My Undergraduate Research Experience and Decision to Pursue a Ph.D.

Frank Wang

Introduction

I would like to start by talking a little bit about my academic background to put this article into context. I received my B.S in Computer Science at Stanford, and I wrote this article the summer before I started my Ph.D. at MIT. I want to emphasize that this is *not* an editorial about whether computer science students should pursue a Ph.D., but rather, I describe my research experiences as a undergraduate and my motivations for pursuing a Ph.D. I also talk briefly about the projects that I worked on. If you want to learn more about these research projects, you can find more details along with my contact information on my home page `http://mit.edu/frankw/www/`.

Let me start by discussing one of many decisions I have made in the past four years: writing this article. It was a random Sunday night in August 2012, two weeks before I had to head to Cambridge, Massachusetts where I will be starting my Ph.D. at MIT. I am anxiously flipping through my recent emails from CSAIL (MIT's computer science department) to understand what I need to do before I arrive on campus to start my journey toward a Ph.D. At the same time, I am worried about the amount of work I have to do in the remainder of my internship at Google. I am surfing the web and come across Philip Guo's *The Ph.D. Grind*, and I start reading. This is clearly an unproductive use of my time (among other things) when I have much to do. However, his memoir as well as his blog made me think about detailing my experiences. Much has been said about the Ph.D. student, but very little has been said about the undergraduate student and the choices he or she has after graduation. I thought I would write down some thoughts regarding my time at Stanford. The article is organized in similar way to *The Ph.D. Grind* where I go through each year of my undergraduate experience. Well, here goes nothing.

Freshman Year: The Beginnings

I came to Stanford in Fall 2008 as the typical college student: confused but excited. I honestly didn't have a worry in the world as the college experience was hyped up by my older friends from high school. College was meant to be fun, exciting, and intellectually stimulating. For the first two quarters, I did not do anything but take classes and meet people. For once, I did not have to worry about doing activities or aiming for some end goal like getting into college. I did do some research in high school, but nothing really significant. I was an Intel Semifinalist, doing research with a professor where he guided me most of the way through. The research was not even in computer science. It was in mechanical engineering where I looked at different oxygenates and how they affected the amount of soot and NO_x it produced. I really enjoyed the experience of doing research, but it was nothing I was really serious about because I had no clue about my major.

Two and a half quarters later, I had some direction, which became the first big decision of my undergraduate career. I decided to go more in-depth into computer science to see if I liked it. This was pretty normal for many Stanford students as we did not have to declare our major until the beginning of junior year, and switching between majors after that is just a matter of filing the right paperwork and fulfilling the requirements. To place things in perspective, I came into Stanford partially wanting to go into finance or law. I had done some computer science in high school, but like most things at that point in my life, I wasn't particularly attached or passionate about it.

I tried to advertise my basic computer science skills to companies for an internship, but by March, I was left without an internship or any plan for the summer. Then, one of the students I shared a research lab with for my Intel project contacted me. I had been in somewhat consistent contact with him since we had gone our separate ways two years ago after working in the same computer lab for about 3 months. He had decided to come to Stanford for his Ph.D. in mechanical engineering. He told me that he needed an undergraduate research assistant because he needed someone who could write computer code at a relatively fast pace. Without an internship, I decided it might be fun to work for him for the summer as well as improve my programming skills.

So, the summer after my freshman year started. Initially, I was kind of bored as research is sometimes slow moving, and I knew nothing about the field I was working in – nanoparticles and electrocatalysis. I had no clue what I was doing, and I showed up the first day and spent most of the day in meetings trying to learn more about the project. I asked for papers to read and background information. Three weeks passed with very little progress with my writing many lines of code that didn't really have any use but was meant to familiarize myself with the code. On the fifth week, we had a small breakthrough. Having a casual conversation with a postdoc in another lab, he gave us the idea for a new method. For the next 3 weeks, we went around talking about our idea as well as trying to find applications for it. We had a method with no application. That is the nature of research. Most of the time is spent

looking for a question to which no one knows the answer. The summer concluded with some progress made, but we still had no idea about the relevance of the result or what question we were trying to answer.

Sophomore Year: Slump and Changes

Many people talk about the sophomore slump at Stanford. What it really means is that the freshman excitement is gone, and people have settled into the reality of college where we have to decide on a major and start taking more advanced classes, which means more work and less time for fun. This is the year that many people have to make initial decisions on majors. I continued to take more computer science classes along with other required classes to fulfill the General Education Requirements (GERs). On the side, I was still doing research for the lab because the professor and the graduate student were interested in seeing the research go somewhere. We also received a department grant to continue the research, and in the first month of the quarter, we submitted two abstracts to this big chemistry conference, and I was going to be one of the speakers. This was pretty exciting for me, and everything seemed to be going in our favor. At this point, I started to consider doing a Ph.D. in an unknown field.

This is where the slump started to happen. It's one of those slumps that there's not much to say in retrospect. Some might say the slump I am going to describe is ordinary or a normal part of research or that research is slow and this was a normal part of the solving difficult questions. Having done more projects later on, I feel that what happened during this year and the early part of the next year dragged on way too long, and it made me really question whether I wanted to do research in the future (let alone pursue Ph.D.). However, it happens to everyone doing research, and you learn from the experience. I have purposely not described the details of the research project because at that point, we did not really know what the project was. However, I'll talk a bit about what we knew at this point. To give a brief overview of the research at this point, we had about 2000 lines of code that allowed us to create random shaped nanoparticles to simulate electrocatalysis like in a fuel cell, and more amazingly, we could simulate them being supported on some form of metal like in a real fuel cell. At that point, no one could do such a simulation. Our vague goal was to find a good model to represent catalysis on nanoparticles as well as supported nanoparticles that matched trends seen in experiments, but our approach to the project changed drastically as time went on.

In March, I ended up going to my first conference where I met with many of the top scientists in the field. I also gave my first talk, but the conference was huge and not many of the talks received attention except ones given by the top scientists. The problem with the conference was that it was too broad and encompassed too many fields. This is very unlike the conferences in computer science, which are very specific and are tailored to specific communities and disciplines.

It was a busy year, and the graduate student and I had not made much progress in publishing any of our results because we were both adjusting to taking more advanced classes. I will discuss my desire to publish frequently because that is the only real way to judge performance because your ideas are heavily scrutinized by experts in the field for correctness and quality. Having talked to many of my friends, this desire to publish is particularly true in computer science and mechanical engineering, but it is not necessarily true for other fields. We gained new momentum after this conference, but like many projects, we just weren't motivated enough.

I have to say I grew kind of tired of this research, and I decided I needed to move on to computer science, which I started to enjoy more and more. I was just a sophomore and didn't know what I wanted to do. I applied for a few internships and received a few offers, but none that really interested me. I looked up computer science professors for two reasons. Firstly, I needed an advisor for my major, and secondly, I wanted to try out computer science research. I really had little hope of getting response from professors because they are known to be busy, but I thought I would give it a try. I found a professor in an area I thought I might be interested in: cryptography and security. I emailed my future undergraduate advisor Dan Boneh and received a rather prompt response from him. Not only did he agree to be my academic advisor, but he offered me opportunities to do research in his lab.

I went and met with him, and he gave me numerous ideas for projects I could work on in the summer. I said I was fine doing anything so that I could get started because I had no idea what research in cryptography or security was like. He also suggested that I take his research class in cryptography. I told him I was worried, especially since I had not taken any class on cryptography, but he reassured me that I would be fine. I decided last minute to take an applied number theory class that had some cryptography to give me some basics for Dan's class.

I must say taking his class was a great and tough experience. The class's format was basically one to two research papers a day, and Dan would lecture on the ideas that he finds important. The class was very interesting, but I was lost many times because of my lack of background knowledge. However, surprisingly with time, the concepts started to sink in. I wouldn't say I understood a majority of it, but I doubt most people in the class did. I think it was fair to say that I absorbed about 30 percent of the concepts, which I reviewed in the summer to the point where I understood around more than 60 percent. The class was good because it wasn't about problem-solving like a normal class. It was understanding something very novel and the motivation behind coming up with that novel idea.

I think it's cliche to say that this is the summer that changed everything, but it was definitely a turning point in my undergraduate career. I started off the summer with a project where we tried to create an authenticated airplane communication system. At the time, airplane communication with ground location systems known as distance measuring equipment (DMEs) was not authenticated, meaning any signal can claim to be a DME without the airplane's computer knowing. Dan was very patient with me as I worked through

the cryptography, and I would say that this project is where I learned most of my basic cryptography. Whenever I was confused, I would read more about the concept in a textbook or online. I really developed a passion for research at this point because I really felt like I was learning and understanding concepts more deeply, which is the point of research. Many times in research, you only use a little bit of what you learn in class, but it allows you to apply it to real situations and understand much better. Sometimes, you have to learn something completely new from scratch. As I like to describe it, it challenges you to be mentally versatile and push yourself because research is about pushing the limits of knowledge.

This was also the summer where I started to commit myself to more projects. During the middle of summer, I took on a new project where I looked at *location based privacy*. The basic problem we were trying to solve was the following: Alice can test if she is close to Bob without either party revealing any other information about their location. Although I did not participate in the main parts of the research because it was pretty far along, Kina, an undergraduate from UCLA, and I implemented it on the Android mobile phone platform. We were in charge of the whole process from design of interface to the systems details. I feel like from every project I learn one lesson. From this project, I learned that there is a huge gap between cryptography in theory and in practice. It took Kina and me basically the whole summer to implement this. I enjoyed the projects that I did this summer, and at this point, I was pretty convinced that I wanted to do research. As far as I know, a person could not do advanced research in industry or academia without a Ph.D. degree, so I started to think more seriously about going to graduate school.

Junior Year: Adjustments

With a new passion for research, I started my junior year. I knew this would be an important year because it was the last full year remaining before I applied for Ph.D. programs. However, I had no idea what I wanted to do. I was in a somewhat unique position to do an undergraduate thesis during my junior year because of units I had carried over from taking more rigorous classes during my first 2 years of undergraduate as well as from high school. Many times, research just becomes a second priority to classes and activities. At that time, I didn't know why, but now, I realized that research is difficult and sometimes a bit undirected. There really isn't any goal, and the project evolves over time, which makes it really difficult in an undergraduate setting where everything is very structured like classes. Also, research is frustrating at times because sometimes, you just feel like you make no progress and really have no definitive end goal. In fact, forcing yourself to make an end goal like a conference deadline or the desire to graduate actually forces you to make progress. At Stanford, the undergraduate thesis counts as about 1-2 classes based on how much time you wanted to spend on research. I knew that if I wanted any chance at a top tier graduate school, I had to focus on making progress in research. This also meant that I had to take fewer classes to have time to do research, which was another good way to make myself spend more time on

research.

I wanted an end goal and something to push me to make progress, so I decided to write up a proposal to do an undergraduate thesis in my junior year. It was ambitious because I had not taken many advanced computer science classes, and I just started to do research in computer science without knowing too much about security or cryptography. For the first 3 months, I spent time discussing with Dan different project ideas. Dan was relatively busy, and research wasn't really high on my priority list. However, one thing research did force me to do was to take classes that would help my fundamentals in computer science. Specifically, at Stanford, we had various tracks, but more or less, the same subset of classes fulfilled all the requirements because of the large number of electives available to us. Basically, we could define our own undergraduate curriculum. I chose to take systems and theory classes that would help me with research in security and cryptography. This was great because I never really struggled with class selection indecision one week before the quarter began like many students. I looked at the available class offerings and chose classes based both on interest and research necessity. Many people will say that classes sometimes help with research, but most of the time, it is better to learn the material on your own. I agree to a certain extent, but at this point in my research career, I had no idea how to do research and did not even have the basic computer science fundamentals.

Finally, after a few discussions, we decided on a rather ambitious topic. Many security operations are being offloaded to the graphics processing unit (GPU) because the GPU has many more threads than a CPU, which makes it perfect for security operations because they were mostly computations. However, a GPU was designed for graphics use not for security, so we were wondering if the offloading of security sensitive data is actually secure. The question was pretty broad but straightforward. I tried for months to find design flaws in the GPU architecture and the languages surrounding it. To cut a long story short, I didn't make significant progress by the thesis due date but learned a lot in the process.

The GPU project was not the only project I worked on during junior year. In fact, there will be a theme of my working on multiple projects simultaneously. I recommend that all researchers, regardless how early or late in their career, work on at least 3 projects at any given time. Ideally, each one would be in a different stage of progress. I feel like everyone has his or her magic number. I usually think 4-5 is ideal (tailored to the number of conferences in a year), but it really depends on how many you can handle. If one project fails, there are backup plans so that they are making progress in research instead of working on one project for a year and have nothing to show for it at the end. I was working on two other projects on the side. I was trying to publish my work on fuel cells and nanoparticles, but we kept on having philosophical arguments with the professor over the best way to write the paper. It was a frustrating project, and many times, we just lost hope that it will make any progress. We kept on delaying this project until we had a better idea what to do with it. We finally decided to submit the paper and see what happened. We would not hear back from the journal until the summer. The other project I was working on was in preparation for my summer project on securing medical records effectively using advanced cryptographic

techniques. I was trying to learn more about the background and progress of the project. This project was with Professor John Mitchell, who I would continue to work with later on. I wanted the experience of working with a different professor and getting a different perspective on different research in the field.

I think this is a good place to take a quick detour before I talk about my summer. I have to explain the difference in publishing research in computer science compared to other engineering fields. Computer science research is usually published in conferences. There are many conferences during a year, but there are a few (usually 3-4) top ones. Usually, you submit a paper to a conference by a certain deadline, and you get reviews back. Sometimes, you get to respond to these reviews without including new work in case the reviewers misunderstood parts of your paper. Other times, you receive an accept, accept but fix some parts, or reject. The program committee, an assembled group of top scientists in the field, decide which papers are accepted or rejected. At top conferences, usually under 20 percent of submitted papers are accepted, so there is stiff competition and your research has to be very top notch to get in. If your paper is rejected, you can submit it to a less competitive conference or go back to the drawing board. Sometimes, it takes a couple of conferences before the work is accepted.

In other fields, like mechanical engineering, you submit a paper to a journal at any time and get a response, which varies based on journal. Most of the time, it involves some revision. You revise and respond to the reviewers' comments. Then, these revisions and comments are returned to the reviewers. Then, you will get a second round of responses. Usually, by the second response, you will only have to make minor revisions, which do not go back to the reviewers, and the editor makes a final judgment call on whether the paper can be accepted or not. It usually takes a few months after the paper is accepted before it shows up in the journal. This whole process usually takes around 6-8 months. Like conferences, journals are also based on quality. The main differences I would say between journals and conferences are the lack of deadlines for journals and the ability to make changes to your paper as a result of reviewer feedback.

The summer was more of the same. I received an intern offer at Microsoft, which paid pretty well, but I turned it down to focus on research in preparation for graduate school. I went to research talks, and I worked on the 3 projects I talked about above but with focus on the securing medical records project. The basic project was that the HIPAA law is difficult to enforce, so we wanted to build a system that would do this automatically using advanced cryptographic techniques. This work eventually was published with me as one of the authors at a conference. I basically finished my project on GPUs by finding a couple of flaws, but none were significant enough to write a paper. Our paper on computational methodologies to analyze nanoparticles finally was accepted (published in 2012) after a couple rounds of tough review. This gave us momentum to start our next paper where we used this method to analyze various nanoparticles. This paper I did on the side mostly in the month of August. I will talk more about the progress of this paper in the next section. My summer was basically routine, but senior year was a relatively chaotic but strong year for me.

Senior Year: Wrapping Up and New Beginnings

My last year at Stanford was pretty chaotic, and little did I know that it was my last year here. I was always pretty adamant about doing my Ph.D., but a set of events caused me to change my mind. I was busy finishing projects. I wanted to do research but still leave time to hang out with my friends because many of them were graduating and moving away. To start off, I took the bare minimum number of classes because I had fulfilled all my major requirements by the end of junior year, and I was done with all my requirements by the end of fall quarter. This gave me time to focus on my social life and research, which were my two top priorities. Of course, there was applying to graduate school, but that was part of research.

I feel like I could fill the next few pages just talking about all the research projects I did senior year. I think it is just best to highlight some research that I did and talk about progress made on previous research. If you want to learn more about the projects that I worked on, you can find them on my home page `http://mit.edu/frankw/www/`. I really want to focus on applying to graduate school and visiting those schools.

I would say that my senior year was my most productive year where I took on and helped on many projects. After much struggle, we finally published two journal articles on the nanoparticle research. One paper was on the computational methods, and another paper was on the optimal particle size for catalysis. These two papers both made it to top journals, so that was the end of that project from sophomore year. Research is sometimes slow, so it's good to work on multiple projects. My summer work on HIPAA enforcement was also published in a conference. I helped out my labmates and other undergraduates on smaller projects that did not lead to papers. I also worked on the key exchange for a new type of Tor, which is an anonymity network designed to skirt censorship and traffic monitoring that exist in certain countries. This resulted in a paper that was accepted at CCS, a big security conference every year. By the middle of the summer, I also submitted a paper, which was joint work with John Mitchell and Google, about security audit tools and developers in an attempt to understand where web vulnerabilities come from. This paper was submitted to NDSS, a big security conference in February. At the time of my writing this, I am still waiting for a decision on this paper. As you can see, I had a pretty productive year in terms of research. Doing research causes you to generate more followup ideas, so you will hear many people say that you can always do more research. The hard part is finding the correct stopping point to publish. In my case, it's good to have many areas to follow up on because I am just starting my Ph.D. My senior year prepared me pretty well to start thinking about important topics for my Ph.D.

That was my research, so I want to talk a little about the graduate school process. There are many articles or posts about the graduate school process, but I wanted to give you my personal account on it. I want to disclaim that I will not talk about the graduate school process in general but rather more specifically about schools that I applied and was accepted.

I feel that it's hard to compare graduate school applications to undergraduate applications. Every application is very similar, so you end up having to only write 1-2 essays. The problem is that every application is "similar," and not the same. There is no common application. You end up filling out your address, contact information, school information, etc. multiple times. The good thing is you can track your application online (most importantly, if your professors have submitted their recommendations). I applied to 6 schools: Stanford, MIT, CMU, Princeton, UC Berkeley, and UCSD. My top choices were Stanford, MIT, and CMU in that order. Honestly, I wanted to only apply to those schools, but Dan advised me against it. It's just a few more dollars to apply. The application consisted of basic information, a research statement, transcripts, and three recommendations (Berkeley required a personal statement on top of that). It was pretty simple and seemed less painful than undergraduate applications where I had to write a series of 5 essays and try to apply them to different schools.

On top of graduate applications, I did fellowship applications so that I can be more independent during my Ph.D. I applied to NSF, NDSEG, and Hertz, but they weren't my main focus. I was more focused on getting into graduate school and figuring out funding later. Either way, I won't talk more about these after this because I didn't receive any, and any information I can provide is available on the Internet or their sites.

After applying to the schools, I played the waiting game. Decisions come out earlier than undergraduate admissions, but they come out unexpectedly. They don't specify a date, and they come out in waves. Usually, if you don't hear from a school by the end of January, you are probably going to get a rejection letter. However, there is no set rule. Some schools like Stanford send out a decision (acceptance or rejection) on a set date. Some schools like MIT send out an initial notice for acceptances in January, and rejections come out in March. It varies from school to school. My first acceptance came from CMU in the beginning of January. Next was MIT and Princeton. I was not accepted to Stanford, Berkeley, or UCSD.

Every acceptance letter comes with an invitation to visit the school with expenses up to 500 dollars covered. It generally isn't an issue because the visit days are coordinated more or less, so the schools will split the airfare. The allowance is more than enough (at least in my experience). I highly recommend going to the visit days. They treat you very well in terms of hotel, food, etc. All the visit days involve talking to graduate students and professors in your area that you usually don't have a chance to talk to. They schedule personal meetings to discuss projects and research. It is a great experience, and it gives you a great feel regarding the school and graduate culture. I was not very keen on switching coasts because I had grown up in California my whole life, but I finally chose MIT for a variety of reasons. My reasons might be different from your reasons for choosing a graduate school. For me, I really liked the lab I was going to work in. The lab came highly recommended from Dan and John, the two professors I worked with at Stanford. Also, one of my labmates from Stanford was also in this lab and she said she really enjoyed it. Also, MIT places a lot of emphasis on their graduate students, and Boston was a great city with many colleges. I thought I would give it a try because I knew I was not bound to MIT or even to graduate school.

Another small detail I wanted to mention was funding. If you don't get an external fellowship, that's generally fine. I received an MIT fellowship for my first year, which just means the school funds me instead of the professor I do research with. This makes my research more flexible and not tied to a specific grant. Funding methods vary by school, so it is important to ask when you go to a school, specifically the graduate students. From what I hear and know, the top schools, such as MIT, Berkeley, CMU, and Stanford, tend to have a better stream of funding because they have larger research budgets.

My senior year was very busy and productive in my mind. I traveled a lot and learned a lot. It was hard to balance everything, so I recommend that if you are applying to graduate school, schedule your classes and commitments accordingly. If you don't, you're just going to have a bad time.

Conclusion

Well this concludes a look into undergraduate research at Stanford. At Stanford, it was easy to be tempted by industry and startups. I would say keep an open mind and consider doing research even if it's for one quarter or one semester. I truly learned a lot from research. However, pursuing a Ph.D. is a big commitment and should be taken very seriously.

I am currently doing my Ph.D under Nickolai Zeldovich at MIT. I am always open to questions and comments about my research, graduate school, and this article. My contact information and more research information is available on my website `http://mit.edu/frankw/www/`.

Acknowledgement

I would like to thank Philip Guo for taking time to discuss this article during my internship at Google and for giving me useful comments on the article.