Debmalya Panigrahi

I arrived at MIT in the late summer of 2007 to pursue a PhD in computer science. I was looking forward to the next few years with a lot of excitement and a little trepidation. I had never lived outside India, had very limited research experience, and was largely ignorant of the cultural and academic traditions of the university and the country that I found myself in. Often, the first task in a PhD program is to decide a research area and an adviser. In my case, these were relatively easy. I had a fair idea of what I wanted to do when I was applying to the PhD program. I wanted to work in graph algorithms, having previously dabbled in it under the tutelage of Ramesh Hariharan at IISc. David Karger was one of the preeminent experts in the field, and he was to be my adviser.

Looking back at the five years that I spent as a PhD student, I can now see three distinct phases: my first year was spent mostly doing coursework with a smattering of research; the next two years were a bit of a struggle as I tried to broaden my horizons and slowly establish an independent research agenda; and the last two years were the most productive and enjoyable as the effort put in earlier started paying off, eventually transitioning me to an independent research career. Let me elaborate on these phases below.

The first problem that I worked on was suggested by my adviser. I was quite inexperienced in research at the time but the problem was well-aligned with other problems I had looked at previously. Luckily, I managed to solve it. (Sadly, that was the only problem suggested by David that I managed to solve in the next five years!) An early result is always a tremendous boost to one's confidence, and I was indeed energized by this small but important success. The paper was written and submitted toward the end of my first year, and the copious comments and suggestions that I got from David went a long way in improving my technical writing. As an aside, I want to note that technical presentation is one of the primary skills that one picks up in graduate school. This includes both technical writing and oral presentation. These are relatively easy skills to acquire in that they require conscious effort and sincerity but relatively less intrinsic ability. This also extends to conversational skills in individual meetings and small groups. The importance of these soft skills cannot be over-emphasized, particularly in the context of employment opportunities after graduation.

The next two years were low on research output. I tried to be more independent and chose problems that would broaden my horizon, but progress was slow. This was partly because I was simultaneously trying to acquire the necessary all-round skills that would make me a more complete researcher. Graduate school is the ideal time for exploration of areas outside immediate research interests, and to build a repertoire of technical skills in these areas. For instance, as a researcher in algorithms, I felt that I was being limited by the fact that I did not possess a deep understanding of mathematical programming and polyhedral mathematics. To overcome this shortcoming, I spent deliberate effort in acquiring better knowledge in these areas. This I did by reading papers, discussing basic tricks and techniques with other students and colleagues, taking relevant courses, etc. I undertook and executed similar plans for some other areas as well, where I felt my existing knowledge and intuitive grasp were not sufficient. This came at a cost: progress in research became slow and output was sporadic. I must say that I was lucky to have David's support at this stage as I tried to carve a niche for myself. In fact, he always stressed that it is important that I develop a taste for problems independent of his preferences. In hindsight, I feel that the struggle of this middle phase was crucial; it taught me to be tenacious and persevering, qualities that are essential as a researcher. I also learned to retain focus on important problems even when they seemed difficult. It is important to remember that one is, and should be, judged by one's best works. A small set of important results is typically accorded much more value than a large body of indifferent work.

The last two years were my best in graduate school. While I continued to learn to do research, I also actually started doing research. I was fortunate to have a tremendous group of collaborators besides my adviser and a very supportive and vibrant research environment in the theory group at MIT. The regular talks, the informal discussions with friends and colleagues on research as well as life in general, the opportunity to interact with a steady stream of distinguished visitors, and my own research visits contributed to a holistic experience that I will treasure for years to come. At this stage, I felt I could choose and judge problems myself, and had quite a few ideas of my own that I wanted to pursue. Once again, David was perfect as a sounding board as I explored these ideas, continuing to give crucial feedback on particular projects, research papers, practice talks, and the larger questions of academic life. As a graduate student, it is important that one identifies one's own strengths and weaknesses, and develops an identity that distinguishes her in the research community. Driving one's own research agenda in the last phase of the PhD program is key to establishing this identity.

Before I end, I would like a say a word about my internship experiences. I spent all my four summers on the west coast, working at Microsoft Research three times and at Google Research once. Since summer internships are relatively short, they present a good opportunity for slightly shifting focus from trying to learn the research methodology to actually conducting high-quality research in a short amount of time. My internships were quite productive and helped enrich me as a researcher, besides providing an opportunity for collaborations many of which continue to this day.

Mythili Vutukuru

My PhD journey began when I got admitted into MIT for their combined MS/PhD program starting Fall 2004. I was in my final year B. Tech. at IIT Madras when I got the admit; I had decided to go to grad school right after undergrad. My choice of advisor was pretty straight forward. Prof. Hari Balakrishnan, one of the faculty members at MIT who was working in the field I had applied in (networking) expressed interest to work with me. And I was more than happy to join his group, having heard great things about him already. I did not have a well-defined thesis problem to begin with when I joined. As a result, my first few years in grad school were quite unstructured. Initially, I started working with some of the senior grad students in my research group, on projects that were mainly driven by them. I worked on a few problems on my own during this time, none of which progressed into a full-fledged thesis topic for various reasons. I also did an internship at International Computer Science Institute (ICSI) at Berkeley with Prof. Vern Paxson after my second year. I ended up working in many research areas (Internet routing, network architecture, network security, and internet measurements) during my first 3 years in grad school, none of which had any connection with my final thesis. However, I value those years greatly, and I strongly believe there could have been no better way to learn how to do research. Working with some of the senior students at the beginning of my PhD helped me in many ways. I got to see and learn first-hand many aspects like how to do research, how to organize your work, and how to write papers. During this exploratory phase, I also read many papers and thought a lot about what kind of research I like and what I don't (as part of trying to find a thesis topic), which helped me become a more mature researcher.

By the end of my third year, I started to worry a bit about graduating and started looking seriously for good problems that would lead to a thesis. I got interested in a problem in wireless networks that one of the senior students was working on. I started working with him initially, and on my own after he graduated. I did not explicitly try to define a thesis topic in the beginning, but only focused on solving interesting problems and publishing them. Once I had a few good publications that were closely related, my advisor and I spent some time thinking and identifying a common unifying theme across the different projects, which eventually lead to my thesis topic. Therefore, in some sense, my thesis was built in a bottom-up fashion, rather than in a top-down manner. Once we defined the topic, my last year or so was spent in working on the smaller pieces that would connect and complete the work in the existing publications, and on writing up the thesis.

Overall, getting a PhD is a challenging experience, both technically and personally. On the technical front, one has to deal with working on new topics without much guidance. For example, some of my research required working with software defined radios, and learning a lot of new communications and signal processing concepts on my own. I did not have any background from courses in these areas, and I had to spend a lot of time trying to find good sources to learn from. This type of learning is very different from the structured education that happens in courses, where the material and syllabus is very clearly defined. On the personal side, a PhD journey adds a lot of uncertainty to your life, and one cannot help asking questions like “Will I ever graduate?”, “Will I ever find a thesis topic?”, “What if someone else has already solved this problem?” and so on. In spite of these challenges, most people remember their PhD days very fondly, and many people help in making your PhD a very fun experience (or not!). The most important person in your PhD experience is your advisor, and your compatibility with him/her. My advisor’s style worked very well for me and I enjoyed working with him very much. Next come the environment in the lab, the work culture, and your relationships with other students in your research group. Finally, and most importantly, one needs to have a good support system of family and friends outside the lab to help you keep your sanity during the (many) lows.

Looking back now, I felt I could have done a few things differently. For example, I wish I had done many more internships in a variety of places, to gain more exposure and breadth. However, these minor issues aside, I think I had a great time during my PhD and I wish the PhD students reading this all the best in their journeys as well!

Ashok Anand

I pursued my PhD in systems and networking under the guidance of Prof. Aditya Akella at the University of Wisconsin, Madison. My PhD journey was quite exciting and it gives me great pleasure to write about it.

Let me start from the beginning. Initially, I was quite confused and uncertain about the area that I wanted to work on. I was generally interested in algorithms and systems, but was unclear on the specific area. During the first two semesters, I explored what I really wanted to work on. Wisconsin’s course requirements really helped me in making this decision. We had a small research project in our courses, which gave me some exposure to think deeply in certain topics. Apart from courses, I also worked on a research project under Prof. Akella in my first semester, and that also helped me. In addition, we used to read papers related to certain topics in our courses, so we also learnt the skill of reading papers, which was quite useful.

My first research project under Prof. Aditya Akella really got me excited. It was on the idea of storing packets on routers and exploring its implications. By having packet caches on routers, we can suppress duplicate content at the packet level, i.e., perform redundancy elimination (RE). We also asked whether we could design better routing algorithms if there were packet caches on routers. I had taken courses on linear programming and algorithms earlier, and that helped me in formalizing the routing algorithm problem as a linear program and solving it. We showed that redesign of routing algorithms can provide significant benefits. This work got published in SIGCOMM 2008. I think that *learning general algorithmic techniques is quite useful as they can be applied across range of domains.*

After this research project, I continued to work with Prof. Akella. Furthermore, I took various courses in networking, and my interest in networking and systems grew. I decided to pursue a PhD in networking and systems. The next big thing was how to go about deciding the problem space and thesis area. I was really intrigued by the idea of packet caches on routers, so I started to think more on the problems in this area.

One of the main questions was to understand how much benefits redundancy elimination techniques can provide in the real world. Fortunately, we had access to packet traces of our University access link, so we could analyze some real traffic. However, studying a single University access link did not seem sufficient. In summer 2008, I did my internship at Microsoft Research India (MSR India), and then I got access to multiple real traces from enterprises. My internship experience also gave me a fresh perspective on thinking about deploying redundancy elimination (RE) not only on in-network routers, but also on end-host systems. Furthermore, interaction with several researchers at MSR India helped nurture my thought process. In my summer internship work, we performed a comprehensive measurement study of packet traces across multiple enterprise locations and established the benefits of redundancy elimination techniques. This work got published in SIGMETRICS 2009. *I think that a summer internship gives you a great opportunity to get a view of the real-word significance of problems* and it worked out quite well in my case.

After internship, I continued my thinking on the problem space of redundancy elimination. It was clear to me that there are at least two sub-problems to solve: a) one was to investigate how to realize the benefits of redundancy elimination technique in settings where routers run at very high speed and b) other was to explore the design space of applying redundancy elimination on end-hosts. I decided to focus on the first problem.

Applying RE in a hop-by-hop fashion at high speed was very challenging, as core routers run at very high speed while we could not perform RE at such high speed (10 Gbps). *Many times, thinking about a problem by looking at the big picture, gives you a fresh perspective and helps in coming up with a novel solution.* In this case, the main goal was to reduce bandwidth on all links by applying RE. So, instead of looking at the narrow problem of improving speed of hop-by-hop RE, I considered the high-level goal of reducing bandwidth on all links and it came naturally that I could think beyond a hop-by-hop solution. I came up with a novel design of network-wide RE, where RE operations can be distributed across routers. In particular, edge routers can perform slow encoding operation of RE, and core routers can perform fast decoding operations in distributed fashion. We showed that such an architecture can lead to effective and practical deployment of RE. This work (SmartRE) got published in SIGCOMM 2009.

Then I started looking at other sub-problems in RE problem space. Applying RE on end-hosts had interesting challenges and had different implications (e.g., last-mile bandwidth savings for mobile devices). I continued to work with my mentor (Dr. Ramachandran Ramjee) at MSR India and Prof. Akella on this problem while being at University of Wisconsin. Another sub-problem was on cache management for RE and I started to consider not only using memory but also fast flash SSDs. *It helps a lot to be aware of recent trends*, and during that time, flash SSDs was getting a lot of traction. By leveraging SSDs, I designed a new index structures for caches, which could give superior lookup and insert performance compared to traditional disk-based indexes. Both of these works got published in NSDI 2010.

In this manner, RE problem space became my dissertation topic. The next big thing was to combine these various pieces of works, and it came naturally as these sub-problems were part of the big research agenda of realizing redundancy elimination as a primitive and exploring its implications.

Apart from RE work, I also explored other problem spaces. I worked on traffic engineering problems in data centers with other students in my advisor's group. During my summer internship at Microsoft Research, Redmond, 2009, I worked on the problem of traffic engineering in inter-datacenter setting. During my summer internship at MSR India, 2010, I worked on virtual machine migration problem. In summer 2011, I worked at Google on high-speed packet processing. While these problems were not directly tied to my research topic, they gave me a good perspective of various research problems and help me broaden my views. *I also feel that during PhD, it is useful to work in diverse topics as it helps build breadth and perspective*.

During my PhD, I worked with various students in my advisor's group on these research topics. I think that working in a team was really helpful for me. Various ideas came out during discussion with team members. In addition, we used to have group meetings where we read papers and discussed ideas. These activities helped in broadening my perspective.

I also had a great relationship with my advisor. He was very encouraging and supportive. Despite his busy schedule, he used to make himself available whenever I needed him. I was also quite enthusiastic about working with him. *I think it helps a lot to build a good relationship with your advisor*. In my case, I was really fortunate to have Prof. Aditya Akella as my advisor.

The PhD journey was not completely rosy. I also had my moments of paper rejections during this period. The paper rejections can be discouraging at times, but *the key thing is to consider the feedback in a positive and constructive way, have patience and continue doing your work*. My advisor was really helpful during these tough times.

I think that PhD journey is a great learning experience, and I feel great to have done this. At the same time, it is important to keep few things in mind to make this journey a good one. I strongly encourage students to pursue PhD in topics of their interest and wish them all the best in their endeavors.

Youngki Lee

Doing a PhD is a strong commitment. It is often a long journey of more than five years, and may decide the whole direction of your life. Especially, doing a PhD in systems research requires even more time and effort, which frustrates many students who started PhD with ebullient enthusiasm. I agreed to write this short essay to share a part of my PhD experiences, hoping that this could be a small help to the students who are dedicating their youth to grow as a great researcher.

I did my PhD at KAIST, Korea and recently received my PhD in mobile systems area. Just like any other people who completed PhD, my PhD journey was full of frustration and failure with just a little successes and achievements, which all together affected my thoughts, capabilities, and attitudes as a researcher. In this short essay, I don't intend to share detailed tips on how to write good papers, do good presentation, and do good research (many tips are already available, and I can hardly do better since I myself am still learning), I would like to share more philosophical thoughts that I believe is important to keep in mind throughout the course of doing PhD, especially in systems research area. I might be biased based on what I have experienced, so just take this as one opinion.

While I have been researching for years, I constantly asked myself questions about what values and abilities are essential to be a leading researcher. From my experiences so far, the most important aspect is in understanding hidden needs of people and society, possibly of the future from imagination, and providing novel systems solution to such future needs. More important, such a creative view and curiosity to the world should be ever-lasting as society gets more complicated and people's needs become more subtle. I do believe building new and novel systems for hidden needs is the most valuable research (compared to improving existing systems). I call such futuristic systems as "experimental systems” from a perspective that their usefulness is still questionable and the completeness of the systems is likely to be premature in many aspects.

To be successful in building experimental systems, I think following three values are critical to keep in mind all the time: *creativity*, *experimentation*, and *collaboration*. In many times, innovation comes from creativity. Such creativity needs be substantiated and supported by real implementation and experiments. Also, the whole process of innovative work is enabled by active communication and close collaboration.

First of all, I believe creativity is a core value to play a leading role as a researcher. Early and creative works are often very significant and influential, but challenging due to the fuzzy and uncertain nature of the problem spaces. Also, it is extremely difficult to deliver the usefulness and importance of the work in the form of the paper and acquire satisfactory academic depth. This is why huge portion of students are solving well-established problems, from which they can acquire stronger academic depth more easily. I have worked on several futuristic research topics, and got frustrated many times since I did not see visible progress for such fuzzy problems. However, I did realize that such challenges can be overcome by accumulating lessons and knowhow about how to do creative research. In terms of creativity, it may not be a good sign if you are simply following popular topics and improving existing techniques under the well- defined problem scope.

Second, extensive experiences on implementation and experiments are indispensable to develop creative ideas into real systems. While it is important to understand the theoretical foundation for the principles and concepts, it is equally or more important to provide experimental understanding on the design and implementation. Thus, it is very important to have strong implementation skills – if you don’t have such skill yet, it is never late; practice for fast and accurate coding skills. With strong implementation skills, it is possible to more deeply and thoughtfully design a novel system and related-technology, and quickly show its potential value through a working prototype. It is very important that creativity does not stop at the imagination stage.

Third, creative systems and implementation need to be strongly supported by active collaboration. Performing good systems research and building a creative system are often a very complicated process involving multi-lateral challenges. Even highly talented researchers could hardly succeed in the works on complicated systems without high level of collaboration skills. Above all, collaboration is a critical factor to bring out creativity. Many creative ideas often come from active discussion with colleagues who have various viewpoints and ideas. I am sometimes very surprised that a great idea is very easily generated even by a short discussion, especially in case to talk to people from other disciplines. In addition, collaboration can accelerate the realization of software systems. It often requires too much time and effort to develop a complicated system alone, while losing a good opportunity to lead the issue in the meantime. Throughout my Ph.D studies, I almost all the time worked with other colleagues and professors collaboratively, having daily discussions and implementing complex systems as a team. It will be great if you could expose yourself to an atmosphere for active collaboration, and seek for many opportunities to work as a team.

Another important axis to remember is international quality and sensitivity. This includes an ability to communicate with international community (through presentation and writing). Although I have studied at one of the leading universities in Asia, I myself have struggled to learn international style of communication due to huge differences in culture and language, and also, to improve the quality of work in perspective of international standards. For the students who have continued to study in the similar local research environment, it is very important to feel and fill the gap if there is any.

Beyond learning all the above capabilities (creativity, implementation ability, collaboration skills and global sensitivity), you should be happy and proud of what you are doing during the whole research process. It is very important to be fond of and enthusiastic about your research, and keep positive energy. In this way, you can learn genuine pleasure of research and have positive attitude toward your future life and career. Also, this will serve as basis to continuously generate meaningful and quality results to the research community as well as society. These attitudes and abilities, once combined with in-depth knowledge about specific fields, will enable you to play a leading role as a researcher in any top global research institutes. I hope that some of you will see the value of building creative software systems from this short writing, and make breakthrough throughout your PhD studies and your future career.

VIJAY GABALE

If I have to describe my experience of pursuing a PhD in one sentence, then it won't be an exaggeration if I say that it was a process of discovering myself. This process taught me a way to live my life, to achieve simplicity through complexity and in the process I discovered a methodology for my living. Hence, one of the objectives of this document is to unfold that methodology for myself and to share it with prospective PhD candidates with a hope that it might prove useful to them to decide on their courses. In this document, I have talked about several aspects of my PhD such as choosing a PhD advisor, choosing a research topic, my approach to pursue the chosen research topic, a set of skills that I learned, the assertion of philosophy and depth, the art of problem solving and the methodology that I internalized through this process.

For each of these aspects, I have a two point summary *in italics* in case the reader wants to skip the details. I would like to clarify that some of the perspectives are out of my personal experience and are open to questions or different interpretations. Hope you enjoy reading my PhD story.

Let me first tell you about my background. I obtained an engineering degree from a reputed engineering college in Pune, India in 2007. To continue my thirst of understanding technicalities, I decided to appear for Graduate Aptitude Test in Engineering (GATE) in my last year of engineering. The test opened the coveted doors of computer science and engineering (CSE) department of one of the prestigious institutes in India, Indian Institute of Technology (IIT), Bombay, India. I completed my masters in 2009 and subsequently pursued a PhD in wireless systems domain from the same department. I completed my thesis work by the end of 2012 and graduated with a degree in 2013. Over next few paragraphs, I have penned down my thoughts to describe my experience with my advisors, my research topic, and what I think looking back in time.

Let me begin with the most important aspect of a PhD, the choice of an advisor. Fortunately, this did not prove to be a big hurdle for me. Since I pursued my masters from the same department, I had worked with both of my PhD advisors during my masters, and I had a reasonable frequency match with both of them. By frequency match I specifically wish to point out the degree to which I and my advisor could brainstorm over an idea and in the process understand and comment on each other's thought process. I vividly remember my meetings with my advisors especially during the inception phase (say first twelve months). This phase was important to me since I gained several important skills (how to think and what to think) and expertise. Hence, I think it is important to choose an advisor who is well known for certain set of skills and who is expert in a domain. Towards the end, I started driving the meetings with my advisors and it is this brainstorming process with my advisor(s) that carried me through.

Another important aspect is the background and the type of problems I and my advisor sought to work on. I was particularly inclined towards solving networking problems addressing the digital divide which overlapped quite well with the research agenda of my advisors. I spent last semester of my masters and first semester of my PhD to chalk out the ideas and problems in this domain. The overlap in the research interests helped me in later parts of my PhD to motivate myself to persist on the research problems. Looking back, I feel that this overlap of interests was a key factor which allowed me to sail and anchor my ship at various stages of my PhD. *Thus, the two takeaway points I would like to highlight* *here are as follows. (1) It is always in the best interest of both (you and your advisor) to have a few rounds of discussion to estimate the frequency match. Think about working on a joint problem; it could be a good starting point. (2) It is quite essential to invest a few months to find the right research areas and open research problems of mutual interest. Spend a few months on the problems of mutual interest before you decide to work with an advisor.*

This brings me to the next logical question of the choice of my research topic. To be frank, the choice of topic for my PhD was partially made by my advisors when I joined to work under their guidance. However, the topic overlapped with my research interests quite well and I spent considerable number of cycles (worth of a semester) to understand the potential depth and breadth of the topic. I could see at that point that there were at least three key research problems which were interesting and important, and which, if solved, could yield in a productive thesis.

That is when I decided to pursue the topic as my thesis topic. It is my opinion that, especially for a PhD in wireless systems, it requires both experience and vision to think over a research topic which has both depth and breadth. Experience helps to decide the scope and potential depth of the topic such that the topic can be pursued for 3 to 5 years and the candidate gains an essential set of skills in the process. Vision determines the significance and potential utility of the topic such that the work gains relevance once the research problems are solved and the candidate evolves as a computer science researcher. I feel that, for a fresh PhD student, attaining this state of experience and vision to choose a research topic is an onerous and demanding task, though it is not impossible. Rather, gaining this experience and vision is one of the desirable outcomes at the end of a PhD. Hence, I think that a model where the topic is partially decided by the advisors is a reasonably sound model.

Another possible model is where advisor and advisee explore interesting problems for initial 1 to 2 years before finalizing a thesis topic. This period can also be utilized to publish the research outcomes out of such exploration. However, it may turn out to be quite challenging to publish a work if there is no adequate relevance or significance, or without an end-to-end story. Also, in my opinion, working in a group to solve 3 to 4 problems to build a thesis or working on 2 to 4 incoherent problems is not an ideal way to gain a PhD. Instead, going by the “wait more, gain more” philosophy, a better way could be to find a coherent story, carve out the chapters and attempt to build those chapters on an individual basis. Brainstorming and collaborating with peers to write research papers is of course desirable; however, one has to go through the pain of painting the story by one’s own hands. The topic may already be available, but the task of carving out the research problems should be pursued by an individual. In my opinion, this teaches a candidate to think on the following aspects: *why is a research problem interesting, i.e, what are the key challenges to solve this problem, why are these challenges difficult to tackle, i.e., what is nontrivial about the research problem, why is the research problem important, i.e., what is the significance of the research problem in present and in future, what is the related work in this problem space; does it address any of the challenges; what does it lack, what are the possible ways to tackle the challenges, is there need to invent a new technique to address any of the challenges.* Especially for an experimental work, where one quantifies a complex phenomenon, addressing following questions at an individual level could develop a set of key analytical skills: *what is the hypothesis to solve the problem, how should one design experiments to test the hypothesis, how should one interpret the results, and how should one present those results. Thus, my two takeaway points here are as follows. (1) You may want to take up a topic of mutual interest with your advisor, or you may want to invest a couple of years to find the right research topic. (2) In either case, it is very important to be patient until you yourself get convinced about the potential depth and breadth of your research topic. Also, it is essential to have a coherent story with individual chapters carved out by you on your own.*

The next logical topic to comment on is the research problems of a research topic. I won’t go into the details of the actual problems that I tackled (I won’t say I solved them completely for I embrace “I have 99% certainty and 1% uncertainty” factor), but I specifically wish to comment on the kind of problems one should be working on their research significance and implications. In my PhD, one of the problems I worked on related to MAC design. In general, in the space of communication networks, MAC design is one of the most celebrated problems. Think of CSMA/CA or TDMA protocols. Working on a celebrated problem is both good and bad. If you find something novel, your work has high impact. However, finding something novel is nontrivial given the tons of prior art in a celebrated domain. Thus, it was quite challenging for me to point out the open design issues, the significance of addressing them, and a few novel techniques in solving those problems. However, it exposed me to a problem space where I could make a generic contribution. I also worked on a system design problem where we had to experiment and show that our idea works in practice; something which the community believed to be difficult to achieve. This problem involved characterization of a technology and implementing a few techniques to show that our idea is indeed realizable. This exposed me to the domain of system design, system implementation and a systematic approach of realizing an end-to-end system. I also worked on a couple of theory problems which sent me in tranquil for a few weeks. I realized at that point that the stories of people locking themselves in a room to crack a problem are indeed true.

With this experience, I feel that, in general, one should choose to work on broader problems than point problems. Let me explain what I mean by point problems (and point solutions). We see a lot of problems around us to which we can relate ourselves to; shoddy state of road infrastructure, annoying traffic jams, depressing waste disposal systems, sorry state of healthcare systems, inefficiency in agriculture systems and many more (the list only grows).

Readers from the west or developed parts of the world may not be able to relate to these problems as much as the readers in developing regions. Some of these problems can be attempted and tackled by computer systems and more specifically by wireless systems. There is nothing wrong in tackling such problems by the application of computer systems (e.g., one can build an image processing software to identify diseases in early stages of a crop development using mobile phones, or one can build a crowd sourced software to identify the traffic jams). In fact, in the course of a PhD, one should attempt to build at least one such system. However, one should also understand the entire solution space to solve problems like these and that there is a difference between research contributions (of your thesis) and solving such problems (by building systems). For example, some of these problems can be solved by efficient execution of plans and policies. In comparison, adopting a research prototype into reality often turns out to be far more difficult. Often, a research prototype with its specific technique or specific algorithm or specific evaluation becomes too specific to a use case. My intention is not be cynical about such research; one of the important contributions of such research is to show previously unknown, nontrivial, and valuable insights and provide directions for further research and development for a better human life. However, often these point solutions fall flat since it limits the applicability of the insights and directions of the work (and it may not be a research contribution), and to be of any practical use, these point solutions leave a lot to be desired. The necessary skill of carving a right abstraction of an important problem and to work on an appropriate solution of a broader theme is often missed in pursuing such research. While developing and evaluating a system is an important skill, my point is that one should pay necessary attention to work on a thorough and generic piece of research. Your topic should either have beauty (a very interesting and previously unknown insight) or utility (an immensely useful and practical work).

Hence, in my opinion, instead of a thesis revolving around such point solutions, one should attempt to solve broader set of problems. For example, if we imagine the motion of vehicles in and around a city, with a road and rail network, traffic signals or transport stations, and the demographics of a city, one can think on a system and a set of algorithms to efficiently route the traffic in the city. Designing such a system poses several interesting and broader set of questions. How should one collect the data for a scale of a city? What kind of networking protocols one needs to build if one wants to enable communication to and from communication nodes for a scale of a city? Is it required to design an entirely new protocol stack? How should one route the traffic? Are there any generic graph theory problems in designing such algorithms?

Are these problems attempted previously? If yes, are the solutions partial? What is the right solution space and how should one design systems and algorithms? How would the civic authorities utilize these generic pieces of research to develop a feasible solution. With these broader set of comments and examples, *the two takeaway points I would like to make are as follows. (1) If you want to seek a PhD in wireless systems domain, plan to have a broader set of problems to work on, be it characterization or quantification of an interesting phenomenon, breaking your head on an analytical problem, or engrossing yourself in designing and implementing a system. (2) Read through blog posts, articles, research papers to identify the right problems and design generic solutions. The goal of research is a search for new knowledge, hence you should be able to derive previously unknown and meaningful insights, or you should be able to show working of a previous unknown but a system of significance. I don't mean to insist on the practicality of your work, but I think it is good to know the value of a puzzle while you crack it. Understanding this aspect of meaning and significance, though it may consume a few months, will be especially helpful when you transition after your graduation.*

This brings me to the next question of the variety of skills I learned in the process. My answer is simple in this context. While hard work is important to get things done, I would rate aptitude as the most essential skill than hard work. In my opinion, aptitude teaches you what to work on and how. Be it a systems PhD or a PhD in theory, it is absolutely essential to be street smart and develop the general aptitude of (1) abstracting the problems and (2) solving the problems. While you may work hard on your PhD problems, and hard work is a good asset to have, in my opinion, it is far from sufficient to build on your PhD and survive in academia or industry once you cross the finish line. What do I mean by general aptitude? You may relate to the following activities.

Solving aptitude oriented puzzles. Solving algorithmic problems. Solving mathematical problems. Also, in my opinion, general aptitude is incomplete without encompassing the curiosity to know how does the world around us tick. For instance, you may want to find abstractions and understand the timings and reasons behind the evolution of digital cameras, or you may want to understand the evolution of a technology related to the use of energy (hydro, nuclear, renewables) and how it is helping human beings to live a better life. In this context, I would like to point out two more important skills. The ability to question and understand the ecosystem around you, and the ability to interact socially while you seek to question and understand. For example, would you rather be blunt in asking a question to a speaker or be humble to seek a discussion to know more about the speaker's work. It is also important to know what kind of problems your colleagues are working on and understand their perspective in solving those problems. If you have the right appetite to know more about the work around you, you will connect with a lot of people around you to build an interesting list of peers or contacts. These contacts come helpful when you seek to find a faculty position in academia or a research position in industry. I would like to keep this section short since I have two simple *takeaways. (1) General aptitude is equally albeit more important than hard work. (2) Social skills are essential if you want to know more about the world around you.*

Let me now describe my internship experience since I believe my internship played an important role in making me realize the importance of above set of skills. I interned with the smarter wireless group in IBM Research, India when I was two years into my PhD. Overall, it was a fantastic period of my PhD. I worked on an interesting research problem which was completely orthogonal to my PhD work. This work exposed me to a completely different problem and I worked on it from its inception. And given my experience, I now vouch for interns to work on a problem different than their PhD work threads. I was mentored by a shrewd and immensely cool researcher. I enjoyed the liberty and open ended discussions with my mentor and the smarter wireless team. It exposed me to the perspective of researchers in an industry research lab.

Further, the routine of my day was quite different from an academic environment, a bit more systematic I would say. I worked for 5 days and enjoyed the weekends exploring the city. I made a set of good friends too. Interacting with them brought new perspectives in my own work. I liked the overall shift in the atmosphere and it was a much needed break from my regular PhD work. I learned quite a few things from my internship. Wireless systems domain by its nature is an applied field. Unless your research involves physical layer advancements, you need to choose your area of application very carefully. In this respect, working in an industrial research lab, where research is neatly tied with a utility driven outcome was an eye opening experience. Often people work on whacky or seemingly cool but point ideas or point solutions which are neither straightforward to realize in real world nor do they result in any meaningful insights. Unless an outcome is a meaningful insight, a useful system, or an interesting future direction, most of such research drifts in void and is a lost cause. People tend to add tons of literature in that domain which ceases over its course. Of course, working on cool ideas is an interesting and enjoyable experience. And a seemingly closed area sometimes opens doors for another area in future.

However, if you don't want be in those less probable outcomes, and if you are looking for an exciting experience in an applied field, working on technology relevant research problems is a way to go for. In such technology relevant space, most of the times, researchers end up becoming the first few to crack the problems and derive meaningful insights which are later hooked on to by various elements in the research and development ecosystem. I would like to close my internship chapter with the following *two takeaways. (1) One should break the shackles during a PhD to go out and work on an entirely different problem. (2) Depending on the possibilities of internship, one should intern with an industry research lab to taste the nonacademic research atmosphere.*

This brings me to the last part of my story as I look back on the 4 years which I invested in working on my thesis. To begin with, my advisor had an interesting idea for a research topic. It overlapped with my interests, I understood the overall problem space, I carved out the research problems, I managed to tackle some of them with some necessary details worked out, and in the process, I evolved as a computer science researcher. Looking back, I think I could have worked on certain set of tasks in a more systematic way. For instance, while focus on one's own research topic is important, one should also attempt to understand the overall research ecosystem. What I mean is that I did not ask and think on the following questions to the extent I wish I should have thought. *What are the fundamental research problems in my domain? What are the game changing research problems? Are there any attempts made to solve them? What are the technical challenges in solving these problems? What is the next wave of advancement? What are the factors that decide such advancement? Are these advancements necessary toward a better human life? What is the next hot area in my domain? What are the problems my colleagues in various universities across the world working on? What do the technews sections or blogs or articles say about the state of the art? Why are those problems interesting? What kind of problems I would like to work on in future? What are the problems that industry research labs in my domain working on? Why are these problems relevant to industry?* I think it is very important to seek answers to these questions to be a good computer science researcher.

In the beginning of this document, I mentioned about a methodology that I learned from my PhD experience. Let me end my write up by describing this methodology. I think, a PhD experience is about learning to understand and harness abstract skills and to use these skills to simplify any problem one encounters in life. Any problem that you may face in your life; whether you have to solve a technology problem, or if you have to buy a property, or suppose you have to take a key decision in your life, these abstract skills are immensely helpful. The first step in tackling any such situation is to ask those familiar set of questions. What is *really* the problem/situation I am into? And why is this problem/situation seemingly difficult to tackle? Answering these questions often helps one to first seek the desired clarity about the problem, and a PhD experience makes these questions and answers an essential part of one's life. The second step is a matter of focus and patience. The journey of earning a PhD teaches one that solving a problem or tackling a situation demands understanding the background or prior art, understanding the overall context and then applying known or new techniques to attempt to solve the problem. One attempts to solve the problem in a systematic manner. And a repeated execution of this task through a PhD prepares a candidate to seek the desired simplicity through the complexity. The third and final step is to discover and execute a set of simple actionable items. The transformation of the initial complexity into the final simplicity is the key in this methodology. At the philosophical level, as Oscar Wilde has pointed out, *Life is not complex. We are complex. Life is simple, and the simple thing is the right thing.* And I believe, the process of PhD is a definite path to learn to seek this simplicity.

Shahriar Nirjon

When I was a child, I thought, PhDs are done by very old people. Those who do it are locked up in a dark room for the rest of their lives and their job is to look at something through a microscope. I don't think I ever had a plan of getting a PhD at any stage of my high-school or earlier life. The first time it felt like I could do a PhD is when I saw the professors at my undergraduate school in Bangladesh. From their body language of those who had a PhD versus those who did not, I understood, there is something important about having a PhD. I had to know this, and so I kept looking.

Eventually I knew about the GRE, the application process, and after applying to about a dozen US schools, one weird computer science department within the University of Virginia seemed happy to have me. I left my job, packed my baggage, and got into a 29 hour flight to the USA. On 16th August 2008, I landed on the beautiful city of Charlottesville. That's when the journey started.

Soon I started to discover a grad student's life. The stories from the PhD comics had begun to make sense. I was taking more classes, spending more hours in TA'ing undergrad courses, running after free foods, and doing absolutely no research. While I was getting used to this new lifestyle, one day, during a departmental luncheon, I met my professor.

The meeting was neither planned nor totally unexpected. Every graduate student at that time was looking for an advisor who will agree to fund him for the rest of his PhD. I knew I too have to talk to someone and I was looking forward to that day. But it wasn't so easy for me to spot a professor among all those unknown faces. The USA is not like my country where I could easily tell apart a professor from a student by his appearance, age or body language. It happened that after talking to someone for minutes, I discovered, he was just another 6th year graduate student. At some point, I started to think, I should meet and greet every single person in the room and ask first: 'Are you a professor or a student?'

I was about to leave as I my lunch was over. Then, at one corner of the room, I saw a wise-looking, nicely-dressed, 50+ aged person, who was sitting on a chair while talking to an eager listener, possibly a student. I asked one of the fellow grad students standing near me, 'Who's that guy?' He said, 'Which one? The one who is talking is Professor Jack Stankovic and the other one is Professor S.' Well, by that time I had become used to it and that was not too surprising any more. Thinking - now I have two professors at target, I went forward, and I spoke.

I was familiar with Jack's work long before coming to the UVA. I probably have seen his picture on the UVA's website too. But that day I was not aware that it was him. But anyway, I was delighted to finally meet someone and I expressed my earnest desire to work with him – which I would probably do to anyone that day. He told me to meet him at his office right after the lunch. I had my first meeting with Jack.

Since then, I guess, I have met him at least a thousand times. And now that I look back, I see the change in me. I was once a shy, less-talking, confused, and an aimless graduate student like anyone else.

Jack's mentoring has turned me into an aggressive, talkative, confident, and a clear-sighted person who knows what he is doing. For me, it was easy. I just followed him. I will miss him big time once I graduate.

ZAHIR KORADIA

My PhD has been a journey of unexpected events, both exhilarating and disappointing. At the risk of sounding "filmy", I present the parts of my journey that seem to have influenced my PhD the most.

Let us dial back to 2007. I was completing my Masters at IIT Kanpur, working on routing in Delay Tolerant Networks for my thesis. The work was supervised by Prof. Bhaskaran Raman. Even though I joined the program with a vague interest in artificial intelligence, by the end of the program I had taken a liking to Computer Networks, thanks to Bhaskar, and had taken all the Networks related courses available. It was placement season and it was time to take a call on what comes next.

I had been influenced by the idea of working for the development of those in need. Voluntary service to the society is a part of my culture – everyone in my family and community contributes to the community voluntarily. However, I felt that contributing professionally would allow me to spend more time on it and also ensure that I demand quality off myself when doing so. I had also loved the liberty of working on topics of my liking. Finally, I had thoroughly enjoyed my relationship with Bhaskar. So I decided to give PhD a go with Bhaskar. Practically, this meant that I had to apply for PhD at IIT Bombay, since Bhaskar was moving there. Getting admission was not as challenging as I imagined and soon Bhaskar, his family, and his first PhD student moved to IIT Bombay.

My initial topic of research was design of communication systems for large scale disasters. I spent time with Humanitarian organizations learning their needs, studied related work, and tried to formulate technical challenges. I could not get very far, mainly because the domain was completely unexplored. Such a situation is ideal for experienced researchers as they get the chance to do foundational work in a domain and potentially gain a lot of recognition. As a student doing so is hard because, one doesn't have the experience to differentiate a field with significant potential for research impact from a field with little impact potential. I eventually gave up and decided to work on looking at road traffic monitoring in India.

By this time it was last few months of 2008. I was getting married in December and what transpired next resulted in a significant change in my PhD and my career. I was looking for married students' accommodation at IIT Bombay. As it turned out, there was a queue for such accommodations. One is required to make a request to the relevant authority and once an accommodation is available one is informed of the same. Some enquiry among the married students indicated that it generally took 1-1.5 years after making the request to get the accommodation. So I would have to wait that long before I would start living with my wife. This was not acceptable to me. Strike One! I talked to Bhaskar to see if he could give me higher stipend, Rs 11,000/- at the time, so that I could afford accommodation just outside the campus. Bhaskar was okay with paying more, but the norms then prevented him from paying me more than Rs. 25,000 per month. There weren't any rules, just unwritten norms. With rents outside IIT ranging around Rs 10,000 per month the stipend was not going to be enough. Strike Two!

The next option I considered was my wife getting a job in Bombay. She was returning from the US after completing masters in city planning from GeorgiaTech, so her qualifications made me hopeful. Unfortunately, she could not find a job in time. Strike Three! I could not continue full time PhD.

At about the same time I came across an opening at a start up company called Gram Vaani that was planning to work on technologies for NGO run radio stations. The company was being setup by a person whom I had briefly met at University of Waterloo, Aaditeshwar Seth. Adi was doing his PhD there when I had gone for an internship in 2006. I thought that it may be possible for me to work at Gram Vaani and use a good percentage of my work towards my PhD. I consulted existing PhD students on the thought of converting to part-time PhD and working along with it. We have this email alias of all PhD students at IIT Bombay where I posted my query. I think such a list is invaluable as a source of information and support. All those who responded to my query said that it was doable but DON'T DO IT. At the time I did not think I had a choice, so I still went ahead. I contacted Adi, we had a few email exchanges and phone conversations and I was a lead developer at Gram Vaani in New Delhi.

Part time PhD at IIT Bombay requires that I have an external advisor at the organization, where I work. Adi readily agreed to be one. The general agreement I had with Bhaskar was that I will visit Bombay every couple of months to discuss PhD progress and we will have weekly con-calls in between. This was not to be. Year 2009 was the first year of Gram Vaani's existence and I had no idea working at a start up was going to be this hectic. It was not fun being the first person to leave office at 9pm and come back to a wife who looked forward to quality time since we had just gotten married. I could only manage one visit every 6 months and very little contact meant Bhaskar was unable to provide inputs for my research. Often, there was no research to give inputs for. Eventually, we all agreed to an arrangement where Adi becomes the primary supervisor and Bhaskar will only play a secondary role in guiding me. Eventually, the hard phase of 2009 also ended with me having more time for PhD and personal life.

My PhD work has revolved around exploring how different computing technologies can help NGO run radio stations in India. The research draws on and contributes to the domains of technology for development, HCI, and computer networks. It wasn't so at the outset. My research progress committee – a committee of professors/researchers who evaluate PhD students' progress annually – were uncomfortable with the amount of "non-research" component and felt the need to have more traditionally accepted computer science research in my PhD for it to be acceptable. Being one of the first PhDs in the domain of application of technology to development in India, I have faced this criticism constantly. The non-traditional domain of research has also often resulted in unfavorable attitude in post PhD hiring processes. When presenting my HCI related work at one of the institutions an attendee said – "But this is social science research!" Without meaning any disrespect to my research progress committee or the evaluators I have to disagree with their definition of what constitutes computer science research. But I am a student, not in the position of power, which means I have been at the receiving end of the disagreement. It is quite a struggle, but what has kept me going is the real impact my work was making in the lives of the underprivileged.

Some of my research has evaluated usage of software at more than 10 radio stations. Another research has looked at strategies required for scaling deployments at over 30 radio stations. Yet another research studies the state of cellular data connectivity at 7 rural locations. All these works have required huge human resources for planning, management, and support of deployments. I have been fortunate to have been able to work with teams of 4-8 people for these projects, which at times included masters' students and research assistants and at other times included colleagues from Gram Vaani. I believe none of my research would have been possible without such astounding support and availability of such large teams. Being a part of Gram Vaani has allowed me to do some particularly unique research – creating case studies of scaling technology in developing regions is hard to do.

I am currently writing up my thesis, and looking back, I feel glad and grateful about how everything has transpired. At the end of it all I must add that the PhD is really a very small part of everything that I have achieved – beautiful relationships and better understanding of my own self. I hope my experiences will help others taking on this journey.

NICOLA DELL

I had no idea what I wanted to do when I started my PhD at the University of Washington. Not knowing that Computing for Development existed as a research area, I started out doing something completely different with a different advisor. Then, during my first year, I started going to a weekly graduate seminar that featured speakers talking about projects that use technology to promote international development. Being from Africa, this particular area intrigued me. I wanted to get more involved and figure out if I could focus my PhD thesis on research that aimed to solve problems in developing countries.

At the seminar I was fortunate enough to meet Professor Gaetano Borriello, who encouraged me to explore the research area and later agreed to be my advisor. This was a key moment for me. Finding an advisor who will encourage and support you is undoubtedly one of the most important steps in succeeding in a PhD program. I believe that I could have been happy and interested in a variety of different research areas or topics. The important thing was to meet someone who I connected with and who would mentor and advise me through the long and difficult journey that is a PhD. Without such a great advisor, I would have undoubtedly pursued a very different path.

The Computing for Development research area is new and exciting, but also extremely challenging. One of the best things about the research area is that many interesting research projects need to be deployed and evaluated in developing countries. This allows me to travel to many places in the world, connect with different people and experience diverse cultures. Being someone who loves to travel, this was a major benefit for me and I during my PhD I have worked on projects in India, Peru, Mozambique and Zimbabwe. Since my projects focus primarily on health related topics, I have traveled to rural clinics in remote areas of these countries and trained people to use the systems that I built. This has been incredibly rewarding.

However, it is also extremely challenging to have research projects that happen very far away from your university, your research lab and your advisor. It can be difficult to find good field partners to collaborate with and also to design projects that will be successful both as computer science research and also in providing tools that people can use to better their lives. It is often difficult to manage target users' expectations and to be a graduate student that doesn't have the time or the resources to fully develop and sustain a large scale project for a long period of time. For me, it has been essential to learn how to prioritize different aspects of the research so that I am able to be successful as a PhD student. Again, I could not have done it without the support and encouragement of a great advisor.

I have had a wonderful experience in the PhD program at the University of Washington. It has been an extremely challenging process that has pushed me to new depths that I undoubtedly would not have reached on my own. It has definitely not been easy, and I still have some important hurdles to overcome before I finish – writing my dissertation and finding a job. However, I am still excited about the research that I do and confident that I will have the support and mentorship that I need to succeed.

Figure 1: Collecting data from an auto-rickshaw driver in India.

Figure 2: Training nurses in Zimbabwe.