The Challenge of Being an Engineer – Reflections from a Security Engineer

Feng Hao

Preface: This article was originally written in June 2014 for a special issue of the IEEE Security & Privacy Magazine on lessons learned by existing editorial board members based on their professional career. The invitation was to provide a white paper containing:

- *"A short description of lesson learned from your experience or your expertise"*
- *"Four or five key points you plan to make along the way"*
- *"What this lesson means for researchers, practitioners, policy-makers"*

In the end, my article was found not to fit the special issue well, so it was not included in the special issue. However, Shari Pfleeger, the Editor-in-Chief, encouraged me to publish it as she thought my lessons could be useful to some young PhD students. It was her encouragement that led me to make this article available on my personal website.

The "four or five key points" I make are:

- *Publishing new ideas is not easy in academia*
- *Be patient*
- *Be diverse; and be ready to dig into new areas with a shovel.*
- *Enjoy and be proud of your work*

And the lessons, which I am aiming at early-career researchers, are that when working on "new" solutions to real-world problems, a researcher should be patient regarding the publication of the results.

New Year's Eve is a traditional time for reflection – regarding both lessons from the past and the future outlook. Hopefully others might find some of my lessons useful.

(Preface written on 31 December 2014)

[Before my PhD]

Before I started my PhD, I happened to come across a series of on-line video presentations[1], recorded on 17th February 2003 in celebration of Prof Roger Needham's fiftieth year in Cambridge and fifth year as the Managing Director of Microsoft's research laboratory in Cambridge (an event called "Roger Needham 50+5"). Roger served as the head of the department at Computer Laboratory, University of Cambridge from 1980 to 1995, Pro Vice Chancellor of the University from 1996 to 1998, and Managing Director at Microsoft Research Limited, Cambridge from 1997 to 2003.

On a wheel chair, having been diagnosed with incurable cancer, Roger gave the final speech to all attendees at the "50+5" event. (He passed away a month later.) In his words: "I have the greatest respect for the people who have built the theoretical underpinnings of our [computer science] subject, and we wish them every success because it would enable the people who want to get on

[1] http://research.microsoft.com/en-us/um/cambridge/events/needhambook/

and make things, to do it better, and to do it more quickly and to do it with less mistakes. And all of this is good, but, at the end of the day … [Roger puts on a construction worker's hard hat], I'm an engineer."

[Applying for PhD]

Roger's words touched my heart. I considered myself an engineer too, because making things work had always been my passionate interest. I felt the Computer Laboratory at the University of Cambridge would provide an ideal environment to develop my interests. So, in 2003, I sent in my PhD application and nominated Professor Ross Anderson as my primary advisor. (Ross had been a PhD student of Roger.) When Ross asked for my research proposal, I expressed my interest in working on a practical engineering problem: how to effectively combine cryptography and biometrics. This subject was later commonly known as "biometric encryption", but at the time of 2003, it was very new. Cryptography and biometrics are potentially complementary but are technically incompatible – while the former requires all data be exact, the data obtained from the latter is fuzzy by nature. As it happened, Ross informed me that an emeritus professor in the same laboratory was working on the same problem and just published a paper. The professor's name was David Wheeler, who had been Roger Needham's PhD advisor. Ross emailed a copy of David's paper to me. With great respect, I read David's paper and was impressed by the insights it contained.

As part of the PhD application process in the Computer Lab, every candidate has to be interviewed. Since I was in a foreign country at the time, the interview was conducted by telephone. Ross (together with another professor, John Daugman, who later became my second advisor) called me on time. I remember the first question Ross asked me: "What's the problem with David's paper?" I was a bit shocked. Having been brought up in an Asian culture that treats the authority of the teacher as paramount, I was being asked to challenge one of the most senior professors in one of the most prestigious universities. After recovering from my shock, I slowly started to explain the basic idea in David's paper and expressed my doubt as to whether the idea was really feasible, since it had not been tested with real biometric data. Ross seemed satisfied with my answer. It later occurred to me that the actual answer was perhaps not so important; what really mattered was whether the PhD candidate was able to think independently and had the guts to critically question an authority, even one of the most powerful and the highly respected in the field.

[PhD research]

In 2004, I was admitted by the University of Cambridge as a PhD candidate and started to pursue my research under the supervision of Prof Ross Andersen (and also Prof John Daugman as my second advisor). I was extremely fortunate to be in an office next to Prof David Wheeler's and found it convenient to discuss my research work with him. Although he had been retired for over 10 years, David still came to the office every weekday early in the morning (usually being the first in office within our research group) and left late in the afternoon. That he did this purely for the interest of research filled me with great admiration and respect. In the paper by David that I had read earlier, he proposed an initial idea of combining cryptography and biometrics by using an error correcting code. However, after testing the code using real iris biometric data, I found the proposed method did not quite work. I was dejected and wondered if I was wasting time on an impossible problem. Then David said something that enlightened me: "You normally change the problem if you can't solve it." This literally saved my research. Through thinking outside the box, I found that the seemingly impossible problem could be changed to an equivalent form and then could be tackled effectively. I was excited by the breakthrough. My two advisors, Ross and John, were also very delighted to see the result, as it was the first in the field. (Unfortunately, David passed away in 2004 before I was able to finish the paper.)

[Paper rejection]

The idea in the paper was new, however, publishing a new idea turned out to be harder than I had expected. Following my advisor's suggestion, I sent the paper to a top security conference. While feeling optimistic about its chance, I received a rejection a few months later containing some harsh reviews. I then sent the paper to a second conference; it was again rejected. On a third attempt, I submitted the paper to yet another conference. But it was still rejected.

As a young PhD student, I felt lost. It seemed the reviewers did not like my work at all. In my most depressing moment, I remember my advisor Ross came to me and said: "Don't worry. I think it's a good paper". That was an enormous encouragement to me. At the time I may have doubted myself, but I trusted my advisor's judgement. Instead of wasting time and energy on getting upset and frustrated, I was determined to move on. While still waiting for the paper to be published, I started looking for new research challenges. The intense energy that I had to put on solving a new problem left me no time spare to feeling low about my paper's rejection. By the end of my PhD research, I had managed to solve three major problems and was able to include them as three main chapters in my thesis. Finally, the paper on combining biometrics with cryptography was accepted by IEEE Transactions on Computers in 2006 (with almost no change from its original version). So all came out well in the end.

[Post-doc research]

During my post-doc research, I happened to encounter an interesting cryptographic problem: how to establish secure communication between two remote parties without relying on any trusted third party. This research problem is commonly known as Password Authenticated Key Exchange (PAKE). The PAKE problem had been widely believed to be impossible, until Steven Bellovin and Michael Merrit's seminal work on Encrypted Key Exchange (EKE) in 1992. However, the EKE solution was found to have some technical limitations. In fact, all other existing PAKE solutions also turned out to have various shortcomings. Furthermore, patent issues were a major obstacle, as nearly all PAKE protocols were then patented. I was therefore motivated to try to design a PAKE solution that would be better than the existing ones and be patent-free. Together with my collaborator (Prof Peter Ryan), I applied a technique that was originally developed in one chapter of my PhD thesis to solve the PAKE problem. The result was a new PAKE protocol called Password Authenticated Key Exchange by Juggling (J-PAKE). At the time, I was learning juggling and though it would be cute to use it to name the protocol, since the working of the protocol was similar to the juggling of balls.

Feeling proud of the work, I submitted the paper to a major cryptographic conference. However, the paper was rejected. The reviewers evidently expected every designer to follow some existing models when inventing a new PAKE protocol, but we had not done that way. After working as a post-doc for a year, I decided to change my career path and work in security industry. I wanted to work in an environment where practical solutions to real security problems were appreciated.

[Industrial work]

Three years working in industry, first as a Software Engineer and then as a Lead Security Engineer, was invaluable experience to me. It not only allowed me to be engaged in many exciting practical security projects, but also enabled me to see more clearly the gap between academic research and industrial needs. In industry, cryptography has become a key enabling technology to protect information security. Practical cryptographic techniques are in great demand. On the other hand, cryptographic solutions published in academic papers tend to be overly complex, hence not meeting real-world needs. Partly this is because many cryptographic researchers do not want to write a single line of code to test out their theoretical designs. When this becomes a culture in the academic community, the gap widens.

In 2010, the J-PAKE protocol instantly received attention over the Internet – Mozilla Firefox decided to adopt the protocol to implement a secure sync service and deploy it to 450 million Firefox users. To an engineer, nothing is more satisfactory than seeing the research outcome used and deployed in real-world applications. I was greatly encouraged. After all, engineering solutions to real security problems are still appreciated by the community; it is just that the process takes a bit longer than you might have expected. At the time I had been considering to return to academia and this development made me want even more. I felt if I were back in a university environment, I could perhaps do more.

[Returning to academia]

In December 2010, I took a faculty offer from the School of Computing Science at Newcastle University and became a lecturer. Having spent three years employment in security industry, I was glad to be back in a university and to become a full-time researcher again.

Back in the university, I started working on a new challenge: electronic voting. I observed that there was a significant gap between theory and practice in the field. The theory of verifiable e-voting had been extensively studied for over twenty years, with many e-voting systems proposed in the literature. However, in practice, few of these systems had actually been implemented and almost none of them used in real-world national elections. Something seemed amiss.

I was motivated to narrow that gap. Together with a colleague (Matthew Kreeger), I started to critically examine the fundamental theory underpinning this 20 years of research on e-voting: particularly, we questioned whether the role of trustworthy tallying authorities that all existing verifiable e-voting systems required was really necessary. To support our argument that it was not, we designed a new e-voting system that provided the same verifiability as other systems, but without involving any tallying authority. We called this new type of e-voting system: "self-enforcing e-voting".

In 2010 we first released the paper as a technical report (on IACR eprint 2010, No. 452), and then submitted it to a conference. But the paper was subsequently rejected. During the next four years, the paper was repeatedly rejected by major conferences, until its recent acceptance by USENIX Journal of Election Technology and Systems (JETS Vol. 2. No. 3, 2014).

Four years of repeated rejections are not easy for a researcher. Fortunately this time, from my previous experience, I had learned to be patient – and extremely patient. I kept faith that one day the paper would be accepted by the mainstream research community eventually. In the meantime, just as I did in my PhD, I kept diversifying my research areas. Besides electronic voting, I also worked on a variety of different research topics so I didn't feel I was putting all my eggs in one basket.

The fate of my e-voting work made an important turn in 2013. After reviewing my proposal, the European Research Council (ERC) decided to award me a starting grant of €1.5m to support my further investigation on "self-enforcing e-voting". The transition from receiving consistently harsh reviews and rejections from top conferences, to being awarded a substantial grant to support my further research and getting the final acceptance of the paper's publication, was an interesting personal experience. And in retrospect, I realise I learned more from the process of receiving, and carefully considering, rejections rather than from the acceptances I eventually received.

[Summary of reflections]

Ten years on, when I revisit Roger's words in 2003, I find they touch my heart even more. Being an engineer is fun and exciting, as there are plenty opportunities to work on practical problems with potential real-world impact. On the other hand, being an engineer is challenging. The challenges are not only scientific, but also cultural. When the culture in academia favours complexity in designing security solutions, an engineer's challenge is to persist in searching for simplicity; whereas the academic culture favours research that is driven by improving existing literature, an engineer's challenge is to focus on research that is driven by tackling real-world problems; whereas the academic culture encourages compliance with established methods, an engineer's challenge is to break the compliance and explore new ideas; whereas the culture moulds researchers into the mainstream, an engineer's challenge is to stay independent; when it takes a long time to publish a new idea, an engineer's challenge is to be patient and carry on.

Finally I conclude by citing Roger Needham's two famous axioms:

- "Good research comes from tackling real problems."
- "Good research is done with a shovel, not with tweezers."

[Acknowledgement]

My sincere thanks to Prof Brian Randell for proofreading this article and providing many useful comments.

Feng Hao (Reader in Security Engineering)

School of Computing Science

Newcastle University, UK