

ADVICE TO MID-CAREER RESEARCHERS

We are starting a new series to provide advice to mid-career researchers. There are a number of programs that SIGMOD organizes for researchers at the beginning of their careers (PhD Symposium and the like) and senior people do not (at least they think they do not) need much help. There are considerable challenges for those who are about to transition from an early researcher to a more senior role. In academia, these are people who are about to get tenured that comes with starting to think of moving from shorter-term research objectives to longer-term ones. In industrial research, this corresponds to the transition from participating in projects to initiating and leading them. As a community we don't seem to talk about these challenges much. That is the gap this series attempts to fill. We will get the views of senior researchers from diverse backgrounds and diverse geographies. We will continue as long as we find original advice and the views are not repetitions.

*M. Tamer Özsu
University of Waterloo*

Mid-Career Academics: What I Have Learned or Wish I Had Known

Qiong Luo

Hong Kong University of Science and Technology &
Hong Kong University of Science and Technology (Guangzhou)

As I agreed on contributing a piece to this series of which Tamer is in charge, I read all the preceding articles to get inspiration. The impression I got was, “gosh, I wish I had known all this back in my mid-career days and even these days!” For example, I did place students on my collaborative projects, but would have made faster progress in some projects if I had insisted on having weekly meetings. Or I should have considered a more detailed list of factors before committing to a task of considerable amount of work.

Now that our goal of the series is to provide original advice and non-repetitive views, I look at myself and think about my experiences that may be different from those in the previous articles and have room for further discussion. As such, my views can be read as complementary to previous articles of this series.

I joined the Hong Kong University of Science and Technology (HKUST) in Hong Kong in 2002, got promoted to Associate Professor with substantiation in 2009, and stayed in this position for 12 years until 2021 when I was promoted to Full Professor. I focused on architecture-conscious databases, especially GPU-accelerated query processing, on my tenure track. Common wisdom is that after getting tenure, one can take on new directions of higher risks and more uncertainty. I did that by switching to working with domain scientists on accelerating genomic and astronomical data processing.

It was fun and fulfilling to work with domain scientists and make an impact helping with their scientific tasks. However, the learning curve of cross-area collaborative projects and the significant engineering efforts took a toll on my research output. Furthermore, my graduate students at the time seemed less interested in working on scientific data systems research than getting an industrial job; as a result, between 2012 and 2014, seven PhD students of mine transferred to Master's programs and went off to work for technology companies after graduation. On the one hand, I am happy that my training helped my students to pursue their careers of interest; on the other hand, I wonder whether it would have been better if I had focused more on research issues in my area of expertise and guided my students better to lead them to a successful PhD. Therefore, in retrospect, the first advice I would give is that, if you would like to climb the career ladder faster, you might want to limit your post-tenure freedom, stay focused, and continue finding ways to motivate your (new batch of) students.

Related to focus, my second piece of advice would be to always have a central thread to connect all your research efforts. As academics, our interests shift from time to time and we are free to explore different topics, especially after getting tenure. I, as an example, work on accelerating genomic data analytics, astronomical data processing, graph queries, and distributed data systems. Understandably, the most frequent question I get is

along the line of whether there is something one size fit all, e.g., “are there common techniques or frameworks that work for all these various systems?” My answer varies depending on the context, but I always explain that systems, unfortunately, are not one size fit all, but common techniques for performance improvement, which is the focus of my research, do apply, for example, data-parallel primitives on the GPU, addressing the bottleneck to speed up the overall performance, cache-friendly access patterns for latency, and bandwidth utilization for throughput, and so on. If you identify a common theme in your work and stick to it most of the time, it will help you build on your strengths and keep a holistic view on your work.

With a clear focus, you will exercise persistence in the following years to go deeper in your research and make bigger impact. Post-tenure careers have no clock ticking any more, so what keeps us going are mostly our intellectual curiosity and genuine interest in nurturing young talent. I have been working on GPU-accelerated data processing since 2006. It was exciting to look for ways to effectively utilize this off-the-shelf specialized device of tremendous raw computation power and memory bandwidth to speed up general data processing. I track the rapid development of the hardware and its software environment over the years, and my research has been in constant update to discover the effects of these technological advancements and utilize them in fundamental data processing techniques as well as various data-intensive applications. Furthermore, my group maintains the tradition of publishing our code and experiments for repeatability tests and facilitating follow-up work in the community. Even though these efforts take a significant amount of time and energy, they make further impact than merely papers and train students better through the process. As a result, most of my students go work in the industry and lead significant product development projects.

Collaborative projects in different application areas took me longer than my other research to publish, and so far, I have limited success in getting them into mainstream database venues. Nevertheless, intellectually it has been a rewarding experience in my post-tenure years. It allows my students and myself to work on real-world problems and see the impact of our work in various domains. For instance, we helped a leading genomic sequencing center in China to speed up their software tools on sequence alignment, genome

assembly, and genome-wide association studies; we also worked with astronomers to build online-processing systems to extract celestial objects from observatory images, crossmatch these data with existing catalogs, and track transient stars in real-time through image subtraction. These collaborations evolve over time, and I now work with colleagues in life sciences, radio astronomy, and organic chemistry. Luckily, along the way, I found students who were interested in such kind of research and successfully supervised their PhDs on these projects. Even though students may not continue in e-Science areas after graduation, the skills and collaboration experience from these projects landed them excellent jobs in technology companies.

Speaking of students, you might get a sudden drop or a big increase in your group size after you get tenure, because your first batch of students graduate, or a new batch of students join. Some colleagues go on and build a huge group of 30-40 people consisting of students, RAs, and Postdocs whereas others supervise only a couple of students at one time. The choice depends on the nature of your research, your preference, and funding situation. I find it most productive and enjoyable with a group of 4-5 students, as I can have weekly one-on-one meetings with each of them and guide them individually. My PhD advisor told me that PhD study is like apprenticeship, so it would be best to provide sufficient opportunities for students to learn from you closely. Even though my work is empirical, and my funding is in good shape, I have always kept my group size a handful. This way, each student is trained by me directly in all aspects, from identifying research problems, formulating solutions, implementation, evaluation, writing up, presentation, and discussions. My students and I treasure the tight relationship we built from frequent exchanges regularly and working together intensively under deadlines. Additionally, I take great effort in recruiting and supporting female students, and keep learning from my experience and others how to best nurture the next-generation women computer scientists.

Outside your research and advising, you will be commissioned for bigger responsibilities in teaching and service immediately after you become an associate professor. Our department’s tradition is that newly minted associate professors will become undergraduate program coordinators or postgraduate program coordinators. I did both, serving two years in each role.

I was also chartered to start the departmental alumni affairs committee, developing alumni networks, assembling resources, and organizing various alumni events for three years. Most recently, I chaired the departmental faculty search committee. At the university level, I worked with a few colleagues across schools to advise the acquisition and maintenance of the university's HPC center as well as participate in other university committees. These service tasks took a large amount of time and energy, but I found them great opportunities in learning about all aspects of university functions and making our own contributions. I am sure you will also find some services that are a better match than others for yourself.

As for teaching, you might want to reexamine the curriculum of your program and develop new regular courses, as opposed to topics courses, in your area of expertise. For instance, I found that my department had no regular graduate level courses on parallel computing, but PhD students needed such training for their research not only in systems but also in other areas such as machine learning. Therefore, I developed a course on parallel programming covering traditional message-passing and shared-memory programming as well as current CUDA-based GPU programming. This course introduces principles and methodologies as well as gives students hands-on training. It has become one of the core courses in the PhD course requirement. Teaching this course not only benefits the enrolled students but also allows me to regularly review the engineering fundamentals of my own research and train my PhD students as teaching assistants.

Before tenure, you have probably served on quite a few program committees and organized some workshops in your research area. Now it is time to extend your professional service to a higher level by not only continuing serving on program committees of major conferences, but also seeking opportunities to become meta-reviewers, area chairs, and program committee co-chairs. Such increased conference reviewing responsibilities are especially important for our community as our major conferences have all moved to multi-round submissions. If you have learned from insightful and constructive reviews on your own papers, you will graciously return the favor by writing high-quality reviews on papers of others; if you have unfortunately experienced low-quality reviews, you will have the motivation to act differently from those

negative examples and make positive impact on fellow researchers. Within conference program committees, you should remember that your job does not end when you enter your reviews; discussions are important and you should remain engaged. I am sure you will learn a lot from senior researchers by reading their reviews and participating in the discussions on the papers you are also reviewing. An important issue to which you and your students should pay close attention is to carefully declare and observe conflict of interest in paper submissions and reviewing, especially if you have had many collaborators. This has become an increasing problem in the community.

During your tenure track period, while you were working hard on building up your research, training students, and securing fundings, you surely received help and advice from senior colleagues in your university as well as in the professional community. Now it is also time to imitate those senior colleagues that you regard as role models to give back to the department and community by taking up more service tasks, helping junior colleagues, and actively serving as journal editors and conference organizers. With the increased responsibilities, you may find that your schedule becomes even more hectic and fragmented, so it is crucial to select the services that you think are most important, and that you enjoy most, and keep yourself organized with deadlines of your committed tasks recorded and met.

Finally, careers are important, but they are only a part of your life. A senior colleague once told me, a child is worth 1,000 papers. Few of us will have that number of publications, but many of us in mid-career raise young families. You might find yourself juggling a lot between a demanding career and personal life. At such times, be honest with yourself, set your priorities right, and enjoy the moment. In my case, my students are used to me disappearing physically after office hours and reappearing online late at night. My children were born when I was tenure track, and my evenings and weekends were filled with their gymnastics training trips and other activities. It felt exhausting, but focusing one day at-a-time we could get by. Looking back, I am thankful for the experience and happy that all the efforts turned out worthwhile.

In summary, mid-career may be your prime time in that the tenure pressure is gone and you are riding on the momentum for a lot more opportunities and

responsibilities than before. You can venture out for cross-area collaborative projects, continue training new batches of students, develop new courses in your areas of expertise, and actively serve inside and outside the university. Accept that new areas and collaborative projects might have a long lead time, and take them as learning opportunities to expand your research and train your students. It is always an open challenge how to motivate and guide research students; close and personalized master-apprentice interactions will greatly help. Your department will count on you on important administrative or academic committees, and our professional community will rely on many of you to form the critical mass in peer-review, conference organization, and journal editor board. Have fun and best luck!