develop a positive professional image to your colleagues and especially the faculty. This is a lesson that holds true wherever you go. It's also good CV filler. Faculties looking to hire candidates also look at how useful that person was to their community.

8 Memberships in Professional Societies

In computer science, everyone who's anyone belongs to either the Association of Computing Machinery, IEEE Computer Society, or both. Those who are specializing in an area also belong to the special interest groups associated with those organizations (e.g. SIGSOFT, SIGCHI, SIGMOD). Student membership is relatively inexpensive and is a nice bit of CV filler. ACM membership also gives you automatic access to the digital library. See if your advisor will help you out.

9 Begin Learning The Composition of your Research Community

Who are the famous researchers in your area? Who are or were the most influential people in the field? Which conferences are first tier? Where are the major centers of work in your area in the world? Which journals should you be reading regularly? What are the current research trends or sexy topics that are getting funding? What are the niches of specialization and who's in them? You must know this like the back of your hand eventually. Some of this you will get by osmosis from your advisor, your colleagues, or persistent visits to conferences and meeting these people in person. It's good to begin taking notes early on these things to help you with your later career and job search.

As a corollary to that, you should also know something about the work being done at your own university. I had the uncomfortable experience once of being at a conference and having someone ask me about the work of a colleague that was strongly related to my own that I didn't know about. Doh!!

Slightly more difficult to learn are the various, often bizarre, political nuances and biases associated with various conferences, publications, and organizations. If your advisor is anything like mine, he or she will have a particular feel for these things and will often throw away statements like "They're not interested in this style of work. We should try this conference instead." Make sure to follow up these statements with clarification questions. Don't take things for granted. Academia is often extremely political. It's good to know the history and climate of your field (and your department).

10 Develop Research Meta-awareness and Skills

One of the most useful skills you can have leaving graduate school is the ability to think objectively and critically about a problem or its proposed solution. You need to learn what is good research from problem statement to validation, how to recognize good research whether it's in your area or not, and how to recognize when you're looking at bad research (either because of irrelevance or implementation). The worst thing I did when I started as a graduate student is to treat the process like an undergraduate student - waiting for someone to tell me what to do and how to do it. Learning how to reflect actively on what was going on around me was one of the most painful and rewarding things that I've learned over the years. Simply being able to recognize intuitively the difference between a solved research problem (even outside my area) and unsolved one is tremendous. Being able to recognize the difference between a tractable and intractable unsolved research problem will help steer you away from the Rocks of Despair when you're looking for your thesis topic. Your advisor will guide you through this a little but you need to learn this for yourself if you want a career as anything other than someone's research assistant.

To develop this skill you need to maintain an extra level of awareness when you're pursuing any of your academic activities - reading, writing, conversing with colleagues, attending conferences, and so on. This 'extra sense' should be asking questions in the back of your mind like "Is this work important?", "Does it tackle the problem in the right way?", "Do I believe the validation?", "What are the implications of this work?", and "What can I learn from this research (or person)?".

I believe the difference between a competent researcher and an innovative one, either at the student or faculty level, is how good they are at recognizing good problems. By the time you leave, you should have the following skills:

- Identifying an important problem
- Identifying an unknown problem
- Recognizing good validation
- Technical Writing / Communication Skills

11 Develop Objectivity

I came across two interesting quotes in a NY Times editorial which address this point.

"One is reminded of Harry Lime's famous quip in the movie 'The Third Man' that 30 years of noisy, violent churning under the Borgias in Italy produced Michelangelo, Leonardo da Vinci and the Renaissance, while 500 years of peace, quiet and harmony in Switzerland produced the cuckoo clock.

'A monolithic framework does not create a critical mind,' remarked the religious philosopher David Hartman. 'Where there is only one self-evident truth, nothing ever gets challenged and no sparks of creativity ever get generated. The strength of America has always been its ability to challenge its own truths by presenting alternative possibilities. That forces you to justify your own ideas, and that competition of ideas is what creates excellence.'"

(Friedman, Thomas "Cuckoo in Carolina", The New York Times, August 28, 2002)

While both these quotes address the disadvantages of stability and monochromatic thinking at the societal level, they indirectly speak to the disadvantages of maintaining a one-sided view of research, research techniques, or a research area. Good research requires the rigor of scientific objectivity.

Scientific objectivity is not simply a good idea to remember every now and then; it is a skill that needs to be practiced from Day 1 of your research career. How do you do this?

a) Get your ego out of the way. Your biases should not interfere with the process of doing research. This is a tricky point. In science or engineering, the facts, experiments, proofs, or logical arguments should be the final arbiter. Defending or challenging research should always be directed at the method, the results, the analysis, or the interpretation. Practice examining your work (writing, research, coding, whatever) as objectively as possible and by getting vigorous feedback from your peers and advisors.

b) Learn to examine all sides equally. Whenever I hear something that I disagree with, I first ask myself, "What if they're right and I'm wrong?" Then I try and construct an argument from their point of view to understand the flow of logic. After I think I've understood it, I test it for possible flaws or holes. Then I do the same with my argument because it's important to be willing to challenge your own ideas with the same rigor as you would challenge someone else's. I practice this skill by reading opinions in reputable newspapers or position papers.

c) Read outside your area. You should actively read books and articles in other research areas outside of your concentration. Does this help you develop objectivity? It might not. But you gradually learn that other people are tackling similar problems in different areas, that there are different techniques for argumentation and research that you might adapt to your own work, and that many people have come before you facing the same or harder challenges.

d) Learn how to discuss other people's research in addition to your own. Being a good colleague requires you to be both a good collaborator but also a good listener. This should apply to work that is outside your immediate area of knowledge. You benefit in two ways. First, by listening to research in other areas, you've increased your knowledge and awareness of your field (or another field). Secondly, if they reciprocate, you will get a different perspective on your research or a problem that you're solving that you may not have been able to reach in isolation. The worst kind of colleague is one that has no patience to listen to any other ideas except their own. The Key Point here is to Listen Objectively! Don't prejudge or correct until you understand what the other person is saying.

Whenever I think about scientific objectivity, I think of Pasteur and Koch. Pasteur was investigating the problem of spoiled wine and beer. He had a hypothesis (caused by microorganisms) but the scientific community didn't believe him. He systematically eliminated every possible alternate explanation (often so he could beat his critics over the head with his results) with numerous experiments until one inescapable conclusion was left. Koch did the same thing - the excerpt of which you can find here:

http://www.cc.gatech.edu/~idris/Stories/The_Microbe_Hunters.htm

12 Document Useful Learning Experiences

I had the painful recent experience of having to relearn programming and Java to build some analysis tools after a long hiatus of doing theoretical research. This is not the first time I've had to relearn something but this time the learning curve finally hit. I now have a document of how to do things like install Java, the stupid system peculiarities that I had to overcome, what I did to solve bugs as they came up, and so on. I'm also developing a web page of useful code fragments - documentation headers, loop structures, try/catch things - that I use over and over again. Eventually, if I ever teach a programming course (or have to relearn Java again), I'll have teaching materials ready. This is also a good thing to do with classes that you TA to keep a memory of what worked and what didn't. An example is below:

http://www.cc.gatech.edu/classes/cs4301_99_spring/TANotes.htm

This document itself is a form of writing up a process that may be useful to me later in life as an advisor or for passing on to other graduate students. What you want to document is anything that took a tremendous amount of searching or effort to learn. It could be as small as how to install print drivers on your Windows machine to interface with the Colleges or as large as how to set up research that involves human subjects. Documentation is very useful for saving you (or someone else) effort later on.

13 Miscellaneous Things

- Design and Practice your elevator speech (1 minute shot at fame and fortune)
- Get Business Cards from the College before going to conferences
- Find out who the senior students are both in and outside your area!
- Find and maintain a stress release mechanism (hobby, social outlet, exercise).

14 Work Daily Towards the Big Picture

I can't stress this last one enough. The most common characteristic across all graduated Ph.D students is that they wanted it badly enough to put up with a lot of suffering to get the degree. You should spend every day making incremental progress towards graduating.

- Remember why you're here!
- Find a topic that you will be excited about when the going gets tough.