How to Be a Good Graduate Student

Frank Vahid, Ph.D.
Department of Computer Science and Engineering
University of California, Riverside

ISBN 978-1-4303-2296-2

Contents

Chapter 1 – Introduction

Graduate school is the best-kept secret in America.

Many students, exhausted from their 4-6 years of undergraduate courses, can't imagine why anyone would subject themselves to 2-6 more years of that torture to get a master's or Ph.D. degree, when a bachelor's degree is sufficient to get good jobs. However, contrary to popular misconception, graduate school[1] is not just "more of the same."

Good Reasons for Going to Graduate School

Academically, the graduate school experience can be very different, and more enjoyable, than the undergraduate years. First, you've already learned how to learn, which is half of the undergraduate experience and the source of much of the frustration of that experience.

Second, and more significantly, professors run graduate classes differently than they run undergraduate classes, with less orientation around tests, homeworks, quizzes, or policies designed to keep problem students in line. Professors run graduate classes differently because they can assume all students are hard working A or B students capable of doing well in college courses, which is mostly true because those are the students that graduate schools admit into their programs. Professors can therefore spend more time just sharing the material with students. Professors tend to flock to graduate courses because they are so much more enjoyable to teach, forcing most departments to have policies that actually *restrict* professors from teaching only graduate courses (with no

[1] Attending "graduate school" in this book generally refers to obtaining a Ph.D., with more emphasis on engineering and science subjects.

such restrictions on undergraduate courses). The professor's desire to teach a graduate course directly translates into a more pleasant experience for the students. Professors tend to treat students more the way good students want to be treated.

A third reason that graduate is more enjoyable is that, when obtaining a Ph.D., one usually finish all coursework in the first two years; the remaining years are just research, which is really more like a job than it is coursework.

In non-academic aspects, the experience can be superb, and can have life-long positive impacts. While you work hard, your time is generally more flexible than at a full-time job, and that flexibility is often more important than total hours. Graduate students often use that flexibility to enjoy activities together, like volleyball, soccer, beach trips, hiking outings, dinners out, jazz gatherings, game nights, and much more. Such social activities seem to come easier to graduate students than to undergraduates because they tend to have more in common with one another than they did with fellow undergraduates, and also because the environment is set up to give students more interactions with one another through its smaller total student population, further sub-grouping into research labs, interactions during teaching assistant duties, smaller classes, etc. My own undergraduate social experience was reasonable but nothing much to speak of, whereas I have *tremendously* fond memories of graduate social life, where I made many friends that are still an important part of my life. Furthermore, because companies seek out graduate students for summer internships, a graduate student can experience both school and work in a nice mix, satisfying the need to both learn and to be treated like a professional. Research conferences provide an opportunity for travel to new cities and countries, expenses paid. A graduate student might live in graduate student housing, which is not only low cost and enables walking to work, but has the student living among highly-educated low-income neighbors – a combination you'll

never find anywhere else. The student can enjoy wonderful conversations with these people who aren't thinking about mortgages and furniture and luxury cars. Through all this, the student may have tuition covered by the school, and may even be earning a monthly stipend from a teaching assistantship, research assistantship, or fellowship, not only avoiding finishing school in debt (which isn't that bad anyways, as school is an investment), but even finishing with some money in the bank. Compare that to law or medicine. Perhaps most importantly of all, the student networks with smart people, and bonds much more closely with them than at the undergraduate level – colleagues in a research lab may become lifelong friends, a sort of extended family.

And at the end of all this, the student gets an advanced degree, which typically translates to higher pay and a more satisfying job and career. It's no wonder that graduate students, far from rushing to complete graduate school, often linger for an extra year or two, hoping to take in a bit more of the good life before moving on.

Of course, graduate school is not a cakewalk. Standards are high. Independence and self-discipline are a must. Finding a good advisor, and maintaining a good relationship with that advisor, can be a challenge. And new experiences, like writing research papers, taking oral examinations, creating a 100-page thesis, and giving talks at conferences, can be intimidating and even harrowing experiences.

Bad Reasons for NOT Going to Graduate School

Hopefully the above discussion will encourage strong undergraduate students to consider attending graduate school. Choosing to NOT go to graduate school is certainly a reasonable decision in many cases, but there are some potentially bad reasons students give for not going.

Some people want to go the graduate school eventually, but decide to work a few years to gain experience, take a break from school, or make some money, before going back to graduate school. To them, I say,

"Good luck." People are indeed correct that the work experience can help setting life goals and defining desired careers, can catalyze learning by providing better context for the material being learned, can provide a needed break, and can get the student some needed money. Some of the best students I've seen, undergraduate or graduate, are older students who have worked for a while, set a new life goal, and come back to school determined and motivated to meet that goal. However, students rarely come back to graduate school after they take on a professional good-paying job – I would estimate that of all the students who planned to return to school, only 1 in 10 actually do. Once the student has started making the big bucks, it's hard for hime to wean himself off of that monthly income, especially if he has a house payment and a car payment, and/or has started a family.

Some students plan to get their master's degree part-time while working full-time. I've seen lots of people try this, and some of them indeed successfully achieve the master's degree, but it can be a challenging path. Do you really want to work 40-50 hours a week, and then come home just to spend more hours doing coursework in the evenings and weekends? It's exhausting, and if you are starting a family, you may find yourself missing out on a lot of family life, and maybe even straining your marriage. And do you really want to spend 8 years of your life like this? (Do the math – if you take 1 course per semester as a part-time student instead of 4 courses as a full-time student, you stretch a 2-year degree into 8 years). If you do decide to work full-time after a bachelor's degree with a plan to return to graduate school, have a very clear plan (e.g., "I will work for one year and I will apply to bio-engineering graduate programs next fall"), and avoid creating a lifestyle that requires a large income (e.g., it's probably not a good time to buy a house or expensive car). Remember that summer internships, or even slightly longer internships (e.g., 6 months), can satisfy some of the goals of that year or two of work.

Some people avoid graduate school because they believe there are fewer jobs available to people with advanced degrees. To them, I say that if it's just the number of available jobs that concerns you, then McDonalds is hiring on every street corner. Clearly, people want more than just a job; they want a satisfying career. A bachelor's degree is indeed sufficient for a satisfying career in many cases, but a masters or Ph.D. can lead to even more satisfaction, due to more interesting work and more responsibility. Today's masters and Ph.D. graduates aren't unemployed en masse; in the worst case of them not finding a job in their specialty, they can probably find a job that required a lower degree. Furthermore, today's job descriptions commonly prefer a master's degree to a bachelor's degree, and many modern high-tech and other knowledge-oriented companies prefer Ph.D.s. If you hope someday to teach either part-time or full-time at a community college or university (and I highly encourage this – "college instructor" often ranks in surveys as being one of the best careers in America), then a master's or Ph.D. degree is a must.

Some students would like to attend graduate school, but don't want to delay getting married or starting a family. To them, I'll point out that graduate school can be a great place to start married life or to start having kids. Many schools have graduate, married, and family student housing right near campus. Such housing is low cost, meaning the big income of a job isn't necessary. Such housing eliminates commuting, leading to less cost, less stress, and more time. Neighbors in such housing are low-income highly-educated people – a combination you'll never find anywhere else. Universities often have low-cost day care and/or preschool programs on campus too. The social network in graduate school can be great for young couples – fellow graduate students aren't (yet) caught up with fancy cars, big houses with accompanying big mortgages, investments, etc. Fellow graduate students generally have more time to enjoy each others' company. I got married the summer

before starting graduate school, had my first child 3 years into graduate school, and my second child a few months before graduating. If I had to do it again, I wouldn't change a thing – my wife and I remember those years with tremendous fondness.

Finally, some students are reluctant to apply to graduate programs because those students aren't sure whether they are really committed to or capable of achieving a Ph.D. To these students, I say to be aware that many graduate students feel that way too. You don't have to decide for sure that you want to get a Ph.D., and you don't have to know that you'll succeed. If you believe there's a reasonable chance of a Ph.D. in your future, you can apply to Ph.D. programs with a clean conscience. If after a year or two you decide it's not right for you, you can always get a master's degree along the way and leave. Or you can just leave. If you try for the Ph.D. and end up not succeeding (perhaps by not passing the qualifying exams), so what? There is absolutely no shame in not achieving your original goal, due to changing your mind or to simply not having succeeded. People cannot be expected to know exactly what they want in life, or to know exactly what they are capable of accomplishing, without trying. You learn as you go along.

Graduate School in a Nutshell

Some readers might not know much about the actual mechanics and details of graduate school, so here I'll provide a brief summary.

Graduate Studies

What most people think of when they here the term "college" or "university" are what are known as ***undergraduate studies***, consisting of about four years of full-time coursework that lead to obtaining a bachelor's degree, such as a bachelor of science (B.S.) or a bachelor of art (B.A.) degree. ***Graduate study*** is more college coursework that leads to obtaining a master's degree or a Ph.D. degree. Schools typically use

different numbering schemes for graduate courses (e.g., 200-299) than they use for undergraduate courses (e.g., 1-199), and graduate courses are typically restricted to graduate students only, namely those students who already have a bachelor's degree and are officially enrolled into graduate school.

Common master's degrees include a master of science (M.S.), a master of art (M.A.), or a master of business administration (M.B.A.), and typically require about two years of full-time coursework, plus perhaps a thesis dissertation.

A master's thesis dissertation usually involves studying a well-defined subject deeply, such as studying whether cell phone usage impacts driver response times (the belief that it does forming the "thesis"), and writing up the study and the results (the write-up forming the "dissertation"). A typical master's thesis is about 50-150 pages long.

Ph.D. stands for "Doctor of Philosophy," coming from the Latin "Philosophiae Doctor." A person with a Ph.D. degree is commonly given the title of "Doctor" ("Dr."), which is obviously different from a medical doctor (M.D.). Ph.D. studies typically involve another 3-5 years of graduate school, but that schooling consists of very little coursework (if any), and primarily consists of doing research and developing an in-depth thesis dissertation (more on those subjects momentarily).

Applying to Graduate School

A student applies to graduate schools typically in the late fall (or maybe early Winter) preceding the year he/she wants to start graduate school. Schools evaluate applications once a year, typically in the winter or spring (some schools evaluate applications year round). The applicant must usually provide transcripts, graduate entrance exam scores, a statement of purpose, and letters of reference.

Regarding transcripts, schools typically have a minimum grade point average (GPA) requirement, and pay much attention to GPA in

sorting applicants. At my school, the minimum GPA this year was 3.2. However, that minimum can be misleading, because we rarely admit anyone with a GPA near 3.2. Furthermore, there may be a higher minimum GPA required to be admitted with funding support.

Regarding graduate entrance exams, a common one required by many schools is the GRE exam (Graduate Record Examination). The basic GRE exam consists of a math part and English part. There is also another part, called the subject part, which is required by some schools too. The subject part focuses on a particular field, such as computer science or chemistry. It's a very good idea to prepare for the GRE exam before actually taking it. There are many books, online sources, past exams samples, and even courses (which can be costly). Your university may offer preparation courses too, at much lower cost. If you are unhappy with your scores, you can retake the exam. Universities differ in how they treat retaken exam scores. But take the GRE exams early enough to give yourself a chance to retake them.

Some graduate school surveys (such as those by U.S. News and World Report) publish the average GPA and GRE scores of applicants to programs at universities.

Regarding a statement of purpose, this is usually a page or two description by the student explaining why the student wants to go graduate school, and specifically why to this particular graduate program. These are usually read with an eye for determining the student's particular topic of research interest. A computer science department, for example, might try to determine if the student is interested in studying the theory of computing, or in making innovations in computing systems. For the latter, is the student interested in database systems, network systems, operating systems, or something else? The department strives to admit students whose interests match those of the faculty – there's little point in admitting a student who wants to do research on databases if there isn't a single professor in the department

whose specialty is in databases. Furthermore, the department strives to admit certain numbers of students in each topic so that those numbers match the needs of the department – in a given year, a department might want to admit 5 students in databases and 10 in the theory of computing, for example. A good letter narrows the topic down somewhat, but need not be extremely specific. So a letter might show interest in computer systems in general, with particular interest in databases and networking, but need not define a research thesis, like an interest in studying how databases larger than one gigabyte can be compressed into 100 megabytes or smaller. If a student has relevant work experience, is particularly interested in working with a specific research professor or group in the department, or has overcome tremendous adversity to get to where he/she is (showing motivation), those things should be mentioned in the statement of purpose. The weakest letters are those that only say things like "Education is important to me" and "I hope to develop a strong educational basis for a rewarding future career." Blah, blah, blah, everybody assumes you want those things – those letters are akin to a restaurant owner asking a customer why the customer came to his restaurant, and the customer saying, "Because I was hungry." Say something nice specifically about that restaurant's food or service – you might find you get better treatment.

Regarding letters of reference, these are typically (and preferably) from professors at the student's undergraduate institution, but may also be other people with a working knowledge of the student, such as former bosses. (Letters from family or friends is not what is being sought here). Letters that say something like "Mary took my course and did well" are not very helpful, though they are better than nothing. A stronger letter talks about the student's intelligence, creativity, maturity, self-motivation, ability to work independently or in groups, etc., and gives specific examples. Good letters are glowing. Obviously, a professor needs to have interacted with the student to be able to say something in

depth or glowing about the student. One form of interaction comes from the student asking questions in class, talking with the professor after class, coming to office hours, etc. Another form comes from having done a project with the professor, perhaps in a project course, or perhaps an undergraduate research project, which is sometimes done for pay, for course credit, or on a volunteer basis. Students planning to attend graduate school might wish to consciously plan to get to know a few professors with the goal not only of learning more or gaining good experience, but also for the reason of getting good letters.

A department that runs a graduate program (e.g., a Department of Computer Science or a Department of Chemistry at a university) typically appoints a committee of faculty to evaluate applications and make graduate admissions decision for the department's program, with the help of a graduate secretary. However, any professor in a department may influence admissions, especially because ultimately it is those professors – not the committee – who will be the ones working with and financially supporting those students. Thus, the applicant might wish to contact key professors in a department, either by email, phone, or even with an in-person visit. Regarding email, be aware that research professors may get hundreds of such emails from applicants, and many professors just routinely delete those emails without ever reading them. If you know a professor at your own school particularly well, you might ask him/her to contact colleagues at your target graduate school to let them know you've applied and that you are a good catch.

The graduate admissions committee makes decisions based on your GPA, your GRA scores, the strength of your letters, the reputation of your undergraduate school, recommendations from department faculty, need and balancing of topic areas, and other factors like country diversity, gender diversity, industry experience, domestic and international student balance, and so on. Don't feel too bad if you don't

get admitted to your school of choice; there are many factors beyond your control.

Admitted students are often given an admission "offer" that goes beyond just "you are admitted to the program." The offer may promise some form of financial support. There are three common forms of support. A "teaching assistantship" means the student will be paid to work as a teaching assistant (TA) in courses, usually undergraduate courses. A TA typically works 20 hours per week, and pay at the time of this writing is usually in the mid-$1000s per month, plus tuition and fees being covered. A "research assistantship" means the student will be paid to work as a research assistant (RA) in a professor's lab. An RA may officially be enlisted to work 20 hours a week, but they may do much more research than just the 20 hours for which they are paid – after all, the research is progress towards the degree, not a job. The third form of support is a "fellowship," which is what undergraduates usually call a scholarship. A fellowship is free money, without a formal requirement to teach or do research, though it is usually expected that you will do research. While universities offer fellowships themselves, fellowships are also available from outside sources, like the National Science Foundation or the Ford Foundation (though those are usually applied for by the students themselves, either when applying to graduate school, or during the first year of graduate school). Offers may combine the above three kinds of support, such as offering a TAship plus a fellowship, and promise of an RAship in the first summer and second year. Support offers are usually for a minimum of one year, though they often are for two, three, or even four years. However, a one or two year offer does not mean that you won't get support for the later years, it just means the department isn't obliged to give you that support. Good students often do find continued support after their promised support ends. .

Once a student joins a program, he/she may take 3-4 courses per quarter for the first 1-2 years. Some of those courses will be "normal"

graduate courses, similar to undergraduate courses (but better, as discussed earlier). Some of the courses will instead by "seminar" type courses, which aren't anything at all like undergraduate courses. Seminar courses usually consist of students and a professor reading research papers in particular topic area, with the students presenting the papers to the class, or consist of just attending a weekly talk given by visitors to the department or professors from the department. A total graduate student's course workload is generally much less than 3 or 4 undergraduate courses in the major.

The student may begin doing research in the first year, but significant research doesn't usually start until the second or even third year. By the third year, the student is doing research almost exclusively. Research may consist of the following cycle: reading research papers; discussing new ideas with fellow students, your advisor, and other professors; formally proving the ideas or building tools or running experiments to validate the ideas; writing up results in a paper that is then submitted to a conference or journal; and traveling to a conference to present the paper to conference attendees. That cycle may be done once or twice for a master's degree, but may be carried out 5 to 10 times for a Ph.D. degree over the course of several years, with the ideas all being related to a central thesis topic.

For a Ph.D. student, a department will typically also have some "qualifying" requirements beyond the initial coursework. A student may be required to take one or more written and/or oral examinations (beyond those given as part of courses) that test the student's breadth and/or depth of knowledge of the field, with some of those exams given by a committee of professors consisting of the student's advisor and 2 to 4 other professors typically in related research fields. In some schools, such qualifying exams may be straightforward, but in other schools, they can be quite grueling, and students may fail the exam as often as students pass it. A student typically is allowed a second try, but ultimately not

passing the qualifying exams and thus being forced to leave a Ph.D. program is not entirely uncommon at some schools. This is not the end of the world; many students transfer to another school and get a Ph.D. there, or go on to industry and have a fine career. In addition to exams, a student may have to write a survey paper (a paper that summarizes a research field), and may have to write an in-depth research paper of publishable quality. All of the above items are usually done in the second or third year of graduate study. After completing all coursework and additional qualifying requirements, the student is said to have "advanced to candidacy" – the student is now an official candidate for a Ph.D.

After advancing to candidacy (though commonly even before that time), the student develops a research topic with his/her advisor, and may propose the research topic to a committee for feedback and approval. Then, the student spends a couple years conducting research on the topic and writing a dissertation. Actually, the research topic may only be vaguely known early on, with the topic becoming more defined as the years proceed. The student eventually "defends" his/her thesis in front of a committee, consisting of his/her advisor and 2 to 4 professors, usually the same professors as on the previous committee. The defense is typically a prepared talk on the research topic, summarizing the contributions of the research to the state of knowledge in the field. The committee asks questions throughout and after the talk. If the committee does not approve of the dissertation or the student's defense of the topic, they typically tell the student what must be changed or done, and the student refines the dissertation and/or does additional work and defends the thesis again in a few weeks or months. It is very rare that a student is flatly turned down this late in his/her graduate studies; in fact, I've never heard of it happening anywhere (but it could have happened somewhere). Once the committee approves of the dissertation and defense, the candidate has completed his/her Ph.D. – the student is then known as "Dr."

Why This Book, and How to Use It

My main goal in writing this book is in small part to encourage people to attend graduate school, and in large part to help make the graduate experience a little bit smoother for those who have already decided to attend. It wasn't that long ago that I was a graduate student myself (from 1988 to 1994), so I still remember the student perspective. Yet I've also been on the professor side of the fence for sufficient time now (1994 to present) to see the challenges and pitfalls facing many students – and ways to deal with such things. Thus, the time felt right to write a short book on the subject. Let me make clear that the information here is likely to be most relevant to students pursuing a Ph.D. in an engineering (including computer science) or science subject at a Ph.D. granting U.S. research university. I suspect that much of the information is useful for other majors and places too, and even for students pursuing a master's degree only.

I have written this book as if I were giving advice to my own students. I have not attempted to be comprehensive, to provide all viewpoints, or to be politically correct – I find that many advice guides that try to do those things are too diluted to be really useful. This book is one person's opinion. Seek out other opinions too.

Probably the best way for a reader to use this book is to build a "to-do" list as one reads the book. The reader may find that certain suggestions in the book are particularly relevant to his/her own lifestyle or goals. For those suggestions, the reader might write down a note to himself/herself, indicating a goal or action item for the reader, e.g., "Read each morning for 15 minutes," "Develop a more regular sleep schedule," or "Try to meet more people." The result of reading the book might therefore be a list of 10-30 items that the reader plans to focus on in their graduate studies. The last pages of this book were intentionally left blank for the purpose of letting the reader write notes there.

Chapter 2 – Choosing a Graduate School

Of the several thousand colleges and universities in the U.S., only a couple hundred of them are research universities where students can conduct graduate level research and obtain a Ph.D. Yet a couple hundred is still a large number, which makes the decision of which school to apply to and to attend a difficult one.

Rule #1 – Don't "Stress Out" So Much

Throughout graduate school, a student is faced with many decisions – in which area should she focus, should she pursue a masters or Ph.D., with which advisor should she work, etc. However, the first decision, after deciding to attend graduate school, is to decide which school to attend. Like all the decisions facing the student, this decision can be the source of great stress. Should one go to the school with the best ranking? Should one go to the school that makes the best offer? How much attention should one pay to the school's location, culture, and climate? How can one pick a school if one is not really even sure what research areas are of interest? These and similar questions are tough ones to answer. What I like to tell students, though, is this. These are the *good* types of decisions to have to make in life. You are deciding between several good things. I'd much rather be deciding between going to the better ranking school versus the one that gives me the better offer, than deciding which type of bankruptcy to declare, whether to stick it out in a bad marriage or to instead get a divorce, or whether to live with chronic pain or to instead

have a life-risking surgery. So relax, weigh out the pros and cons, make your decision, and move on. Know that there really is no "best" decision and there especially is no "right" decision – each school represents a tradeoff compared to other schools. Plus, there are so many random factors involved after you make your decision – your advisor might suddenly die, a hurricane might hit your school (or earthquake, or tornado, or fire, or flood – you can't escape!), you might win a great scholarship, you might meet your future spouse – that your success is not really entirely dependent on your decision. Those other factors, which can affect your future just as much as your own decisions, are largely out of your control. You influence things a bit, but your future is not entirely in your hands. Realizing that one does not have complete control over one's future prosperity is not a bleak perspective. In contrast, such a realization can be a great stress reducer for some students.

Ranking

With the idea that you shouldn't stress already established, let's discuss some of the metrics you might use to evaluate graduate schools. The most obvious is school ranking. Numerous organizations, the best known probably being U.S. News and World Report, provide rankings of U.S. graduate schools and programs. These rankings are typically based on some weighted combination of factors, such as total research funding, research funding per professor, average student GRE scores, Ph.D. students produced each year, number of professors having prestigious awards or fellowships, evaluations of the school by department chairs nationwide, and many other factors. Note that the factors typically do not attempt to measure the actual quality of the education received, and for this reason many people view such rankings with some disdain.

While I would love to say that I think students should ignore rankings and instead look more at the quality of education being provided, I can't say that. The fact is that if you wish to become a

professor, having a Ph.D. from a highly ranked school makes a huge difference. Many top schools simply will not interview faculty candidates who's Ph.D.s are from schools not ranked in the top 5 or 10. Even schools not at the top of the rankings seek out faculty candidates from top ranked schools, to help improve their own image. Their brochures list their faculty by name and Ph.D. granting institution. Even if you land your first job at MIT, if you later move to another university, you are not listed as "formerly at MIT," but rather you are still known by your Ph.D. granting institution. Top government positions are often given to people with advanced degrees from well-known schools. Top companies may target their recruiting activities at top schools, and may in fact have close relationships with those schools, perhaps even having offices right on campus.

Furthermore, a lot of the benefit of graduate schools comes from networking with fellow graduate students, as you will form close lifelong bonds with those people. Not only do you get to know those people intimately well, but you also have this shared experience of graduate school, which is similar in some ways to growing up in the same household. You know these people well, and you trust them – they are almost like extended family. They will help you get jobs at their firms, and you will help them get jobs at your firm – look how often a new U.S. president will surround himself with cabinet members known from his graduate school days. If you go to a top-ranked school, your classmates are probably more likely to become important people in important companies, for a variety of reasons that include intelligence, drive, family connections, etc. – that's just the way it is. At a higher-ranked school, your circle of friends will become more powerful people, and that may ultimately translate to you having a better career.

Thus, ranking isn't the only factor, but I personally would give it a lot of weight. It can make having a successful career somewhat easier.

Research

If you have an idea of the area in which you hope to perform research, then you should obviously seek out schools that conduct research in that area, and that preferably are known leaders in that area. Ideally, the school will have at least two or three professors working in that area.

Many students don't have much of an idea as to which research area they would like to work; that's perfectly normal and nothing to worry about. Choosing a research area after starting graduate school, even a year or two after, is not uncommon.

Advisor

You'll eventually be working with a particular professor who will serve as your advisor. Sometimes students choose to attend a particular school because they have some knowledge of a particular professor. That knowledge may come through professors at the student's undergraduate school, who might know and recommend working with a specific professor or group of professors at another school. That knowledge may come from former undergraduates that went on to graduate school. There are also summer university internship programs for undergraduates designed specifically to attract top undergraduates to a graduate school, and to help professors and students get to know one another.

Most professors maintain web pages too, but the amount of information on professor web pages vary drastically, and no web page can help you determine if an advisor is good to work with or is instead a jerk. An advisor's personality and character are very important factors influencing whether a student/advisor relationship works out.

A department with more professors will offer more choice of advisor than a small department with only a few faculty, though ultimately you only need one advisor, so the smaller department may work out fine.

Most students really don't know much about the professors when selecting a graduate school. If you know something, that's great, but it's common to choose a graduate school without that knowledge. You simply can't know everything when making a decision.

Offers

Graduate schools will often go beyond just admitting students; they admit students with an "offer" that includes some form of financial support. The support typically includes a reduction or waiver of tuition and fees, and a way to gain money, via a teaching assistantship, research assistantship, and/or fellowship. The support may be guaranteed to last for one year, but often is guaranteed for two, three, or even four years. Personally, I don't see much advantage of an offer for four years over an offer for two years, because you really should be working with a professor by the end of two years. Thus, you'll be supported by that professor, not the department, in your third year and afterward; in fact, many students switch over to professor support as early as their second year, and sometimes even in their first year. Most departments that offer "four years" of support know this, and they are counting on the fact that they probably won't have to follow through with support after the second year. Thus, while an offer for three or four years of support may look much stronger than a two-year offer, and certainly provides some amount of added security, it should probably not sway the student too strongly.

Student Success Rates

Some graduate schools fully expect to graduate each student they admit; failures are rare and treated as serious problems. In contrast, some schools regularly expect a significant percentage of students to leave the program, either due to not passing exams along the way, or due to not affiliating with a professor after a year or two. Those schools aren't necessarily any better than the other schools. Instead, they might simply

be using the first two years of graduate school as a filtering process to find the best of the best students (I believe this approach to be unethical). Or, they may have policies that effectively allow any of a large number of professors to veto a student's moving forward, by failing students in required courses or on the students' "Ph.D. qualifying exams" that are taken in the second or third year. Such policies, when coupled with a few professors that like to fail students, can lead to low success rates. Some professors seem to enjoy failing people. I think in some twisted way it makes them feel better about themselves. I obviously don't think that's right; failing students should be very rare if appropriate standards were applied during admissions. A swimming instructor of an advanced swimming class should not take tolerate having students drown; if a significant percentage of students regularly have problems staying afloat, the teacher should not attribute this to his/her high standards of quality. Rather, the standards for admissions into the swimming class should probably be re-evaluated.

Many departments seem to be moving towards policies that do not allow such rogue professors to impart so much control. For example, many schools have switched from department-wide qualifying examinations, in which dozens of faculty from 4-6 different research areas are involved, to subject-specific qualifying examinations, in which just a few professors from the student's specific research area are involved.

If possible, you might ask around a bit to see if you can find out about the success rates of graduate students progressing through in a department. You might have to ask a professor, existing students, or a department staff member for this information, and the information usually isn't published or even written down. One doesn't have to avoid departments with lower success rates, but it might be something to keep in mind.

Culture

Some students need an active social scene surrounding the college to be happy – music, theatre, movies, restaurants, bars, etc. Other students are fine without many those things. If you need the active scene, be careful about going to a remotely located school; I've seen several students switch schools after a year or two because of this reason.

Switching Schools

You are not locked into a school once you join. If you join and find that you are very unhappy, you can apply to other schools and later switch. So on the one hand, you need not be too stressed out when choosing your graduate school. You can make the decision you feel is best, realizing it is not a permanent unchangeable decision.

ON THE OTHER HAND, it is in my opinion unethical to accept a financial support offer from a school with the *intention* of switching schools later. The support offer is made on the assumption that the student intends to complete his/her studies at that school. Having accepted the support, the student has some moral obligation to stay if possible. Some students didn't get into the school of their choice, and thus accept the offer of a different school, with the intention of trying to switch to their desired school in a later year. In some ways, that's like marrying someone now, with the intention of divorcing that person next year if you can finally woo your real love to marry you.

Such a situation is a problem that many departments are trying to deal with (with some adding clauses to their support offers to help deal with the situation). I certainly also understand that for many students, they have to accept that first offer to get into the U.S., yet they want to transfer into the best school possible to maximize their chances of a great career later. Tough situation – I don't know of a good solution. Students should at least realize that such intentional abuse of support offers is not

a victimless abuse. Instead, it hurts professors, who are counting on those admitted students to eventually become researchers that work on the professor's research – which is his/her livelihood. It also hurts other students, who would have been admitted to the school if the abusing student had not taken the spot.

Chapter 3 – The First Years

The first years in graduate school will likely be filled with coursework, teaching assistant duties, and perhaps even adjusting to a new country or city. A student probably won't get a lot of research done that first year, though some students do manage to do so. There are a few things a student can do to make the first year a more enjoyable and fruitful one.

Classes – Start Easy

In your first semester of graduate school, take the easier classes first. If you think that sounds cowardly, let me explain.

When first learning to play tennis, you wouldn't start off by playing against a super-aggressive power player. You'd start off with someone mellower. Once you developed your skills, you might then challenge that power player, and show him a thing or two.

Likewise, when first attending graduate school, it's unwise to immediately jump into one of the harder courses that gives a lot of work and/or is known for its low grades.

Many graduate students find this advice somewhat offensive. They are smart, darn it, and they can handle any course just fine. Other students prefer to take the hard course right away, to get the course done so they don't have to worry about any longer, like swallowing that nasty-tasting broccoli first before enjoying the rest of your tasty meal.

However, there are some good reasons to put off those harder courses until later. Most importantly is your grade point average (GPA). A graduate program may require you to maintain a total GPA above a

certain number, such as above 3.5, where a B counts as 3.0 and an A as 4.0. If you take a hard course in your first semester, you may get a low grade. Because you haven't taken any courses before, that low grade will have a big impact on your total GPA, and thus may bring your total GPA under the minimum – jeopardizing your status in the program or jeopardizing the teaching assistantship. If you take easier courses first and get As, a low grade later on is a smaller portion of your total GPA and thus doesn't bring it down much.

A second reason to take easier courses first is to give yourself a chance to get adjusted to your new life. By the time you take that harder course, you will already have taken care of apartment contracts, bank accounts, visa issues, workspace issues, forming a social support network, etc., and you can spend more time and mental effort on the course.

Finally, there's your own motivation to think about. Getting a low grade in your first course may demoralize you, causing you to get lower grades in later easier courses than you would have gotten had you taken those courses first.

You'll need to talk to other students, and maybe the graduate secretary, graduate advisor, or your own advisor (if you have one), to figure out which are the harder courses. At the time of this writing, a website that publishes grade data on professors, with that data obtained directly from school records, is just coming online (www.pickaprof.com), so public data on courses and professors may be available too.

You might also be wise to talk to fellow students about professors themselves. Most professors are reasonable and competent people. But there are a few that are unreasonable and/or less than competent. They may give too much work, create tests that don't measure what's being taught in the course, teach in an incoherent way, have grading policies that rely almost as much on random factors as on evaluating your work,

have overly-harsh grading standards, or have interpretations of grade meanings that are harsher than the norm. That some professors are bad at their jobs should be no surprise; some lawyers lose many cases and some doctors don't diagnose illnesses well, but they are still in business. Why should professors be any different? Talk to graduate students who have been around a while or to others in the department, and see if you can avoid these bad apples.

None of the above advice should be interpreted as encouraging laziness or taking the path of least resistance. Learning requires hard work – you need to put in the time, and you need to experience some mental anguish, to make progress. But the fact is that some courses are harder than others, and its best to take those later rather than sooner. And the fact is that some professors stink at their job, and it's best to avoid them entirely if you can.

TAing

Working as a teaching assistant (TA) is perhaps one of the most rewarding, enlightening, and enjoyable parts of graduate school. Some programs require each graduate student to TA for at least a semester or two, even if the student has a research assistantship or fellowship, in part so that every student can benefit from the TAing experience (and in part to obtain TA coverage of the department's classes).

TAing usually involves running one or more discussion/recitation sessions or lab sessions, preparing homeworks and exams, and grading, among other things.

The strangest thing about TAing is that young graduate students, who just last year were undergraduate students themselves, are suddenly placed in a teaching role as TAs but with almost no training at all. Some schools do offer a small amount of training, but you shouldn't count on much. There's just not enough time or resources to provide thorough TA

training in most cases, and the usefulness of much of that training is questionable.

One way to adjust to and learn the job of running a discussion or lab session, in my opinion, is to partner with a more experienced TA for a few weeks or months. I believe that departments should require such an apprentice-based introduction to TAing, though few do. In fact, I have found that TAs are happiest when they work side-by-side with another TA when running challenging discussion sessions or lab sessions. The camaraderie fulfills a social need, and the undergraduate students are less likely to cause trouble when there are two authorities around – there's strength in numbers.

Prepare before the semester begins. Make sure to get the textbook and/or notes for the course, and read through them before the semester. You'll already have a leg up on your students, who will likely put off the reading or will just skim the material (undergraduates today generally don't seem to like to read very much). Meet with the professor before the semester also. If a class is well established, prepare as many labs or homeworks as possible early. Remember, your semester will get busier as the semester moves along, so the more TA work you got done early, the smoother of a semester you'll have.

Communicate with your fellow TAs and your professor. Miscommunications are often the source of problems. Your professor or other TAs may have thought that you were maintaining the grades, or preparing the homework, or grading the quizzes before next lecture, while you thought someone else was handling the task, for example. You may grade an entire stack of quizzes following one standard, only to find out that the professor was expecting a totally different standard. Try to meet regularly to talk about the course. Not only will this reduce miscommunications, but it also often leads to improvements based on detected issues and new solutions. As *a general rule in graduate school as well as your eventual job, it is better to over-communicate than to*

under-communicate. Send emails to confirm things, just to be sure. "I plan to grade the quiz tomorrow." "So you are preparing next week's assignment, right?" "As we discussed in person, you'll be preparing the exam and I'll be grading it." Don't assume things. Saving a bit of typing or a bit of breath is no big savings really; in contrast, miscommunications can be a big loss.

Many TAs are too hard on the students when grading. A little bit of knowledge is a dangerous thing. Some TAs know the material very well, and expect the students to know the material just as well, giving low grades otherwise. Remember that the students are just learning, and your grading should be cognizant of that fact. You wouldn't expect a beginner swimmer to be able to swim 10 pool laps without stopping – swimming 2 laps for a beginner would be an "A". You are advanced, but the students are still beginners; don't evaluate beginners as if they were advanced. Also remember that grades impact motivation; overly harsh grading early in a course can demoralize students, who do more poorly in the rest of the quarter than they would have done otherwise. I'm not suggesting handing out As to everyone and thus inflating grades. Instead, I've observed that many young teachers (TAs as well a professors) often err too far on the side of being harsh, lightening up years later after they become more confident in themselves and more understanding of the wider context of education. Try to skip that first overly-harsh phase of teaching, and move straight to the later more-enlightened phase.

Good TAing is a delicate balance between giving students answers and having them learn to figure out the answers themselves. The adage "Give a man a fish and he eats for a day; teach a man to fish and he eats for a lifetime" applies here. Clearly, it's more important for us to teach students to learn to figure out the answers themselves. But students can lose motivation if they aren't making any progress, in which case giving them the answer (or part of it, or even hints) can keep them moving

along. Finding the right balance that keeps the students highly motivated is really the art of teaching.

Be aware that motivation is a key part of effective learning. If a student is motivated to learn a subject, they will learn it far more easily and quickly than if they don't see the relevance of the subject. Thus, a key part of teaching involves convincing the student that the subject is worth learning. Furthermore, a teacher who displays enthusiasm for the subject may find that the enthusiasm motivates the student. Discussing why a subject should be learned, and showing enthusiasm for a subject, may be something that some graduate students feel uncomfortable with. Many graduate students are from other countries or school systems, where there was no such "coddling" of the students. But those graduate students should strive to realize that they aren't in their home country or teaching in the same system in which they learned. The undergraduate students in your new graduate school may be very different from what you are accustomed to. Yet, you are being paid to teach *those* students, not students from your own country or school. Comments from graduate students (and professors) that some people find very annoying are comments along the lines of: In MY country or in MY school we did things differently. That's fine, but we are not in your country or in your school, we are here. While it's fine to share what you believe are positive aspects of your previous school, your priority must be to adapt to the situation at your present school.

An area where many TAs need to improve is how they speak to a class. Many tend to talk in their normal voice, as if they were talking to a friend in a coffee shop. That doesn't work. Speaking to a class requires projecting your voice, and speaking slower than normal, as if you were talking to that same friend in the coffee shop but the friend was sitting clear across the room. Projecting your voice feels awkward to new teachers, especially those on the shy or quite end of the personality

spectrum. Practice talking loud, perhaps at home in front of the mirror. Practice pronouncing your words clearly and slowly.

Speaking slowly is especially important if you have an accent. Be aware that many undergraduate students (especially in America) have little experience listening to people with an accent. The complaints from undergraduate students about not being able to understand the TA can baffle a young TA with an accent, whose fellow graduate students and whose professors have no problem understanding the TA's English. The undergraduates are just not used to accents; many American undergraduates cannot even understand English spoken by someone from England or India even though that person uses perfect English grammar. The accent just somehow shuts off the student's attention. The TA can help by speaking slowly, pronouncing each word, and projecting his/her voice.

American undergraduate students act cool and aloof, but I believe that deep down, most of them either want to get excited about learning or have the capacity to do so. However, high school in the U.S. is fairly brutal for smart kids. Being smart is not cool, and can lead to teasing and estrangement. Only a kid with a tremendously high self-esteem and/or a very supportive family can tolerate or avoid being treated poorly. Thus, many smart high-school kids learn to hide their intelligence, and to never admit to anyone (including themselves) that learning is fun or interesting – academics is for nerds. In college, some students are able to quickly free themselves from this self-created cage, but others take much longer, and some never do. Be patient with them, and try to have some compassion. If you stay positive and provide an example, rather than condemning them, perhaps you will be one of the people that helps them break from that cage and embrace learning.

Pay some attention to your grooming and clothing. People do tend to have more respect for well-groomed well-dressed people. Think about airline pilots – they wear uniforms for a reason (would you want to fly on

a plane piloted by a guy in shorts and a tank top?). No need to go overboard with fancy clothes, but some attention to this area may earn you a bit more respect in the classroom or lab.

Help the students help each other. Use study groups explicitly, by requiring such groups or giving extra credit for group work on assignments like homework or projects. Students learn a lot from each other, and they pay attention to each other. Remember, the undergraduate's goals sometime relate as much to social accomplishments as they do to academic ones. In a busy technical lab section, pairing or otherwise grouping students may not only lead to more learning, but may make the TAs job much easier. At my university, our TAs long complained that running a 28-student programming lab section was nearly impossible – students would be waiting for long periods of time with their hands raised to ask questions, as the TA would go from computer to computer trying to help. Then, a few years ago, we began using "pair programming," in which we paired students up in lab. The TA complaints stopped almost immediately – one TA even said he was bored. The students answered each other's questions. Some people believe this approach hurts students, who might rely too heavily on others and thus not learn themselves. However, test scores in lecture did not drop, retention rates improved, students seemed happier, and students did just as well in follow-on courses. Several studies at other universities have shown that group learning is extremely effective. Grouping students seems to be a win-win idea – the students win, and the TA wins too.

Remember that lecture is generally known to be an ineffective mode of teaching. Keep your lecture segments to a minimum, perhaps 10-15 minutes. Then give the students something active to do, perhaps a problem to be solved individually, or even better to be solved in groups. People learn by doing. You can watch somebody tie a shoe all day long, but until you try it yourself, you'll never learn to tie a shoe. Let students present their solutions to the class – that not only gives them the

opportunity to learn to present and to build confidence, but also lets you and others see different ways of solving a problem and to even learn from each others' mistakes. This more interactive, non-lecture method of running a class means you will necessarily cover less material than if you lectured the entire time. But realize that students won't remember all that lecture material anyways, so you're really not losing much. Better for students to actually learn a little material than to not learn a lot of material.

Casual Conversations and Socializing

In several of the above-discussed subjects and in subjects that will follow, I recommend that you talk to fellow students to get the needed information on courses, professors, and more. Most of that talking is done in informal conversations. So you need to be having informal conversations with lots of different people. That means you need to be social, to step out of your comfort zone and actively talk with people you don't already know.

Some graduate students don't like the idea of getting to know people for what they claim to be selfish interests. "You should befriend people for friendship purposes only, not to get information." But the reality is that there is nothing at all wrong with making friendly contacts with people in order to expand your network of contacts with which you share information. You're not pretending to make blood buddies out there; you're just making friendly contacts. Additionally, those contacts are mutually beneficial – you'll be providing information also.

Socializing can be especially hard for students from other countries. In many graduate schools, the Indians stick with the Indians, the Chinese with the Chinese, the Greeks with the Greeks, and so on. This is natural – a person will clearly be more comfortable being with people that share their own culture, language, and background, especially when that person

suddenly finds herself in a new foreign land, whose unfamiliarity can be rather intimidating. I strongly encourage students to reach out and get to know people outside their cultural group. People new to the country should strive to get to know the local graduate students. Likewise, the local graduate students should get to know people new to the country. There's a lot of information out there on which no one group has a monopoly.

Socializing doesn't mean you have to go out to coffee shops or go to the movies with these other people. It minimally just means that you talk with them when you see them. You may see them in the hallway, at the sinks in the bathroom, at the water fountain, in the mailroom, or in the department office.[2] When you see someone, you have a choice. You can just say nothing, or you can say hi and carry on a friendly conversation. Do the latter, and start from the day you step foot on campus; every time you pass someone without introducing yourself, the harder it becomes to introduce yourself later ("Hello. I've passed you in the hallway silently for 6 months, but today I thought I'd say hi" is a lot more awkward than just introducing yourself on the first day). Each time, your conversation gets a bit longer, a bit more personal. And once in a while, very useful information gets exchanged.

Which brings us to a simple point – never underestimate the importance of casual conversations with colleagues. Those conversations are where much important information gets passed on. I remember reading about a company manager who was concerned that employees were spending too much time conversing around the coffee machine rather than doing "work" at their desks. So he removed the coffee machine. Soon, he found that there was an increase in calls and emails to the computer systems support group, asking for help with minor

[2] A wise department ensures that the department environment encourages casual interactions, via a central mailroom, co-location of research labs, a central graduate student lounge with snacks, etc.

computer problems. Eventually, the manager realized that employees had been solving those problems during those casual conversations at the coffee machine. That's just one example of the type of information exchanged through casual conversations.

Don't assume that you have to ask specific questions in order for these conversations to be useful, like "How can I succeed in class X?" Instead, often times the useful information comes from seemingly random conversations. One day, you might see a fellow student, and in having a brief casual conversation you might mention, "Class X is getting hard." The person you are talking with might then share some simple information that makes a huge difference – "Did you know that the professor draws a lot of his information from textbook Y, even though he doesn't require that textbook?" "No, nobody told me that." "Oh, it makes all the difference. He usually tells the students, but some quarters he forgets, and then he wonders why all the students don't do well." Casual conversations are like a mine filled with gems at random places slightly below the ground surface. Just wandering around the mine can kick up a beauty of a gem.

The above may seem obvious to people who are already quite social. But among highly intelligent people, one seems to find a higher share of extremely unsocial people. I remember a fellow graduate student in our research lab that wouldn't even say hi when you greeted him in the hallway (sometimes he would manage to grunt, though). When I eventually asked him why, he said, "What's the point? You say hi, then I say hi. Nothing has been accomplished. It's just a waste of breath." This student certainly never received the benefit of casual conversations. He was there for years, but he never finished, and while the dozens of other graduate students stay in close touch and have even gotten each other jobs, people hardly remember that one person.

When you become a more senior student, be conscious of how helpful the information you have may be to the newer graduate students,

and strike up casual conversations with them for their own benefit. It's a simple way to help people.

Some students make excellent use of lunchtime to interact with fellow students. They make a routine of going out to a local eatery a few times a week, or eat together with fellow students in a graduate lounge. Some even develop a routine of going to lunch with their advisor along with fellow students. Sharing a meal has long been known to create a bonding effect among humans (look at the importance of shared meals in many families and cultures). Mealtime is also a perfect time for casual conversations. I personally thoroughly enjoy my lunches with my graduate students (I hope they feel the same way!). Not only does it result in a bond of friendship among us, but I learn a lot of technical information from some of the conversations, and I find that about half of the information my students and I need to exchange for work gets exchanged casually during these lunches, decreasing the need for formal meetings. Regarding money, don't just think about minimizing cost of lunch; consider lunches with fellow students or faculty as an investment into your social and academic future. Plus, it's just fun.

Read, Read, Read

As you embark on your new life in graduate school, get into the habit of reading every day. The reading habit will greatly aid your graduate studies as well as your lifelong career. Even with modern media technologies, reading is still the most effective way to acquire part of the vast accumulated knowledge of mankind. Reading helps form your perspectives on the world, broadens your horizons, exercises your brain, grows your knowledge base, and helps lead to insights and ideas.

Read textbooks, research papers, magazines, books, and similar materials (hardcopy or electronic). Read textbooks in your field and in related fields. Reading a textbook without taking a course may sound

strange, but you are now an accomplished learner and critical thinker – you can learn a significant percentage of the textbook material without a teacher, homeworks, and exams. Textbooks give you solid foundational information. Read research papers in your field. Knowing what others are doing in your field is crucial to knowing how your work relates to others, to learn what is considered a research contribution by your research community, to gain new research ideas, and to be able to speak intelligently on your research subject. Don't just read, but try to organize the concepts too; after reading 7 papers on a related subject, think about what criteria distinguishes the papers, how some of them are similar to some of the others, etc. Read magazines published by professional societies related to your field – those magazines give you a broader perspective of the field, both across subjects, and across history. Read magazines that have nothing at all to do with your field, but that discuss subjects that may someone help you form a vision or insights. Read books on things inside and outside your field, on technical and non-technical subjects.

As will be discussed later, avoid reading things that are fleeting, like the daily news. Such items are designed to grab your attention, but provide little long term benefit.

That's a lot of reading, and most of it is not associated with a deadline. Thus, you will be tempted to put off reading and instead work on more urgent tasks. Yet, reading is fundamentally important, and you should make sure you read a lot. One way to make sure you read sufficiently is to make it part of your daily routine. Brushing your teeth if part of your daily routine, and you don't neglect it just because you have a stack of exams to grade (at least I hope you don't, for your colleagues sake).

So make reading something you do every day, part of your routine. Have a number of items in your "to read" pile. Determine a time and minimum duration when you'll read something from the pile every day.

Make the minimum duration something that you can actually maintain every day, like just 10 minutes. I know that sounds like a small amount of time, but even just 10 minutes a day adds up to about 4 hours a month, which is more discretionary reading than many people do. But you'll find that you'll often read for longer than the minimum, when reading a particular interesting subject or wanting to finish off a chapter. So this 10-minute a day habit may result in 5-10 hours a month of voluntary reading – nearly a full work day's worth. (By the way, this 10-minute a day habit also works well for exercise – some people even combine the two, and read while they ride an exercise bike or run on a treadmill). You might select your 10 minutes to be first thing in the morning before leaving for campus, the time right after arriving to your office or just before leaving the office for the day, the time just before, during, or after lunch, the time while you exercise on a treadmill or stationary bike, the time you take a break in the afternoon, and so on.

If possible, keep a pen and notebook handy as you read. Reading will likely cause you to think of new ideas or to decide to do something. Write those things down so you don't forget, in the notebook, or even on the item you are reading if that makes sense (e.g., on the back of a research paper).

Keep in mind that humans are creatures of habit. Developing a new habit can be difficult, but after you do something regularly for a few weeks, it becomes easy and natural. So don't worry too much if it seems hard the first couple weeks; it gets easier as time goes on. And again, if you want to develop a habit, make it achievable. If you tell yourself you will read for one hour every day, you probably won't achieve the goal, and frustrated with your failure, you may not read at all. Take small steps first.

Read in Multiple Passes

Reading a research paper, scientific article, textbook, or similar informational piece is much different than reading a story. A story, such as a mystery novel, is specifically intended to keep the ending a surprise. A story takes the reader on a journey, through lows and highs, periods of suspense, and climatic endings. When you read a story, you should start at the beginning and read the story in sequence, just as the author intended, to get the full effect. Reading an information piece is very different. You usually should NOT start from the beginning and read the information in sequence. Instead, you should make a 5-10 minute first pass over the paper to determine if the item is worth reading at all. For a paper, read the abstract and/or introduction, skim over the details, look over the results, and read the conclusions. The information you gain on such a first pass may actually be sufficient for your purposes.

Your first pass might lead you to conclude that a more thorough reading is warranted. In this case, you might want to do a *time-constrained* reading. Decide approximately how much more time you are willing to spend reading this paper – 15 minutes, 60 minutes, maybe several hours. That decision is based largely on how relevant the subject matter is to your needs and interests. Then, read according to that time constraint, skimming sections as necessary. Most young graduate students are unaccustomed to time-constrained reading. Many students feel that, if they are going to read the paper, they should read it "properly." "A job worth doing is worth doing right." True, but "properly" and "right" are relative terms. Spending more time on a paper than makes sense based on the subject matter is not "proper" or "right."

I remember my wife and I taking a road trip vacation during winter break of graduate school, driving across the country to see some sights and visit some friends. An older couple along the way chastised us for not "properly" seeing the sights, spending only a few days when clearly

weeks were needed to properly take in all those sights. Well, we only had a few days, not weeks – sorry. Seeing what we did see was better than seeing nothing at all. What was proper for that older couple was different than what was proper for us. It's all relative. Experienced graduate students and professors have mastered the skill of time-constrained reading.

Be aware that the literature is filled with lousy papers. The pressure to publish lots of papers results in some papers being written in a rush, with the authors themselves confused about their main point. If the authors don't understand their main point, how can you, the reader, be expected to understand it? The peer-review process doesn't filter out all such papers; sometimes time-strapped reviewers accept the paper based on a high-level reading rather than a careful reading. Sometimes the reviewers simply aren't qualified to review for a particular subject, and thus err on the side of the authors by accepting the paper. Papers in highly-respected conferences/journals may have a higher rate of good understandable papers, but even those forums aren't immune to incomprehensible papers or papers without a clear point. So if you are having a hard time reading a paper, it's not necessarily your fault. You might wish to put such a paper aside for a while, and move on to the next paper.

Chapter 4 – Choosing an Advisor

Your Ph.D. advisor has a major influence on your graduate school experience and your ultimate success. Again, don't stress – this is another good kind of decision to have to make in life, and you can't possibly predict everything anyways.

The term "advisor" is really a misnomer. In most cases, your advisor doesn't give you advice. Your advisor may tell you what to work on, when things are due, what papers to submit, what hours to keep, whether you can work as a summer intern in the summer (which he/she may help you find through his/her connections), and he/she may be paying your salary. In other words, your advisor isn't your advisor, your advisor is your *boss*. Ultimately, your Ph.D. advisor determines whether you will get a Ph.D. or not; it's nearly impossible to pass your final defense if your Ph.D. advisor doesn't sign off. Graduate studies are somewhat unique in that you, the "employee," get to shop around a bit for your own boss.

The most obvious criterion for choosing an advisor is to find an advisor in your desired research area. But there's more to it. We now discuss some of the other, less obvious, criteria.

Know the Students

One way to judge a tree is by the fruit that it bears. Likewise, one way to judge an advisor is by the students he/she bears. Do the students successfully complete their Ph.D.s., and do they do so in a reasonable amount of time? Do they get good jobs after completion? Do they

publish papers? Do they get credit for their work? Are they happy while they work for the advisor? Do they stay in touch with their advisor after graduating? You might find some of information about a professor's students on the professor's web page, which is usually a good sign – a professor proud of his students may one day be proud of you too.

However, the best way to get information about a professor's students is to talk directly with students, and even with department staff. While on rare occasion this might mean arranging a formal time to talk with a student, in the vast majority of cases it instead involves casual conversations. These casual conversations may take place in classes that you are taking, through your teaching assistant duties, through department-organized social events, through student organized events (like soccer or volleyball), through informal conversations with the department's graduate secretary and other staff, and through a myriad of other settings. Thus, the task of picking an advisor is tremendously aided if you are a social person who can comfortably talk with other students. If students and staff feel comfortable with you, the information that a graduate secretary or senior Ph.D. student may give you informally in the hallway may be worth 10 times more than the information you meticulously gather through more formal means.

Choosing an Advisor is More like Dating than Making a Decision

An important part in choosing an advisor is to find an advisor that you work well with. In this regard, there is no such thing as an ideal advisor, any more than there is such a thing as an ideal spouse – what defines a good advisor or a good spouse is highly dependent on your own personality, interests, and needs. It follows that, in some ways, finding the right advisor is more like dating than it is making an objective decision. Some students are highly motivated, while others need to be inspired. Some students can tolerate being lectured and criticized, while

others need a gentler, more constructive advisor. Some students may do better with a highly-structured advisor, while other students will thrive if left more on their own. Some students want a topic defined for them, while others wish to have more control over their topic. An entrepreneurial student may prefer an advisor who is well networked with top people in the field and with industry, while a more theoretically-oriented student might prefer a professor willing to spend hours discussing research problems in detail.

The best way to determine if you work well with a professor is to simply work with the professor. After narrowing down the set of possible advisors, some students will then try to get involved in some sort of short-term project with the professor. This project is often done on a volunteer basis. Another way of working with a professor is as his/her teaching assistant. Yet another way to work with a professor is to enroll in a course taught by the professor. The wise student will not just do well in the course, but will go beyond the course requirements, turning the coursework into a short-term project itself. A surprising number of papers published in research conferences are actually extended projects from graduate coursework. Through a combination of the above methods, students can spend their first year or so of graduate school getting to know a variety of professors. Often, "choosing" an advisor evolves naturally, with the student expanding on the project, the professor beginning to pay the student, the professor providing lab space for the student, etc. In my experience, more often than not, there's a natural pairing up of students to their advisors over time, as opposed to an evaluation period followed by a "decision." Indeed, finding an advisor is much more like dating (American-style dating, that is) than a decision.

While most professors are good people, it's important to be aware of the fact that some professors are – for lack of a better term – jerks. Some professors have something to prove, and they try to build themselves up by pushing their students down, as well as others around

them, down. Some are just people who never had to learn how to behave properly. In most professional careers, such behavior is not tolerated, but professorial tenure (which essentially means a job for life as long as the professor doesn't break the law) tends to allow strange behaviors to evolve and persist. I recall reading about a study on the "Most obnoxious career-people" or something like that, with doctors and professors being at the top of the list. Of course, not all rude professors are bad advisors – some are brilliant people that produce great Ph.D. students, and behind that rough exterior is sometimes a very caring person. The key here is that it's very hard to select an advisor from a web page or a list of research topics, just as it hard as it would be to pick a spouse solely from information on an online dating web site. Working with a professor, talking with his/her students, and especially looking at the success rate of past students, can help you determine if this is a good advisor for you.

Much of the above discussion relies on a professor's past performance. However, many professors are young professors who have not been around long enough to have much of a track record. This should not stop students from considering young professors as potential advisors. A young professor may provide several unique benefits. First, a young professor may pay more attention to his/her graduate students. This is because he/she may not be as busy with other commitments, such as organizing conferences, serving on campus committees, taking part in department administration, managing a large group of graduate students, and so on. In fact, young professors are often given course releases (they teach fewer courses per year than more senior faculty) so that they can focus on research and proposal writing, which are key components of getting tenure. Second, a young professor is highly motivated to have his/her students publish, and this motivation can work in the student's favor – the professor will likely encourage paper submissions, spend a lot of time helping to write those papers, and may even give some of the paper presentation talks at conferences. Third, a young professor may be

more directly involved in the research than a more senior professor. He/she may spend hours discussing extensive details with the students, and in some cases may even help carry out the research in the lab. Fourth, a smaller age difference between advisor and student might enhance the advisor/student relationship. Of course, attention, publication-orientation, direct involvement, and a good relationship, may be present with senior professors too, but the point is that young professors are likely to be especially strong in those categories. A drawback is that the young professor is still learning how to work with students, and may make more mistakes; but this drawback is often compensated for by the younger professor's unwillingness to be too hard on students due to the professor realizing his/her own weakness. Another drawback is that a professor may be denied tenure before his advised students have finished their degrees. Generally, young professors are given 5-6 years to prove themselves, after which they are either granted tenure (job for life), or denied tenure (he/she has one year to leave). The students of a professor denied tenure might have to scramble to finish up quickly (if nearly done) or to find another professor willing to take over as advisor. Of course, a professor leaving is not a phenomenon unique to young professors; senior professors are more likely to take on a time-consuming administrative position, take a leave of absence to start a company, move to another university, or die. Thus, looking at all the factors, one concludes that young professors should be considered as potential advisors just as readily as senior professors.

Be Aware that Professors are Competing for Students

Be aware that while you are looking for an advisor, the professors themselves are looking for students, resulting in competition among the professors. This competition can result in a delicate dance, in which

professors try to woo strong students, but at the same time try to avoid stealing students that are becoming affiliated with other professors. There's a sort of unspoken (or sometimes spoken) territorial marking that goes on, with the marking occurring as early as when graduate student admission decisions are being made. Professors generally try to respect other professors' territory. Of course, students are not forced to stay with a particular professor, but students should be aware of the inter-professor competition and thus act wisely when evaluating professors. Mini-projects should not be presented to the professor as having chosen an advisor. Once a student has worked long enough with a professor, switching without an obvious reason becomes more difficult, as other professors do not want to be seen as stealing another professor's good student. And if a student "abuses" a professor, perhaps by saying he'll work with professor A in order to gain admission to the university, but then switching to professor B without good reason, the student may find that professor B is not willing to take on the student simply because he does not want to offend professor A. Remember, graduate students come and go, but the professors of a department must work together for decades. The key point is to be wise in your interactions with professors, communicating your intentions clearly, and never burning any bridges.

Working a bit with multiple professors sometimes has another advantage. Many professors would like to collaborate more with their fellow faculty, but finding the means for collaboration is sometimes hard. If you really like a couple of professors, you might be able to serve as a point of collaboration between them. Of course, you have to pick one of them as your main advisor. But the other one can serve on your committees. He/she can participate in joint research meetings and co-author papers with you and your advisor, often done along with one of his/her own students.

Chapter 5 – Time Management

Time management is the job of balancing multiple tasks that you must do.

The Key to Multitasking – Don't

The key to balancing multiple tasks in graduate school is: DON'T.

Let me explain. Before your graduate studies, undergraduate studies require balancing multiple tasks -- different assignments in different classes, dorm activities, professional club activities, part-time job, etc. After your graduate studies, your job will also involve multiple tasks, like preparing a report for your boss, finishing up some data analysis, and preparing slides for a meeting. In between these two multiple-tasked times lies graduate school.

Graduate school is not a time for doing a dizzying array of different tasks. Rather, it's a time for in-depth focus on one task. *One problem.* For this unique and special time of your life, you can look at one problem in more depth than anyone else in the world. You will be an expert on that problem. You will have deep insight into that problem. You will have a respect and appreciation for how a problem that looks simple on the surface is really quite complex when you dive deeply into it. Maybe that problem is something you will continue to work on in your future career, or maybe not. In either case, the experience of solving a problem in great depth is one of the things that shapes a Ph.D.'s intellectual perspective, and that leads to his/her intellectual maturity.

Allow yourself have this in-depth experience. To allow it, you have to avoid having many tasks going on. Free yourself from other work/school commitments – teaching assistant duties, non-essential committee work, etc. Focus on your research problem. Think about it when you take a shower. Think about when you walk to school. Read everything you can on it. Talk to people about it. Analyze it, ponder it, breathe it! Going deeply into a problem is really a wonderful experience, and something that will teach you that real progress does not just come from meetings and fancy presentations and brainstorming sessions and polished reports. It also comes from deep insights and reflection.

The above doesn't mean to throw away your personal life. You must stay balanced to be happy. The graduate school years are important years socially – you are building lifelong friends, and maybe even getting to know your future spouse, or perhaps starting your own family. But in your school life, try to focus on the main task of solving your chosen problem in great depth.

For these few special years of your life, let yourself focus on one task in great depth. You'll likely never get another chance to do so again in your future experiences. These graduate school years will very positively influence you in future experiences.

Stay on Path

With the above said, we must recognize that you will likely still have multiple tasks competing for your time, especially at the very beginning of your studies when you are taking several courses and acting as a teaching assistant while trying to get started on your research, and at the very end when you may be interviewing for jobs or teaching a course while trying to finish your research.

The key to time management in the presence of multiple tasks is to recognize two points:

(1) You can't do everything, and
(2) Tasks on the path to your main goal should get priority.

For example, suppose it is Monday and you have two tasks pending. One is to write a research paper for a conference that your advisor says is important and is due on Friday. Another is to grade a lab assignment for a class you are TAing and that is due on Wednesday. You have 4 hours available for each of the next few evenings. You estimate each task will take 4 hours. You might be tempted to do the task with the closer deadline first, namely to grade the lab assignments tonight. But, grading a lab assignment is only indirectly on the path to your goal – you need to do well as a TA to keep your TA support, you want to get good teaching experience, and you of course care that the students get good feedback. In contrast, the paper is much more directly on the path to your goal – you need to publish papers to focus your own research ideas, to disseminate your research, to establish your Ph.D. research in the community, to strengthen your resume to get a good job, and to keep your advisor happy with you. Giving priority to tasks on the path to your main goal, you might decide instead to spend this evening writing a first draft of the paper, that you then can give to your advisor for feedback. You'll grade the labs tomorrow evening. In other words, you do the more important task first, *not* the task with the closer deadline.

Not only is this decision probably best for you in the long run because you give your advisor more time to provide feedback, but you might also find that the paper needs two hours more than you expected. Tasks *often take longer than expected.* You finish up the paper on Tuesday night and then send the paper to the advisor. That leaves only two hours for grading the labs, which are due Wednesday, and a task that really required four hours. So, you have to do an adequate but not great job of grading of labs to get them done on time, or you may have to return them late.

Now if you are like many readers, you are thinking: "That's not right!" Shouldn't you always do quality work? Shouldn't the lab assignment have been graded better? In a world with infinite time and resources, yes. But in the real world, people MUST MAKE TRADEOFFS. Those aren't just words. Making tradeoffs *translates into some tasks not being done as well as possible*, being done late, or not being done at all. That's life, and there is no way around it, whether looking at graduate school, work, raising a family, etc. Time is not infinite. You cannot have your cake and eat it too.

The mistake that many graduate students make is to spend time on tasks with earlier deadlines. Tasks often take longer than expected, and thus by choosing to do less important tasks first, they are reducing time available for the more important tasks. I've seen many students go overboard being great TAs, only to let their research falter, resulting ultimately in the student not completing a Ph.D. While that person was likely a great TA, he might have had more impact on students overall had he merely been a good rather than great TA, enabling more focus on research, finishing his Ph.D., becoming a professor, and thus helping even more students over a lifelong teaching career. It's like the difference between a sprinter and a marathon runner – if you watch both of them for one minute, the marathon runner looks lazy. But if the goal is to run 26 miles and not just 1 mile, then the marathon runner proves to be the wiser of the two – he is still running after two hours, long after the sprinter has dropped out of the race.

Another common mistake that graduate students make is to choose to do a well-defined task first (e.g., grade a stack of homeworks) rather than a vaguer task (e.g., figure out how to improve on someone else's research), even though the vaguer task is more important. We all enjoy the satisfaction of completing a task, and a well-defined task is more likely to have a clear completion point. But that should not mean that we do those tasks first.

Steven Covey, author of numerous books on self-improvement and making effective use of time (including the well-known book "Seven Habits of Highly Effective People"), refers to these concepts as the distinction between *importance* and *urgency*. Some tasks are important but not urgent, some are urgent but not important, and some are both. (Obviously, such task categories are not clearly separated but rather a continuum.) He stresses that people should focus more energy on the important tasks, even if they are less urgent.

With respect specifically to tasks related to TAing, be aware that being a great TA doesn't necessarily mean spending long hours on TAing. The best TAs seem to be the ones that inspire the students, and that provide just the right challenge to the students. Don't interpret my earlier example as suggesting that graduate students should not strive to be good TAs. The point is that being a good TA isn't necessarily achieved by putting in long hours. My own graduate students have been frequent recipients of the department's "TA of the Year" award. I believe those awards are because of their attitude towards the student more than anything else. They don't put unusually long hours into TAing.

Focusing on important tasks means making time for those tasks. For example, if reading a stack of papers is important, even if not urgent, then you might plan to read for 15 minutes at the start of every work day, before beginning to work on the more urgent tasks of the day.

Get Organized

Managing multiple tasks can be greatly aided by keeping a list of things to do, known as a "to-do list." I find it helpful to keep such a to-do list of tasks. I am in the habit of checking the list a couple times a day to see what tasks are pending. I rearrange the list regularly (it's just a text file on my computer) based on my constantly adjusting the priorities of the tasks as time goes on and as tasks come and go, and I try to work on important non-urgent tasks for one or two days a week. The list usually

has only 2-5 items at the top, and then a list of longer-term tasks further below. (By the way, when I complete a task, I don't delete it, but I move it to the bottom of the file where I then have a list of completed tasks organized by month; I have found such a completed-tasks list to come in very handy several times). Other people make more detailed lists or different types of lists. Some keep the list in a Personal Digital Assistant or some other portable electronic device. Whatever works best for you is what is right.

In addition to keeping a to-do list, it's a good idea to write down what you have done, results of meetings, new ideas, etc. There are many organizational aids that promise to help you accomplish more in less time, like online calendar programs, Personal Digital Assistants, tablet PCs, scheduling programs, and much more. Even with all these electronic aids, I've found that the good old paper notebook seems to satisfy most of a graduate student's needs. It's cheap, easy to carry, easy to write in, and reliable. In this regard, the only thing distinguishing a student of today from the student of a century ago may be that the student of today may occasionally scan in the pages of the notebook, to prevent complete loss of the information if the notebook is lost. Otherwise, like blue jeans, the trusty notebook seems to be just as useful now as it was a hundred years ago.

Stop Reading Email So Much

To be more productive, stop reading your email so much. That goes for instant messaging, text messaging, blog reading, and similar endeavors. Those things not only take up too much of your valuable time, but they interrupt your important work.

Many people set unimportant goals (explicitly or implicitly) of keeping their email inbox small, responding promptly to messages, and keeping up with blogs. They thus spend a lot of their time on those tasks, and interrupt their work to do those tasks. When an email arrives, they

read it. When someone instant messages them, they respond. However, those tasks are not very important, with the exception perhaps if the item comes from your advisor or key students or staff. So spending time on those urgent but unimportant tasks is generally unwise.

It's better to do those tasks at particular times, such as just before heading out to lunch, and before going home for the day. By doing those unimportant things once in a while and also just before another event, you minimize the time spent.

A key productivity killer is the interruption. When you are interrupted, you lose your train of thought. After the interruption is over, it takes time to recollect your thoughts, and you may actually never get back on the same track. Working for 12 5-minute chunks separated by 1-minute unimportant interruptions is not the same as working for a 60-minute block followed by 11 minutes of unimportant tasks, even though both involve 60 minutes of work time total. You might get a lot done in those uninterrupted 60 minutes, but get almost nothing done in the interrupted 60 minutes. Reading emails, responding to messages, and similar things, cause repeated interrupts. You need long blocks of uninterrupted time to be productive.

Stop with the Daily News too

Many graduate students are hooked on daily news. They read news website, listen to news, watch news one or more times a day, and try to keep up with subject-specific news sites like Slashdot. However, to be frank, most news is not very important. I realize it sounds almost blasphemous to say this – shouldn't we be concerned about the world we live in? Shouldn't educated people be knowledgeable of current events? Allow me to explain.

News outlets in the 21st century have one goal in mind – to keep your attention. Their goal isn't to inform you, or enlighten you, or make you a better or wiser person. They are simply trying to get you to pay

attention to them, so that they can charge their advertisers more money. Thus, they focus on stories likely to grab your attention, and they make every subject they cover seem urgent and important. We respond with the same sort of attention that we give to a traffic wreck on the side of the road – we just can't seem to look away. News providers appeal to a base side of humans that wants to know what's "going on" around us – a key aspect of survival in the wild. While talk of a pack of lions seen near our cave required our immediate attention, talk of the President trying to raise support for a bill for some particular program does not. Much of the information provided in the news is simply not relevant to us in its specifics, though we might wish to be aware of things in general. For example, it doesn't really matter that there was a specific shooting in some other state, though we might want to be generally aware that such things do happen. Much of the information is totally irrelevant, such as nearly all celebrity news and many scandals, or is only very marginally relevant, such as the media reporting on how the media is covering a particular event. Much of supposed "news" is actually prediction and conjecture, like guessing what the President will say in a speech – why in the world do we need to know this beforehand? Are we going to sell all our stocks or escape from the country before the speech, based on what he plans to say? Furthermore, much of the information need not be heard every day, but rather can be distilled into a weekly or even monthly summary. The day-to-day operations of our government, and even worse, speculations on what our government might do or name-calling going on among government members, is simply not something we need to know today – a summary at the end of the month, which will be freed from speculation since the events already happened, and which will be more encompassing – would suffice. To test this idea, read a news magazine from last year. Many of the articles will be useless to read now because they were mere predictions. Of the other articles, ask yourself how important it would have been for you to know the information at the time

it became news, or if learning the information at a later time and in a summarized form would have been just as good if not better.

It's not just a matter of the time spent reading those news stories. Another problem is that reading those stories every day clutters one's mind. Reading the daily news – fires, killings, fighting, controversy – is an awfully depressing way to start the day. I find that I am generally a happier person if I keep my news intake limited to key sources that summarize, reflect, and go into depth, and if I get my news at specific limited times, usually once a week or even monthly.

My suggestion is to stop reading the daily headlines (and to cut out "local news" almost entirely, which these days is just a few minutes of pointless freak-show stories connected by 20 minutes of commercials and "coming up" announcements). Go online maybe once a week, perhaps to a news magazine rather than a daily news site, or subscribe to a print copy of such a magazine.

When a friend of mine passed away suddenly in 2000, I was so sad that I did not want to hear of any outside news – none of that stuff mattered to me at the time. I ended up going an entire year without watching or reading a single news source – literally. Only the terrorist attacks of September 11, 2001, got me reading news again. But ever since then, I barely pay attention to daily news, though I do read the Economist every month or so, going through a stack of the magazines when it's convenient, and I'll listen to National Public Radio when driving to work to hear a few more in-depth and though-provoking stories. I've survived just fine with this low-quantity approach to news, happier and more relaxed, and with more time for thinking about tasks related to my longer-term goals, plus with more time focusing on my family and friends.

Get Some Sleep

Sleep is important. Findings during the last decade have shown sleep to be important for the brain to learn. Sleep also impacts mood and motivation. It impacts weight gain. It impacts overall health. Make sure you get a good amount of sleep.

The amount of sleep is not the only important aspect of sleep, but the *time* at which you sleep is also important. Try to go to bed and to rise at about the same times every day, and rise such that you are on campus when other students and your advisor are also on campus. This sounds obvious, but because the graduate student's schedule may become quite flexible once the student has finished classes, many a graduate student has embarked on the experiment of simply going to bed when they are tired and rising when they wake up. Such a sleep approach inevitably results in the student sleeping later and later in the morning, and working later and later in the evening, until their schedule is almost completely the opposite of normal people – they come in to work in the afternoon or even evening, and work until the middle of the night or early morning. That schedule has two major drawbacks. First, you aren't on campus with your advisor or with most other students, thus negatively impacting the relationships with those people that are so critical to success. When you do have to show up to campus in the middle of the day, you are tired and not at your best. Second, you may be spending more hours in bed but not getting any better rest than if you were on a normal schedule. Those last hours lying in bed are not a deep regenerative sleep, and aren't very necessary.

While variation in sleep time obviously will occur, strive to keep a schedule. Set your alarm if necessary. You can adjust it on the weekends, but don't sleep in too late on weekends because that will throw off your weekdays.

Chapter 6 – Research

Your advisor is the boss when it comes to how to go about your research. Here, I'll share a few lessons I've learned that may be helpful to some students as well as to some advisors.

Baby Steps

When a child first learns to walk, she takes tiny steps first; she doesn't take long steps, and certainly doesn't run. Those first tiny steps are known as baby steps.

The best way to start doing research is to take a baby step approach. Rather than trying to address a deep challenging problem from the start, a better approach is to address a simpler better-defined problem initially. I like to try to define a clear 1-2 month research study for young graduate students. Nothing complicated. The study should involve a simple idea, be testable using straightforward experiments, and make use of existing infrastructure. Defining such projects is a key task of the advisor, who should know much better than the student what ideas represent a new research contribution and what projects are doable in one to two months. After the study, I like to mostly write the paper for the student, with the student filling in parts of the methodology and experiments sections. This baby step research experience provides the student with a first look at the heart of the research cycle – define a problem, develop a solution, conduct experiments, and write up the results – but without overwhelming him. The student may also be motivated when he sees himself making clear progress, having contributed to a paper with a clear idea and results.

Usually, the early simple problem addressed is related to the bigger problem that I hope the student will eventually address as his Ph.D. topic.

I usually have the student take on larger and larger problems as the years go by, all related to the bigger Ph.D. topic. By the end of the student's Ph.D. studies, the student is tackling tough problems and developing robust creative solutions, and writing papers mostly on her own, in consultation with me. By the end of the student's studies, the Ph.D. student is teaching the advisor as much as the advisor teaches the student, if not more. (A colleague at another school once commented, tongue-in-cheek, that he graduates his Ph.D. students when they know more than him.)

A common mistake is for the student to try to tackle too challenging of a research problem too early in the student's graduate studies. The student has big ideas for solutions, starts to try one solution but discovers it doesn't work or would take too long, starts on another solution only to discover the same, and so on. Years may go by without any successes. Humans need positive reinforcement – small successes – to be motivated to pursue a goal. Without reinforcement, the student loses motivation. The student may start spending time on things that have clear outcomes but aren't related to research, like teaching, video games, sports, pet projects, etc. The student may become depressed.

Sometimes the advisor is to blame for these lack-of-progress situations; the advisor really should have broken down the big problem into smaller ones. Sometimes, though, the student is to blame – the student was too stubborn to take the advisor's advice. A new graduate student thinking he can define a good research problem is much like somebody who just learned the basics of cooking thinking he can create a gourmet meal – better to start with small tasks assigned by an experienced chef.

I've seen two common mistakes by new graduate students when trying to define their own research problem. One is to suggest a research problem that's been studied to death. (If you know Computer Science, I've seen several students, after taking a graduate architecture course, think they've come up with a breakthrough cache replacement policy, which of course has been studied to death for three decades). Researching such a problem is like mining a gold field that was previously mined for a decade, and has long been deserted. At best, you might find tiny scraps of gold that were overlooked by others before you, but even that's highly unlikely. Better to find a new field to mine that hasn't yet been stripped. In research, there are always new directions that require research attention; why not focus your attention there? At least focus it there for a few years, so you'll be more likely to have success and stay motivated as you learn what research is all about. Then, if you really want to address that other problem that everyone else has already studied, and if you still think you have something to contribute there, go for it.

The second mistake new graduate students make is to propose to address a problem that is much too hard. I recall a young student 15 years ago wanting to do his Ph.D. by developing a robot that could roll around a room looking for any object it was instructed to find (a particular book, coffee mug, pen, etc.). A wonderful idea, except that just detecting the edges of an item was hard to do at that time, let alone recognizing that an item was a coffee cup or pen. Decades of research has been done on image recognition by thousands of people, and we are still a longs ways off from that student's proposed robot. A good Ph.D. contribution to solving that bigger problem might solve a very specific sub-problem, like finding the handle of a coffee cup. Sounds trivial, but think about how many different angles from which a coffee cup much might be viewed, how parts of the cup may be obstructed, how patterns on the cup (like a photo) may make the cup look like something else,

how lighting can cast shadows on the cup, how background colors may cause the cup to blend into the background, how cup are of different shapes and sizes, and so on. Just that one problem is very, very hard.

Take baby steps when you first embark on research. Let your advisor define a simple manageable problem. Get some success under your belt to help with motivation, and some experience to build insight and wisdom. The time to take more control over defining the research problem will come; be patient.

Make a Contribution

The goal of research is to make a unique contribution to the state of knowledge in the world. That goal is usually quite different from the goal of most jobs, which may instead be to process items coming in (be it mail, or light bulbs, or elementary school students), or to gain a competitive edge in the marketplace by producing a slightly better product (or just making people think you did). Thinking about the "state of knowledge" in the world is not something most people do. Yet you, the graduate student, must do so. That means you must first become familiar with that state of knowledge, usually by doing lots of reading, talking with your advisor and other professors, attending conferences, searching online literature, and more. Then, you and/or your advisor must come up with a research plan that adds to that state of knowledge – research that makes a contribution.

Sounds simple enough. It's not. I can't possibly explain in this book how to come up with a contribution; I'm not even sure I know how to do so myself.

In thinking about research contributions, I find it helpful to consider the criteria for patenting an invention in the U.S. To be granted a patent, the invention must be deemed to be three things: (1) New, (2) Useful, and (3) Non-obvious to a practitioner.

Similar criteria might be used to evaluate research. When considering if research represents a contribution, we might ask:

(1) Is it new?
(2) Is it useful, is it potentially useful, or could it eventually lead to something useful?
(3) Is it hard?

Regarding (1), obviously research must be new. When you embark on a research problem, spend time surveying the existing literature to see if you can find similar work. Don't just search within the narrow limits of your own research field's conferences and journal. Think about whether the problem might exist in a similar form in another field, and check that field too. You might also find it useful to search online patents, something not typically done in research 10 years ago. With patents now searchable on the web, some researchers find that what they thought was a novel research idea was actually patented years ago by some big company, and never published in the research literature. Searching patents can also give you ideas for new research.

Research differs somewhat from patents with regard to (2) and (3). With regard to (2), research doesn't have to have to be immediately useful. For example, the development of the first transistor was the product of much research, but there wasn't a clear use for transistors back in the mid-1900s. Only after more research showed how to make transistors economically and reliably and how to use them for digital computers did the tremendous usefulness of the transistor become clear many years later.

With regard to (3), research typically should be something more challenging, requiring more cleverness, knowledge, and insight, than something an undergraduate might carry out as a senior project. The graduate student has taken another year or two of courses and done extensive additional reading. The definition of the problem and/or the

solution should be clever. Not just non-obvious, but something requiring breadth and depth of knowledge.

Not every research contribution will excel in all three of the above areas. But they should probably exhibit some amount of all three.

A common mistake is to confuse having done lots of work with having made a research contribution. Doing lots of work is admirable, but it doesn't necessarily advance the state of knowledge.

In defining a research project that makes a contribution, it is sometimes helpful to think of the audience for your research. When you publish your research, how will the papers make things any better for that audience? Will your research catalyze their research, by showing a promising new path (or a dead end)? Will your research educate the community in some useful way? Even better, will your research open up a new direction of further study? Could your research lead to new companies or products that don't exist today? If you can't answer yes, or at least maybe, to any of these questions, what is the point of your research? What is its contribution?

Some researchers believe a good research contribution can be clearly summarized in one or two sentences, meaning just a few dozen words. That contribution summary might indicate what you did that was new and why it is useful. As you embark on new problems, it may be good to take a step back every so often and think carefully about the contribution. Taking that step back may help guide your research to result in real contributions.

You might do this at multiple levels. Do it for the specific problem on which you are presently focusing. Do it also for your Ph.D. thesis topic, where you will compose solutions to numerous sub-problems to create a larger single contribution.

It's not easy. But taking the time to think specifically about your contribution – what are you doing, and why is it useful to others – may help insure your research makes a real contribution.

Meeting with your Advisor

Try to meet regularly with your advisor. Some advisors schedule regular meetings. For those that don't, *you* should seek out meetings with him or her. Don't just quietly work away, waiting for the advisor to check on you. Your advisor may be very busy with multiple tasks, and may need help getting graduate students bumped up on his/her list of priorities. As a general rule, you should probably be meeting with your advisor at least once per week.

If the meeting is going to a technical one, come prepared. Bring your notebook, notes from previous meetings, printouts of data, previous papers – whatever could potentially be relevant. Get in the habit of spending some time preparing for meetings, with your advisor or anyone else, as such preparation makes the time together more fruitful. During the meeting, as your advisor makes suggestions or asks you to do things, write them down – humans aren't very good at remembering a big list of things to do. If a meeting resulted in many action items, I ask my students to send me an email summary.

Ask your advisor questions – pick your advisor's brain on things. "What do you think about this problem?" "How do department faculty meetings run?" "Do you think we should emphasize A or B in this paper?" Some advisors give advice freely; others have much wisdom but need to be asked questions.

I personally have found that informal lunches are a great form of interaction between student and advisor. Plenty of subjects end up getting covered in casual conversation, ranging from topics involving research and administrative matters, to topics having nothing to do with school. Not every student and advisor will be comfortable with such interaction, but when it works, it seems to work wonderfully.

Chapter 7 – Writing a Paper

Writing a good paper has very little to do with writing, and much more to do with two other things: (1) Having a clearly defined contribution, and (2) Organizing the subject for the reader.

Make a Point

Some papers are beautifully written. They use perfect grammar, a delightful mix of sentence structure, clever metaphors, and powerful convincing language. Yet at the end of reading the paper, the reader can't determined if he has learned anything. What happened? The goal of writing is not to make good use of the language – that would be like saying that the goal of driving is to properly hold the steering wheel. No, just as the goal of driving is to get to your destination, the goal of writing is to *make your point understood to the reader*. To achieve this goal, you must yourself understand the point you wish to convey. Not understanding your own point is the single biggest contributor to weak papers.

This sounds strange – why would people have written the paper if they didn't have a clear point? Yet it happens all the time. A student may have been asked to study a particular problem. The student comes up with an idea, develops a framework, and runs experiments. She doesn't find the answer, but notices something unusual, so goes on a tangent to investigate. More development, more experiments. Months go by. Another tangent. More experiments. Now a big conference paper deadline approaches. It's time to submit a paper. She's worked for months, so there must be a paper in there somewhere. She creates a paper

that describes her framework, her experiments, and various results. There are a lot of ideas and impressive amounts of data. But there is no clear point – no clear contribution.

The fundamental mistake is thus to confuse having done lots of work with having made a contribution. A good research paper is not a report of "Here's what I did for the summer." Instead, a good paper asks a clear question on a subject that is important, and provides a clear answer. The important question and its answer represent the *contribution* of the research.

Thus, if you want to write good papers, the first thing you must emphasize is to keep your research focused on making a contribution. You can't bake good bread without good flour, no matter how superbly laid out the recipe. And you can't write a good paper if your research didn't make a clear contribution.

Before getting anywhere near a word processor and laying out glorious words and sentences, think long and hard about the actual contribution of the work. Talk with your advisor, talk with your colleagues. Get a clear idea of the point that your paper will make – you should be able to state your point in just a sentence or two. Then, keep your paper focused on making that point.

Don't Start Writing Yet – Organize Instead

O.K., you've come up with a clear point, a well-defined contribution that will be presented in the paper. Time to start typing away at a word processor, right? Wrong! Words, sentences, and paragraphs are the last stage of writing, and in my opinion, the easiest. The next step after defining your main point is to *organize* how you will present the point.

Organization is really what makes or breaks a paper. To create a paper that makes a clear point, you must decide what information you will provide to the reader, and in what order.

Suppose you wish to write a paper that makes the point that driving very fast is safer than driving the speed limit because driving faster means you spend less time on the road[3]. Compare paper organization A with organization B below.

Organization A

- A1: Establish that driving is dangerous and that improving driving safety is thus an important problem.
- A2: Introduce the hypothesis that the less time you spend on the road, the less likely you are to have an accident.
- A3: Provide experimental data showing that driving fast reduces time on the road.
- A4: Conclude that driving fast therefore decreases the likelihood of having an accident.

Organization B

- B1: Provide data on the numbers of people that drive cars every day, and on the numbers of accidents.
- B2: Discuss the fact that many drivers drive fast.
- B3: Provide data on what percentage of drivers drive above the speed limit, on how fast they drive, and on how driving fast reduces time on the road.
- B4: Conclude that people should drive fast.

Most people would agree that organization A is superior to organization B, and may come to believe the (rather silly) hypothesis.

In organization B, sub-point B1 never actually makes clear that improving driving safety is an important problem and a goal of the research. B2 is leading us somewhere, but where? B3 provides lots of data, but why? B4 finally attempts to make the point, but falls short, leaving it to the reader to go back and figure out how the presented data suggests that people should drive fast.

[3] Please don't do that, of course.

This example may look silly, but many research papers on far less obvious subjects are essentially organized in a way similar to B. The importance of the problem is not clearly stated. The main point of the paper is not clearly defined. Obvious information is provided to the reader, which tends to annoy readers. Subjects are discussed that aren't relevant to the main point, thus misdirecting the reader. Data is provided that doesn't have an obvious purpose. Conclusions are drawn that don't clearly follow from the earlier parts of the paper.

The most likely cause of organization B is that the author did not have a clear understanding of the main point to begin with.

From this simple example, we can derive several guidelines to follow when organizing a paper:

- Clearly define the main point of the paper, summarized in a few sentences early in the paper. Don't write a mystery novel that slowly unveils the subject matter. A research article is more like a news article – a headline, the main fact up front, and then elucidation. There should be no surprise endings in a research paper.
- Make sure that every sub-point has an obvious relationship to the main point. Avoid the temptation to write about a sub-topic just because you know about it or spent time working on it. A key to good writing is to not be attached to presenting everything you know, but rather focusing on what the reader needs to know. Thus, what does *not* appear in a paper is as important as what does appear.
- Know your audience, and don't talk down to them. Your research paper is likely intended for experts in the field. Computer Scientists don't need to read that: "Computers are revolutionizing our world." Chemists don't want to read that: "Chemical advancements are influencing many aspects of modern society." Many young graduate students start papers that way. Think about the target audience, and adjust your writing to them.
- Introduce data as a means for validating a clear hypothesis. The reader should never wonder why the data

is present. The data should provide a clear answer to a clear question.

- Conclusions should follow obviously from the paper. Be as explicit as possible. Connect the dots for the reader; don't assume they are connecting those dots by themselves.

Use Pictures

A paper is designed for a reader other than yourself to understand what you did, not for you to document what you did. Text is great for documenting what you did – it helps you record details in an orderly and complete way. But when describing something new to somebody else, a picture is really worth a thousand words.

Consider explaining a process for saving money on electricity in the home. The process could be explained in words:

> *To reduce household energy consumption, first shut down all lights, computers, appliances, and other electrical equipment that you can think of. Go to the electric meter and record the power per minute – this is the base power consumed by electric devices like the TV, the clock display on a microwave oven, etc. Create a list of the most likely large energy consumers, ordered from highest to lowest. Turn on each likely energy consumer one at a time, recording the power per minute of each, which is the meter power minus the base power. Estimate the time that each item will be active. Multiply that time by the recorded power to obtain energy. For the highest energy item, devise a plan to reduce the item's energy usage (e.g., install a vacancy-sensing switch, buy a more efficient refrigerator, etc.). Do this for each item, from highest to lowest, until satisfied.*

Figure 1: Graphical depiction of a process to reduce household energy consumption.

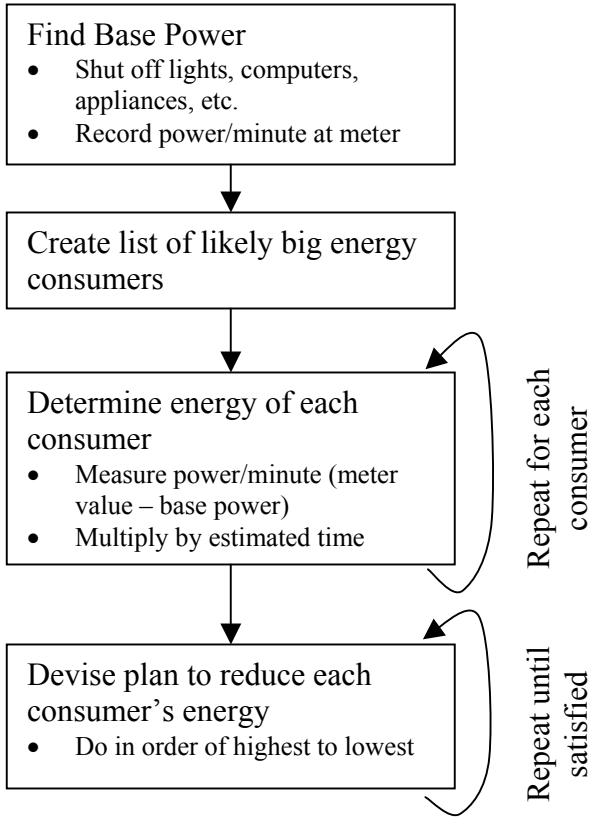

Alternatively, the process could be summarized with a picture, as in Figure 1. The figure would require some additional explaining in the text. However, most people would agree that the figure quickly conveys the main idea of the process more easily to the reader, than does the text-only description. A person skimming the paper (which is the most common way that papers are read) is more likely to glean the process from the figure than the text.

When writing a paper, consider how to make your main point and your sub-points using figures, even though you could describe those items accurately using text. Think of how to present your results

graphically too. Thus, after coming up with the organization of the paper, you should determine the key figures that will be used in the paper.

Refine the Outline

Now that you have defined the organization of the paper, and determined its key figures, it's almost time to write. However, before really writing, you should probably further outline the paper. Lay out the sections, and the main topic of each paragraph in each section. Each paragraph should have a bullet point in the outline. Ideally, each paragraph would even be broken down into a few sub-bullets in the outline.

Your advisor should be involved in reviewing this "draft" of the paper, if he/she hasn't already been involved. Your advisor may want high-level changes made, perhaps asking you to spend more time covering previous work, or having a better idea for illustrating a main point. It will be much easier for you to make changes now than it will be after you've written sentences.

Finally, Write

Finally, you can begin writing sentences. If you've done the proper preparation work, then this last stage of writing should be straightforward. The preparation work can be frustrating, as you may not feel like you are accomplishing as much as if you were to actually write. When you write, you see paragraphs and then pages being completed, and you feel like you are making progress towards completing the paper. But that's not necessarily progress. If you hop in a car and start driving fast without properly planning your route, you may be going fast, but in the wrong direction.

A good analogous situation is that of painting a room in a house. Inexperienced painters want to just start slapping paint on the walls. But experienced painters spend most of their time not painting, but rather

preparing to paint. They sand the walls, clean the walls, and fill in small holes in the walls with spackle. They move furniture away from the walls, cover the furniture with sheets, and cover the floors. They remove electric outlet covers and light switch faceplates, remove lighting fixtures, and remove doorknobs. They meticulously place masking tape along edges. All that preparation work may take longer than the painting itself, and after hours of prep work, not a single drop of paint has yet appeared on the walls. But good prep work leads to a far better paint job.

In addition to preparation, a few rules practiced early can help lead to easier to comprehend papers.

Rule 1: Avoid Vague References

When writing, avoid words that often lead to vague references, like "it," "they," or "this." For example, consider the following sentence:

> *"Shrinking transistors has led to an increase in chip power density, leading to heat problems. This has led to higher costs per system."*

What does the word "This" refer to in the second sentence – "shrinking transistors," "an increase in chip power density," "cooling problems," or something else? Replace the word "This" by the actual item being referred to, such as "*The resulting required cooling equipment* has led to higher costs per system." Such replacement requires the writer to think precisely about what he/she is trying to say, and makes the paper more comprehendible to readers.

The most abused vague reference word is "it." I've completely banned my students from using the word "it" in their papers. The students grumble at first, but most have later said the ban helped.

Once a writer has more experience, words like "it" and similar such references can be introduced such that the items to which those words

refer are always completely clear. Even then, keeping such words to a minimum makes the reader's job easier.

Passive voice often results in vagueness too. If a paper says: "15 benchmarks were developed," there may be ambiguity as to who developed those benchmarks. "We developed 15 benchmarks" is clearer, or if someone else developed those benchmarks, perhaps say "Researchers at XYZ developed 15 benchmarks."

In avoiding passive voice, be careful not to carelessly use "we." In describing a process, don't say "We first ___. We next ___." Rather, say "The first step of the process is to ___. The second step is to ___."

Rule 2: Use Simple Sentences

Use simple sentences as much as possible. A simple sentence consists of a subject, a verb, and then an object or a prepositional phrase. The sentence "There are two approaches to frying vegetables, one of uses a pan, and another uses oil" is a bit long. Even worse is the sentence "Frying vegetables using a pan, and frying vegetables using oil, are two approaches to frying vegetables." One could instead use simpler sentences: "Two approaches exist to fry vegetables. The first uses a pan. The second approach uses oil." Simpler sentences force the writer to think clearly.

Don't write papers the way you talk. For example, when people talk, their sentences often begin with phrases like "There are" or "It is". Avoid such loose language use in papers. Sticking with subject, verb, object/prepositional-phrase sentences helps avoid loose language usage.

Using simple sentences may contradict some of the training you received in writing classes. But we aren't writing poetry here. The ideas are more important than the artistic expression of those ideas.

Experienced writers start each paragraph with a sentence that summarizes the rest of the paragraph. Such a start helps readers organize the upcoming information.

Rule 3: Have Your Paper Read and Proofread

It's hard, but try to write papers early enough so that your advisor and perhaps lab mates have time to read the paper and give substantial feedback, and so that you can revise the paper with this feedback. My students and I have a habit of completing papers several days, sometimes a week, before the deadline. We sometimes end up polishing the paper right up to the deadline. Rare are the late nights before the deadline rushing to finish a paper. It's not easy, but it's definitely possible.

Once you think your paper is polished and ready to submit, have it proofread by somebody else. My students use each other heavily for this purpose, and each gladly helps the others, knowing the others will later help him/her. Having somebody else proofread is a crucial step – a different set of eyes will see things that the author completely missed. Remember, a paper with sloppy English or formatting will immediately bias reviewers to reject the paper.

Reviews

Knowing a bit about how papers are reviewed can help you prepare papers more likely to be accepted, and to deal with the inevitable rejections.

Papers submitted to conferences or journals are typically reviewed, often anonymously, by a program committee or by other people in the field. In an ideal world, the people assigned to review a paper take the time to carefully read the paper in an effort to truly understand the authors' work and point of view.

Alas, we do not live in an ideal world.

Authors should realize that reviewers often do not read the paper entirely. This is not because reviewers are lazy or evil. Reviewers may have dozens of papers to review in a short period of time. Thus, reviewers often seek to determine if a paper can be rejected without a complete reading, in order to finish all their reviews on time. They'll look for easy-to-find shortcomings, such as:

- A low-impact contribution
- The work having been done by others before
- Insufficient experimental results
- Sloppy formatting or English
- A lack of details
- Insufficient references

Reviewers will typically read the abstract and most of the introduction, gloss over the middle pages while looking at the figures, scan the results, and read the conclusions. From that first pass, perceived shortcomings may be found.

Once detecting a shortcoming, the reviewer decides to reject the paper, without ever having read the paper entirely. The reviewer then finds a few other points to criticize in the paper, writes up a quick review, and moves on to the next paper. This may sound unfair, but it's often the reality of things.

Authors can help reduce the likelihood of such an early rejection. A good approach is to read several papers previously published in the target conference or journal, especially best papers. The author can learn about typical lengths of various paper sections, level of detail, previous work coverage approaches, experimental data expectations, number of references, etc. The author can then create a paper consistent with such norms for that forum.

Using plenty of figures helps to convey the main points to reviewers who are skimming the paper.

Staring the contributions explicitly in the abstract, introduction, and conclusions helps the reviewer see the impact of the work. Be sure to be consistent in stating the contributions – I've seen many papers where the abstract lists one contribution, the introduction another, and the conclusions yet another.

Including discussion of previous work shows that the author is an authority, and explicitly distinguishing from that work make clear that the present work is novel. Avoid bashing other peoples' work in order to make yours look good, a common mistake made by young graduate students. Other people do good work too. Just point out what they did, and how you've improved on that, without suggesting that the other peoples' work had shortcomings or problems.

Don't expect the reviewer to give your paper the benefit of the doubt. It is much easier for a reviewer to reject a paper than for a reviewer to accept a paper, because accepting a paper is high risk, while rejecting is low risk. Accepting a paper is high risk because many other people will see an accepted paper, so if it turns out to be a weak paper (already done, insufficient experiments, low impact, etc.), many people will notice. Nobody ever sees a rejected paper, so rejecting is low risk to the reviewer.

The above discussion isn't just relevant to preparing papers, but also to dealing with paper rejection. Don't get too upset by paper rejections, especially by hostile or very negative paper rejections. Reviewers generally don't like to reject papers that are right on the border of accept/reject – it feels bad to reject a paper that is just about as good as papers being accepted. It's much easier to reject papers perceived to be farther from the border. Thus, there's a strong psychological incentive for reviewers to perceive rejected papers as much worse than those papers really are. Reviewers thus tend to provide reviews skewed to the

negative side. Try not to take this personally. I can't tell you how many times we've had papers strongly rejected, only to have nearly the same paper later accepted at another forum, or even at the same forum the next year.

Of course, just because a paper is rejected doesn't mean that the reviewers did a bad job. A quick read of a paper by a reviewer can indeed uncover real shortcomings of a paper, and thus the reviewer comments may be perfectly valid criticisms. Furthermore, many reviewers do spent a lot of time reading an entire paper, and provide thoughtful in-depth reviews of papers.

When your paper is rejected, first give yourself a few days to get over the disappointment of the rejection. Then, read through the reviewers' comments looking for areas of improvement. Especially look for criticism common across two or more reviewers – those are aspects of the paper that almost certainly should be improved. The feedback from reviewers can be a tremendous aid in improving a paper. Make a list of things that you plan to change in the paper, and incorporate those changes when you try again at another time.

Very often, rejected papers eventually get published. Over the past 15 years, my graduate students have a 25-30% success rate for any submitted paper. However, about 90% of their papers have eventually been published. Either the original paper gets improved and accepted to a later forum, or the work is written up combined with other work in a more comprehensive paper. The quality of those later papers is typically higher than the earlier ones. I won't go so far as to say that I'm glad that our papers get rejected initially – I'm not. But rejected papers can eventually lead to better quality papers, so the cloud does have a silver lining.

A paper submitted to a journal is sometimes rejected, but with encouragement to submit a revised version. A revised submission is typically reviewed by the same set of reviewers as the original

submission. This is not really a "rejection," but rather a request to improve the paper before publication. In fact, it is quite rare that a paper submitted to a journal is accepted on its first submission. Many papers go through one, two, or even three revisions before finally being accepted.

When submitting a revision, authors should respond to reviewer comments in a separate document. Each reviewer comment should be listed, with a response to each comment. Young graduate students often make the mistake of defending their paper in those responses. That's not what reviewers want to see. If you try to argue with a reviewer in your response, the reviewer is likely to just say "Fine – the paper is rejected." Reviewers really want to see their comments incorporated into the paper. So try to refine the paper in response to the comments, even if it means just adding a couple sentences to address the comment, and then state in the response what you did. For example, "We've added a paragraph explaining the subject on page 3."

Furthermore, don't give lengthy answers to reviewer questions in the response. Again, the reviewer is really trying to improve the paper, so your answers should be in the paper, not the response. A short answer in the response, with reference to the longer answer in the paper, is appropriate. Or just a reference to the paper itself ("We've addressed this issue with a new paragraph on page 6") is often fine.

Chapter 8 – Giving a Presentation

Graduate students give a variety of presentations during their studies, including teaching a class, summarizing research papers in graduate seminar classes, summarizing their own research for professors and fellow students, passing graduate level exams, and presenting a paper at a conference. Learning to give effective presentations is thus a critical skill.

Being Nervous

Thinking of giving a talk commonly makes the graduate student quite nervous. That's perfectly normal. The best way to overcome this nervousness is to simply give plenty of talks. Each time will get a little bit easier. If you do something frequently enough, you just get tired of being nervous.

You also learn that the outcome is usually just fine – even if something terrible happened during the talk, you learn that life goes on, so you are less scared of terrible things happening in the future. I've had plenty of things go wrong during talks. Once I gave a talk at a big conference – not just a regular talk, but a paid tutorial talk. Somehow the slides being presented were an extremely old version of the slides, totally different than the slides I'd practiced on the past weeks and missing most of the key information. It was like a bad improvisation exercise. The audience was as confused as I was, as each slide that appeared was a surprise to me. I was so upset afterwards. But life went on, and I'm still

standing. My Ph.D. advisor at the time gave me the best advice – "You win some, you lose some." Cliché, but how true it is.

Be aware that even the most seasoned speaker still gets nervous sometimes. New subjects, new audiences, and new situations can all evoke that nervous reaction. A little bit of nervousness is actually a good thing – it keeps you on your toes.

It may also help to think of giving presentations as a growth process. Don't expect that you will be a good presenter from the start. While some people seem to be natural presenters, most of us mortals start off as terrible speakers. We stutter, our throats get scratchy, we completely blank out mid-talk, we don't look at the audience, etc. That's all perfectly normal. Nobody except maybe Mozart starts playing beautiful music the first time they sat down at a piano. Getting good at anything takes time.

So face your nervousness head on. Let yourself be nervous those first few talks, with knowledge that by giving those early talks, you are getting yourself through the hard part.

Knowing the Audience

Perhaps the most common experience of an audience member is being bored stiff by a talk. We've all sat through lectures or research talks where the entire audience is falling asleep, fidgeting in their seats, and watching the clock, while the speaker continues to plow through his prepared presentation. Clearly the speaker didn't intend to inflict such punishment on the audience. What went wrong? Typically, the problem is that the speaker did not understand the audience's needs and desires. If you go to a funeral and tell jokes, people won't laugh, no matter how funny the jokes or polished your delivery.

Thus, perhaps the most important part of giving a presentation is to first understand your audience's needs and desires.

Consider presenting as a teaching assistant to freshmen in a lab section or discussion section. Freshmen typically need concrete information, like drawings, equations, animations, and plenty of examples. They typically desire motivation, namely to know why learning this material is relevant to their future careers. They also desire a bit of entertainment – a few jokes here and there can really win young students over. When teachers of freshmen focus instead on presenting abstract information and just assume the students are already motivated to learn the subject, they usually find the students are bored and inattentive. The information may be great, but it doesn't match the students needs and desired. That same information might be very well received in a graduate seminar, where the students have a deeper background and thus can understand the abstract discussion, and where the students might already be highly motivated to learn the material. To those students, spending time trying to motivate the students may be viewed as condescending.

When presenting research, two common mistakes are going into too much detail, and assuming the audience members are experts in the area. Audience members in a graduate research seminar or at a conference typically want to understand the big picture. They'd like to understand how the problem being discussed fits into the world, and why the solution is relevant and important. They'd like to understand the intuition of the solution. They'd like to see the obtained results, and especially how these results are relevant. Of course, they typically want to see some details of the solution, but here is where many presenters go overboard, giving excruciating details of the solution as if the audience were going to take notes and go home and try to directly reproduce the solution themselves. Those people will read the paper; people go to talks for a more generally perspective. The presenter needs to give just enough detail to convince the audience that he/she really implemented the solution; but no more.

The second common mistake is to assume the audience members are experts. The presenter makes incorrect assumption after incorrect assumption about the audience's knowledge until everybody is completely lost. Of course, if you start with "2+2=4", everyone will also be bored and annoyed, and may think you aren't all that intelligent. Most young presenters are more afraid of being considered unintelligent than they are of being criticized for talking above the audience's knowledge level, so they tend to err towards the latter. To be frank, that's probably the wise direction to err.

Knowing the audience is an art developed by experience. To improve, one has to consciously try. Before each presentation, take time to ask yourself questions about the audience: "What does the audience already know? What does the audience hope to gain from this talk?" During the talk, watch the audience to try to gauge their attentiveness. After the talk, don't just ask your friends or advisor how the talk went – they'll usually say fine. Instead, ask them: "What could I have done to improve the presentation for this audience?" – and don't let them get away without giving you some sort of an answer.

Make a Point

As with writing papers, presentations should have a clear point, and be organized around making that point. A key mistake is for the presenter to look at the entirety of their work, and try to describe that entirety in their presentation. That's a *presenter-centric* view of a presentation – "I did a lot of work, and I want everyone to know it!" But the audience is usually more interested in learning something new. Focus on them, not you – give an *audience-centric* presentation. Of the topic that you are generally covering in this presentation, what specifically do you want your audience to learn? Focus on that specific point in your presentation. Resist the very strong temptation of covering material tangential to your main point, just because that material is something you did or know.

Once you've decided on your main point, tell it to your audience up front. This is not a movie or TV show, designed to take the audience for a thrilling ride or to keep them from changing channels. No suspense, please. Tell them the main point up front, even including it in the title if possible. Give them the necessary details supporting the point, and then conclude by summarizing the main point yet again. As the common advice to speakers goes, "Tell them what you are going to say, say it, and then tell them what you said."

For example, suppose you researched the effects of smoking cigarettes on worker productivity, and you found that smoking decreases worker productivity. Some sample titles, from least directly to most directly to the point, are: "A Study on Smoking," "The Effects of Smoking on Worker Productivity," and "Smoking Decreases Worker Productivity." The first title is awfully vague, yet similar such titles are ubiquitous in various disciplines, like "A Low-Power Design Methodology for Microprocessors" – for people familiar with low power and/or with microprocessor design, that title tells them almost nothing. The second title gives more information, but does not make the main point – it's more of a teaser title. Did productivity go up or down? Stay tuned to find out. The third title tells you the main point right up front. The presentation would then provide the details relevant to this main point – who was studied, what variables were controlled for, and what was the resulting data. The data showing that smoking decreases worker productivity should be the focus; other data, even though the presenter spent time and effort gathering that data, should not be included if it doesn't support the main point. Finally, the paper should conclude that, based on the study, smoking decreases worker productivity by X%.

The above may seem obvious, but I would estimate that only about 10% of presentations by graduate students effectively focus on a main point.

Don't Compete with your Own Slides

Many presentations today involve using PowerPoint (or similar) slides. Making good slides is thus a key component of a good presentation. In making slides, it's helpful to think of how the audience will process the visual and auditory information being provided to them. Consider the following fact – the audience will be trying to listen to you and will also be looking at your slides. This fact brings us to an important rule of preparing slides:

Don't create slides that you'll be competing with.

If you create slides that consist of a list of bulleted text items, and you then speak similar information during your presentation, you have set yourself up to compete for the audience's attention. Should they read the text on the slide, or listen to you? You, the speaker, have set yourself up to be the bad guy – the guy who is talking while people are trying to read the slides.

Instead, the information on your slides and the words you speak should *complement* one another. The most straightforward complementary relationship is a relationship where the slides contain pictures while the speaker explains those pictures. Remember that: "A picture is worth a thousand words." People can process a lot more information by seeing than they can by hearing. In Figure 2, which of the approaches, that in (a) or (b), do you find more effective at making the same point?

Most people would much prefer to attend a talk where the presenter showed the picture in (b) and spoke the information in (a), versus a talk in which the presenter showed (a) and then spoke similar words as in (a) too.

Text does play some role on slides. Text on a slide should mainly be used for two reasons:

- To organize the information for the audience.
- To remind you, the speaker, what to say.

Figure 2: Two ways of presenting information: (a) textually, and (b) graphically.

- Items were defined as widgets produced by a single worker in the factory being studied.
- Average productivity for non-smokers was 200 items per day
- Average productivity for smokers was 160 items per day
- Productivity decrease for smokers was thus 20%

(a)

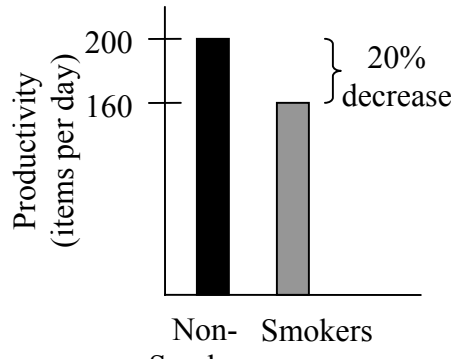

(b)

So if you'll be discussing two types of fruit, apples and oranges, you might have a bullet for "Apples" and one for "Oranges," possibly with sub-bullets briefly summarizing what you'll be discussing ("Crunchy," "Soft," etc.). Other than famous quotations, complete sentences should almost never appear on slides.

When presenting at a conference, put your name at the bottom of the slides. In such a presentation, you want people to remember who you are – you are presenting yourself nearly as much as you are presenting the researching findings.

Audience members will appreciate if you include on each slide the current slide number and the total number of slides, such as "Slide 8 of 20." That way, the audience has some notion of how much longer your talk will last – a key question in many audience members' minds.

Leaving that question unanswered can be somewhat distracting to an audience member, even if they aren't going to change their behavior based on that information. People just like to know.

Engage Your Audience

Countless articles describe how to give good presentations. One common feature in those articles is the important of *engaging your audience*. When you begin your presentation, audience members will be deciding (consciously or unconsciously) whether this will be a talk worth listening to, or whether it is time to tune out. If you start the talk in a boring or incomprehensible way, the audience assumes the rest of the talk will be that way too, and they slip into a more passive mode of listening. If you instead start with something attractive or interesting, the audience perks up and becomes engaged with the speaker.

It is hard to think of a more boring way to start a talk than starting with necessary acknowledgements or a bulleted text outline of the talk. "I'd like to first acknowledge John Smith and Mary Jones for their contributions to this work, and funding from agency XYZ. The outline of my talk is as follows. I'll begin by introducing the subject and providing the necessary background motivation." Not very engaging.

Common ways to engage an audience include telling a story, saying something humorous, asking a provocative question, or giving interesting facts. A classroom of students might be engaged by any of those methods; an experienced teacher uses a variety of methods during a semester. Attendees of a technical conference might especially appreciate interesting facts, such as facts highlighting a problem ("Microprocessor chips have become hot enough to fry an egg on.") or summarizing interesting findings of the speaker's research.

To appreciate the importance of engaging the audience, imagine going on a first date with someone you never met before. If the first words out of that person's mouth at the beginning of the date were: "I

want to start by saying that I am thankful for the fact that I have a car, which I used to drive here, and without which my commute would have been twice as long," what would you think? You'd probably predict that this person is rather boring, and hung up on formalities (and even a little weird). That early prediction would likely influence how you viewed this person for the rest of the evening. Giving a talk is certainly different than going on a date, but first impressions do make a difference, on dates, in talks, and in many other things.

Practice, Practice, Practice

Inexperienced presenters should practice their presentations before the actual presentation. That means setting up a timer, standing up, and then going through the slides one by one and talking on each one. The presenter should practice their presentation many times. For a 10-minute subject introduction by a new TA, a single practice run might suffice. For a 30-minute presentation in a graduate seminar course, 1 or 2 practice runs might be wise. For a 20-minute conference talk, 5 or even 10 practice runs might be appropriate. For conference talks, practicing once in front of fellow graduate students and your advisor (known as a "dry run" talk) is also a good idea. It's a good idea to measure times for each talk segment (e.g., Introduction – 4 minutes, Previous work – 2 minutes, etc.), so that you can later determine mid-talk whether you are running long or short than when you practiced.

Practicing has several advantages. First, it helps identify portions of the presentation needing improvement. For example, if you can't seem to explain a particular slide smoothly, perhaps you need to divide that slide into two slides. Second, it helps determine the actual length of the talk, and where most of the time is being spent – young speakers are often surprised how much longer a talk lasts than they expected. Third, it reduces the level of nervousness experienced by the presenter – the more familiar they are with the talk, the more confident they are that the talk

will go smoothly. Finally, it helps create a presentation that is much more polished and succinct.

As you gain experience in giving presentations, you'll find that practicing becomes less necessary for some talks, such as for a TA presentation. And you may find that fewer practice talks are needed before a conference presentation. However, even seasoned presenters will typically practice their conference talks at least a couple of times before the real thing.

Respecting Allotted Time

Have you ever been to a talk where the person has exceeded his allotted time and the audience really wants the talk to end, but the person keeps showing slide after slide of information that nobody is digesting? There are two common causes of this disconnection between speaker and audience: First, the presenter is thinking more of himself than of others' time, and second, the presenter doesn't know how to shorten the talk.

A speaker should think about other peoples' time. Time is to be respected. If a talk during a conference or seminar is supposed to be a specific length, then going over time means that the schedule will be thrown off, which can be upsetting to the organizers of the meeting and uncomfortable for the attendees. It may mean that subsequent speakers may have to shorten their talks, which isn't very fair to those speakers. For some talks, going over time means that audience members are missing out on activities they had planned for afterwards – perhaps doing some work, lunch with friends, attending another talk or a meeting, or going home to family. Those audience members are then faced with the difficult choice of missing these activities or walking out in the middle of a talk, something people try to avoid out of respect for the speaker. Thus, when a speaker goes over allotted time, it sometimes means that he/she is not thinking about the meeting organizers, other speakers, and audience

members. Instead, the speaker is thinking about himself/herself – "I prepared these slides, and by gosh I'm going to present them!"

Sometimes, as a speaker goes over the allotted time, the problem isn't that the speaker doesn't care about other peoples' time, but rather that the speaker simply doesn't know how to shorten the talk.

Thus, the two rules for respecting allotted time are: (1) Be aware that going over time may be somewhat disrespectful to organizers, to other speakers, and/or to audience members, and (2) Be prepared to shorten the talk on the fly.

Shortening a talk on the fly can be aided by two things. The first aid is the realization that the audience will survive even if you don't cover everything you had hoped to cover. They likely will have learned most of what there was to be learned. The last bits of information that you try to cram in while going over time probably aren't essential, and also won't be digested by an audience staring at the clock and fidgeting in their seats.

The second aid is a bit of thought beforehand about what you could cut out if necessary. Going long is a very common problem in talks, and thus you should expect that it will happen to you. There are many possible causes. People might ask questions during the talk, you might end up taking more time to describe a subject than you had anticipated, or you might go off on a tangent for a while. Whatever the cause, it does happen. Thus, for any presentation, you should look over the last part of it, perhaps the last half or third of the talk, and decide which slides you can skip if necessary.

In fact, experienced presenters often look at this issue from the opposite perspective. They start with *fewer* slides than necessary for the allotted time, with those slides focusing on the main point, and they then insert additional slides to cover the situation of the talk being short. Such an approach is especially useful if the audience is allowed to ask questions during the talk. If there are lots of questions, the presenter

merely skips some number of the optional slides, confident that the main point will still come across. This approach means that the presenter prepares a core talk that is shorter than the allotted time, perhaps two-thirds as long, and then extends it with optional slides.

I once gave a talk as part of a seminar at a major computer company. Before my talk were three other professors also giving talks; mine was the last of the seminar. After the seminar was a separate event at the company in which Stephen Hawking, the world-renowned physicist, would be speaking, and to which everybody would be eagerly rushing off to attend. The other professors each went over their allotted time a bit. I'd been originally asked to talk for 30 minutes, but by the time it was my turn, there were only 15 minutes until the scheduled ending. The organizers and audience kindly told me to proceed with my original 30 minute talk. That was a wonderful gesture on their part. Realizing that nothing I was saying was earth shattering, and not wanting to be remembered as "The Guy Who Prevented Us from Hearing Stephen Hawking," I covered key slides only, and gave only the key points on those. I finished the talk in 10 minutes. People seemed very grateful.

Stop Apologizing

Sometimes a presentation is given as one big apology. The presenter keeps telling us what he did not have time to include, what animation was supposed to appear but didn't work on this computer, etc. The audience leaves feeling cheated. When giving a presentation, remember that:

> *Nobody misses what they didn't know was supposed to be included*

I once went to a magic show. I was impressed by the magic tricks, and really enjoyed the show. It was quite well done. But then at the end

of the show, the magician said: “Our final trick is normally to have a lion appear right here in the middle of the audience. But the lion is sick today.” Up until that comment, everyone was having a great time. After that comment, everyone was wondering what that trick would have been like, and there was a clear feeling of being cheated out of something.

If your talk is missing some things you’d hoped to include, remember that only *you* know that. Rather than creating a presentation where the audience feels cheated, instead focus on what information you *do* have in your presentation, and do the best job you can to present that information.

Chapter 9 – Travel

Graduate students disseminate their research by publishing papers. A common publication forum is a conference, symposium, workshop, or similar meeting. If the paper is accepted, one of the paper's authors must travel to the conference and give a talk that summarizes the research to the conference attendees. Graduate students may also make presentations to funding agencies on behalf of or along with their advisor, may give talks at companies, and may attend an occasional seminar or similar informational meeting. Thus, graduate students may make several long-distance trips during their studies. Because graduate students are often new to such traveling, I'll offer some tips.

Get a Passport Now

When you start graduate school, do yourself a favor and get a passport if you don't already have one, long before you actually need one. Many graduate students wait until they get they get a paper accepted in another country before looking into getting a passport. However, the passport application process can take months. Expediting the process is possible for an extra (hefty) fee, but even then, one is never guaranteed to get the passport in the week or two that is promised – a small mistake, by you or the passport processing agency, can hold everything up. Remember, passports are issued by the government, and the government is not known for providing particularly good customer service. Thus, apply for a passport long before you need one. That will be one less thing to have to worry about, and to possibly go wrong, when you do eventually have to travel internationally to a conference.

There are only a couple of passport offices in the U.S., and most people don't live near one. Thus, most people use a passport processing service in their city. Although you can download a passport application from the web, you cannot simply mail the application to the passport office – you must go to a physical office that offers passport processing services. Many (but not all) U.S. post offices offer passport processing services, where you can turn in your application, have your identification checked, and pay the fees. The post office will then forward your application packet to a passport office, who will then mail the issued passport directly back to you.

Something that I've learned over the years is this – avoid government offices, like the post office or the Department of Motor Vehicles, whenever possible. Going to those places is like going back in time 30 years – long lines, rude attendants, poor service, sterile settings, etc. The alternative to the DMV is an AAA office, if you are a member (and if you drive, joining AAA is a great idea), which can take care of many matters. By the way, if you do have to go to the DMV for any reason, make an appointment if possible – you'll avoid the long lines. As an alternative to the post office for passports, some public libraries, and even some universities, offer passport processing services. Check the web or phone book, your university's web pages, or ask around. When you do have to go to take care of government matters, bring every document you can think of that might be relevant. I can't tell you how many times I've had to return just because some document (like a passport, birth certificate, proof of residency, etc.), which was never listed as required on the forms, was required.

Make a List

Some of the stress of travel involves wondering if you are forgetting something, and of course actually forgetting something important. I have

found that making a checklist, which you maintain and revise from trip to trip, helps.

Besides the usual things to bring, I recommend bringing two alarm clocks (I use my watch as one of them), and setting both of them at your destination – you may sleep through the first one, set it wrong (pm instead of am), its batteries may die, etc. Bring business cards, and hand them out – going to conferences and meetings is as much about networking as it is about the sessions and formal information. A commonly forgotten item is the power chord for a laptop. If bringing your presentation on a flash drive, you might consider keeping a backup in a separate piece of luggage, and even putting a copy temporarily on your web page just in case. Bring a photocopy of your passport, kept separately from your actual passport. You might want to keep a credit card separately from your purse or wallet also, in case you lose the purse or wallet. If checking in your luggage, bring essentials in your carry-on bag that would get you through a day in case your luggage arrives a day late (which happens more often than you might think). Because meals are not be regular when traveling, I like to bring nutritional snack food, such as cereal bars (also know as power bars or granola bars), nuts, and/or trail mix (a mix of nuts, raisins, dried fruit). I also find that having a bottle of water on long flights is helpful. Bring comfortable shoes, as you'll likely be doing a lot of walking. Bring something to read or do for when you are waiting in terminals or while flying.

Travel Funds

Traveling is important, but also expensive. Even if your advisor is paying, he/she will likely appreciate if you can find some travel support. Check the conference website for information on student financial support. Sometimes participating in a special conference aspect, such as a Ph.D. forum, might result in free registration or a bit of travel money – again, check the website carefully. If you don't find anything on the

website, try sending an email to the conference organizer asking if there is financial support for students. Conferences are often sponsored by professional organizations, and sometimes those sponsors have travel grants for students attending sponsored conferences, so check those organizations' websites too. If you are supported by a fellowship, check if your fellowship includes additional support for travel, as many do. Finally, your own university may have small grant programs to encourage students to present papers; ask other graduate students, your advisor, other professors, staff, or others.

Many travel grant programs impose very early deadlines (e.g., 2 months or more before the conference) as a way to reduce the number of applicants – by the time most students think about applying, they have already missed the deadline. So start your search early, as soon as you determine that you will be attending.

The key is to be aggressive in searching for money. There is money out there, but nobody is going to knock on your door and offer it to you. You have to seek it out. Some advisors have plenty of money, so looking for travel grants may not be necessary. But most advisors will deeply appreciate your efforts.

Flying

Your advisor will likely appreciate if you can find a low fare airline ticket. However, convenience (or more accurately, freedom from severe inconvenience) is also important when flying, so the low fare should be considered in the context of the total cost of the trip – saving $100 by taking flights with a terrible schedule is not worth the savings if the hotel and conference registration costs total $1000. In addition to using the online travel search engines to find good fares and schedules, consider using a travel agent that works with consolidators. A consolidator is a company that buys tickets from airlines in bulk at deeply discounted rates, and then sells the tickets for prices typically lower than the lowest

published airline fare. These are perfectly legitimate tickets; the only problem is that if you have to cancel due to illness or something similar, an airline will typically give you credit for a future flight if you bought the ticket from them or through a normal travel agent, but if you bought from a consolidator, you may be out of luck.

If taking a flight with a connection, avoid flying through cold-weather states, like Denver or Colorado, in the winter; they are more likely to have delays due to bad weather. Try to minimize the number of stops, even if it means slightly higher price – each stop is another chance for a delay or a missed connection. I always try to find a reasonably-priced non-stop flight first, before looking at flights with connections. If you are flying using federal funds (like National Science Foundation funds), be aware that you may have to fly on a U.S. airline (or at least on a flight that is code shared with a U.S. airline) to get reimbursed for the flight.

If flying within the U.S., the major travel sites are a great place to start (e.g., expedia.com, travelocity.com, or orbitz.com). However, be aware that certain important airlines, like Southwest Airlines and JetBlue, do not list themselves on those sites. You need to go directly to their sites (southwest.com and jetblue.com). I love flying Southwest, not only because they are so cheap, but because they have very flexible change and cancellation policies. They also have one of the best on-time success rates of any airline. Check them out.

When you book your ticket, you can usually select a seat. I like to choose a seat on the aisle, so I can get up whenever I please, to use the bathroom or just to stretch my legs– walking can reduce the fatigue caused by a long trip. Sitting in non-aisle seats means you'll have to disturb the person next to you to get up. Of course, sitting on the aisle means you may be disturbed by other people yourself.

Many people don't realize that they can order special meals, such as low salt meals, low cholesterol meals, vegetarian meals, lactose free

meals, gluten free meals, kosher meals, and much more. Not only can this help ensure you get a meal that you enjoy (or perhaps more accurately, that you can tolerate), but another small perk is that they usually bring the special meals out first. Special meals must be ordered in advance, perhaps 48 hours before the flight leaves, or more. You can usually order the meal right when you buy the ticket.

At the airport, I like to avoid checking bags, in order to save time at the airport both when checking in and when arriving, and to minimize the chances of a lost bag. But even if you have carry-on bags only, be prepared to check one of them in, as overhead bins often fill up and you may be required to check in your larger bag.[4]

Don't try to time everything such that you minimize time spent waiting in the airport. Delays are common – traffic, long airline counter lines, long security lines, etc. Why sit around at home waiting to leave when you can leave early and sit around at the airport instead to be safe? Bring plenty of reading material – books you've wanted to read, magazines to catch up on, work material, etc. If you have things you want to read and enjoy reading, time spent waiting goes quickly and is actually fruitful. An unexpected delay thus has a small bright side; yes, you'll be late, but you may finish that book you've been reading.

If new to a city and your hotel doesn't offer a shuttle, it's often safest to take a taxi to your hotel, even though a taxi is more expensive than bus or train. That way you avoid walking through bad parts of town, trying to negotiate trains, buses, stairs, escalators, and so on with heavy bags, or wandering around with heavy luggage trying to find the hotel

[4] I was once on a flight where a man boarding the aircraft refused to check his bag when instructed. He said "I'm not leaving this bag" and defiantly boarded the aircraft, and made space for his bag in an overhead bin. About 10 minutes later, a police officer arrived and escorted the man (and his bag) off the plane. Many non-U.S. residents, coming from countries that are more relaxed about rules, don't believe me when I tell this story. America is a much more rule-oriented place than many other countries.

(and perhaps getting lost). By the time you have to return to the airport, you may have a good enough feel for the area and transportation system to take a bus or train back to the airport.

Hotels – Safety First

If a conference is held in a hotel, and you are not familiar with the area, it is usually best to stay in the conference hotel if the price is only moderately higher than surrounding hotels. Unfortunately, the price is often much higher. In this case, sometimes the conference has suggested alternate hotels; if they don't, some students will ask a conference organizer. If you do search for a hotel yourself using an online travel engine, think about safety. The lowest-priced hotel may be low-priced for a reason – it may be in a bad part of town. A friend of mine took the best deal on a hotel in Paris, and got mugged outside the hotel. I once reserved a good-priced major-brand hotel in San Francisco for a 5-night stay, only to check out after the first night because the area I had to walk through to get to the conference was so dangerous. I ended up having to stay 10 miles away and take a train because all hotels were sold out by then. In short, when thinking about hotels, don't just think about cost – think about safety and convenience too.

Reducing Jet Lag – Avoid the "Nap of Death"

When flying to far off countries, adjusting to the new time zone can be difficult. Many people say that flying from the U.S. to Europe, for example, is very difficult, and many articles say flying east is harder than flying west (the latter just lengthens the day). Being out of synch with the local time is known as "jet lag."

I've learned to minimize the difficulty of flying to Europe by flying such that I arrive in the afternoon or early evening (which usually means leaving the U.S. in the evening), rather than arriving in the morning. If you leave the U.S. in the evening, after a few hours of flying, you may

doze off for a few hours. You certainly don't get enough sleep, but you then arrive in Europe in the afternoon or evening – just enough time to get some dinner, and go to bed around 10:00 p.m. You'll probably wake up at 3:00 a.m. and not be able to go back to sleep – I think that has something to do with your body temperature cycle still being on U.S. time. But you are at least making progress on becoming adjusted.

In contrast, if you fly such that you arrive to Europe in the morning, you are going to have a hard time staying awake on that first day. You didn't get a good sleep on the plane, and now it's 9:00 a.m. and you are ready for sleep because it's midnight (or later) back home. That first day will be one of the longest days of your life.

In either case, do *not* take a nap. You will most certainly be tempted to take an afternoon nap on your first day. Do NOT take that nap! Take a brisk walk, swim in the hotel pool, go out for coffee or ice cream – do anything you can to stay awake. Why? Because that nap is not a normal afternoon nap. It is the nap to end all naps, the mother of all naps. A former student of mine called it the "Nap of Death." You will enter a deep, profound sleep that usually begins your 7-8 hour nighttime sleep cycle, and waking up after an hour or two will be very difficult – you'll feel terrible. You might even sleep through your alarm and end up sleeping 4-5 hours or more. In either case, you won't sleep as good later that night, and you'll have the strong urge to sleep again in the middle of the next day. Avoiding napping will help you adjust to the new time zone much more quickly.

Adjusting to a new time zone takes time because your body maintains certain 24-hour rhythms, relating to body temperature, kidney function (your kidneys slow at night to let you sleep), stomach acid production (you stomach produces less acid at night to protect its lining), etc. You cannot change those rhythms instantly; it can take many days or even a week or more. Experienced travelers try to adapt to the local schedule as soon as possible. That not only means to avoid sleeping in

the day, but also waking at the proper hour, and eating meals on schedule too. Sunlight plays a significant role in body rhythms. Open the shades to your hotel room so that your body sees the sun when it rises. When hit with a wave of sleepiness, get outdoors in the sun to tell your body that it's the middle of the day and not time to sleep.

Network

One of the main reasons for attending a conference is to network, i.e., to get to know other people in your field. Knowing people will be helpful when you are looking for a job – those people may be involved in hiring, or may be able to serve as a reference. If you stay a part of the research community after you graduate, they will be people you see regularly at conferences, who you appoint or who appoint you to serve on program committees, who review your research proposals to the government, who write reference letters for you, who you collaborate with on joint research projects, who send you good graduate students and vice-versa, who invite you to give talks, who fund your research, who hire your students, and so on.

When you talk with people at a conference, you never know which person will eventually become a long-term contact. Maybe 1 out of every 20 people ends up being someone you remember and with whom you maintain a relationship of some sort. But if you hadn't talked with those 20 people, you wouldn't have met that 1, so get out there and talk.

One easy time to talk with people is after their presentations. Presenters appreciate questions about their work. You can usually walk up to presenters and talk with them after their session is over. Some conferences allow attendees to formally ask questions immediately after a presentation, in front of the audience – that's a different type of interaction entirely. If you do ask questions during that formal question time, be careful – too many hot-shot young graduate students ask hostile

questions that seemingly debunk the presenter's researcher. Such questions really aren't fair, as its easy to ask such questions in 10 seconds ("Wasn't this work already done by Smith?"), but answers may require many minutes to carefully explain. Furthermore, the presenter has to think of an answer quickly, in front of a big audience, while asking the question is easy. I recommend to my students to never publicly attack someone else's work, either during formal questions, or in their own presentations and papers. Spirited discussion certainly has its place, but young students should likely err on the side of having such discussions in one-on-one or small group settings rather than publicly at conferences.

Another time to talk with people is during meals. Look for people you'd like to get to know, and try to sit at their table. You might ask your advisor before the conference if there are people you should meet.

Conferences typically have numerous breaks for socializing. To a young graduate student, these breaks can be very frustrating. Everybody seems to be talking with somebody else, and the graduate student feels like he/she is the only person not talking with someone. That's a normal feeling. Rest assured that things get easier every year. Little by little, you meet people, and so you'll recognize more people at each subsequent conference. As you talk to those people, you'll be introduced to people that they know, and thus you'll get to know even more people. And so on. Keep trying – it gets easier. If your advisor is at the conference, stand by him occassionally, and see if he introduces you to colleagues.

If traveling with fellow graduate students, resist the temptation to spend all your break time and meals with them. That's certainly the most comfortable thing to do, but you really need to get out and meet new people, as you'll see those fellow students back in the lab at home every day. Of course, bonding with your fellow students is also an important part of traveling to such conferences. The key is to find the right balance.

Chapter 10 – Preparing for an Academic Job

Some graduate students plan to seek out a career as a professor. Those students should pay special attention to a few things during graduate school, to strengthen their resume (known as a Curriculum Vitae, or CV) and thus to be a strong candidate for a faculty position.

Publishing

Academic-bound graduate students should be sure to publish in numerous conference and journal forums, especially high quality forums, throughout their entire graduate studies. While people in academia like to say that they don't just count papers when considering potential candidates, and instead say that they try to look at the impact of the candidate's research, the fact is that the number of published papers in quality forums is still an extremely important aspect of a candidate's CV. This counting is in part due to the difficulty of determining the impact of a paper just a year or two after being published, and also because people reviewing your faculty application usually aren't experts in your field.

Your advisor might not be driving you to publish a lot, especially if your advisor already has tenure, so you might have to take the initiative to publish papers. I've seen many cases of students having a hard time getting academic jobs due largely to the advisor not having actively encouraged publishing. You might wish to start with publications in conferences/journals that are a bit easier to get accepted into, rather than top-tier conferences/journals, to give yourself some early success and some practice.

Publishing an outrageous number of papers usually isn't necessary – beyond some number (perhaps 10-20, depending on the field), there is some notion of diminishing returns – if the student is publishing so many papers, perhaps each paper is just an incremental improvement. But, I've seen many faculty dazzled by very high paper counts too. However, most professors do seem to look for a reasonable number of publications specifically in forums known for being high quality in a particular field.

Teaching

Teaching is an important part of any professor's job. Whether a person will eventually be a solid teacher is thus something that is also considered during faculty hiring. TA experience helps, and if evaluations were strong, they should be highlighted on the CV, including even some quotes from student evaluation forms. Even better than TA experience is experience serving as the instructor of a course. You might therefore seek out such an opportunity in your department, especially in the last year or two of your graduate studies. Teaching a summer course is a common way of getting instructor experience.

Networking

While networking is important for all graduate students, networking can be especially important for academic-bound students. Try to get to know faculty at different schools, as they may be involved in hiring, or at least have influence, when you are applying for a position. Ideally, try to get a letter of reference from at least one professor from a different university, especially a well-known professor. To accomplish that, you would minimally want that professor to have read your work, discussed it with you, and gotten to know you a bit through conversations. Those things can be accomplished at conferences. Preferably, you might have even done a collaborative research project involving that professor, or more

likely with that professor's student. It's really more your advisor's job to set up such collaborations, but as a graduate student, what you can do is keep your eyes open for such opportunities considered by your advisor, and jump at them when you get a chance.

Proposals

The ability to obtain grants seems to be more important today than ever before in academia, as universities place increasing importance on well funded research programs. Having some experience writing grant proposals is thus a plus for a faculty candidate. Tell your advisor that you want to help writing the next proposal your advisor works on. Ask your advisor about funding – To whom does your advisor apply for money, where does he/she get money, from which places is money easier or harder to get, etc. Not only does a good CV list some proposal writing experience, but also a good candidate can speak competently about funding during an interview, having some ideas of where he/she will be applying for funds as a new professor.

The CV

Unlike a resume for a job, which typically should be only one page, a longer CV is usually better than a shorter one. As one professor told me when I was preparing my CV for graduation, "Pump it up!" List educational background, work experience, awards and honors, publications, talks, teaching experience, tutorials, company visits, paper reviewing experience, etc.

List your references (don't say "available upon request."), and make sure you pick people who will write very strong letters. In academia, a one-page pleasant letter is considered a weak letter ("damning with faint praise"). A strong letter needs to be multiple pages and should be *glowing*. It's OK to ask a professor: "Would you be able to write a very

strong recommendation letter for me?" Letters from professors, especially well-known professors, are generally preferred over letters from industry managers or colleagues, though 1 or 2 of those can be helpful too.

Applying

Be sure to apply to plenty of schools. You never know which ones have openings in your field, at your level, for your gender, etc. Don't restrict yourself to only those schools listing openings in your specific area – sometimes the hiring needs change after the ads have been posted, or hiring outside the posted areas is considered due to not finding anyone within the posted areas. Depending on your background, it's not uncommon to apply to 30-50 schools but only get a few interviews.

If your advisor doesn't do the following on his/her own, ask your advisor to email or call people he/she knows (or knows of) in the departments to which you are applying. Keep a neat list of such departments, and provide that list to your advisor. Those emails and calls can make a *huge* difference. Departments get hundreds of applications for a single faculty position, and the recruiting committee may overlook good candidates. An email or call, stressing how great a candidate is, often plays a key role in bringing a candidate to the attention of the recruiting committee.

Chapter 11 – Conclusions

Graduate school can be a fun and stimulating time of life, with the added reward of leading to a higher paying and more satisfying job. Most undergraduates don't realize this. That's why I think graduate school is the best-kept secret in America, and why I wrote this book – to help students improve their experience.

The Key to Success is to Tolerate Failure

However, a danger in writing a book like this is to imply that if you do all the things suggested in the book, you *will* have success. That's not true – you can increase your chances of success, but nothing you do can *guarantee* that you will do well in all your courses, find a good advisor, make great research progress, publish lots of papers, complete your degree, and ultimately find a great job. There are too many external factors beyond your control. Perhaps the most important thing I've learned from my experiences is that:

The key to success is to tolerate failure.

No matter how hard you try, you may encounter some courses that you do terribly in, or choose an advisor that ends up being rather unhelpful, or get numerous conference paper submissions rejected, or experience similar such failures. "Failing" at things is perfectly normal. Some people don't handle failure very well – they lose their motivation, get depressed, or just quit trying. It's quite understandable – we're all just human. However, if you can learn to tolerate some amount of failure, you may encounter successes too, and ultimately the successes may

compensate for the failures. Ideally, you can also learn from the failures, thus improving your chances of future successes.

For example, you might do fine in most courses but do terribly in a particular course. Well, you tried – see if there are any lessons to learn, and then move on (retaking the course if necessary, perhaps waiting for some time so you get more knowledge, or waiting for a different instructor). You might have a few strongly negative comments from student evaluations about your TAing. Well, you are still learning, you can't please everybody, and even award-winning teachers still get some negative comments – determine if there are any changes you might make next time in response to the comments, and move on. You might end up with a less-than-ideal advisor, and kick yourself believing that you should have known better. Well, nobody can predict everything, so don't beat yourself up over it – just deal with the situation as best as you can.

Perhaps the most common failure for Ph.D. students is getting papers rejected. You worked so hard on the research, spent a week or more writing the paper, and waited for months, only to have the paper turned down – you feel like you've wasted your time. Furthermore, those reviews were really hostile, making you feel extra bad and doubting the quality of your work. But getting papers rejected is perfectly normal; even professors at top universities and having hundreds of previously published publications get plenty of papers rejected. Give yourself a few days to get over the rejection. Then see if any of the reviewer comments can help you improve the paper for a later submission, and/or can help you write better papers in general. Be aware that there are many random factors in the paper review process; the same paper could be either loved or hated by the very same conference or journal, depending on which handful of reviewers happened to be assigned to review that paper.

This ability to tolerate failure – learn what you can learn, and move on – is absolutely critical to success in graduate school, and in a career too.

A Final Thought

People come first. Getting good grades, publishing papers, strengthening a CV, getting a good job, and eventually moving up in your career, are all laudable goals. But all those things may not mean much if there are no good people in your life. Pay careful attention to the social side of your life, and seek out friends. Don't argue over little things – most things really aren't worth arguing over. If you have family and friends in your life, nurture those relationships. Be good to those people, help them when they are down, and seek out their help when you are down. Success in graduate school, and success in your career, may be much more meaningful and satisfying when you have good people in your life.

Chapter 12 – Notes

It's a good idea to get in the habit of writing things down as you read. You'll probably want to write down ideas that you get, things that you might want to do, and perhaps summaries of important points. You can use the space below and the next blank pages for that purpose while reading this book, or keep a notebook with you while you read.