



The Changing Culture of Computer Science

An Interview with Marianne Winslett, Part 2

by Richard T. Snodgrass

Editor's Introduction

In this second interview with Marianne Winslett, we explore her diverse experiences as a faculty member, as a prolific international collaborator, and as a director of a research laboratory in Singapore. Our conversation mines the temporal and spatial aspects of the culture of computer science research and development within academia, labs, and startups, delving into recent trends that are fundamentally changing our profession.

The Changing Culture of Computer Science

An Interview with Marianne Winslett, Part 2

by Richard T. Snodgrass

Richard Snodgrass: The [Data and Information Systems \(DAIS\) group at Illinois](#), which you founded 25 years ago, is now quite large, with more than 50 graduate students, two physical laboratories, and an active weekly seminar. How has the culture evolved over those two-plus decades?

Marianne Winslett: One of the big changes over the years is how much more competitive computer science research in general has become. I think that the main cause is that the number of computer science researchers has been growing exponentially—and I emphasize literal, not figurative, exponential growth.

Think of it this way: For our field to retain its current size, each faculty member would have to produce exactly one new Ph.D. researcher over his or her entire career, on average. Instead, we have so many academic offspring! My colleagues at Illinois probably each produce at least a dozen new researchers during their careers, and that branching factor gives an exponent for how fast our field is growing at the Ph.D. level. But our publication venues and funding have not been growing at the same rate. They've grown very slowly, and that means that the competition to get funding and slots in those top venues is also increasing exponentially.

We should have seen this coming, since we are computer scientists and know a thing or two about exponential growth. Yet we don't seem to have really anticipated the consequences of exponential growth and recognized that ours is not sustainable.

Both good things and bad things have happened because of this growth. One good thing is that we have been able to keep up with the expansion of topics in our field. For example, data mining is an extremely important area that didn't exist 30 years ago. Another good thing is that the average quality of accepted papers has gone way up over the past 40 years.

On the other hand, we all know that there are too many safe, incremental delta papers, and not enough systems papers proposing big new half-baked ideas. It is easy for authors to get stuck in a cycle of resubmission, in a culture where reviewers are just looking for the first reason they

can reject a paper because they know the overwhelming odds against acceptance, given the canonical 10-20 percent acceptance rates for top conferences. After rejection, the authors make a few fixes and submit their paper again to the next conference. That wastes everyone's time and slows down scientific progress, and maybe our tentative moves toward a journal culture with no target acceptance rates will help solve that problem.

As a general chair for CIKM'17, I've learned how the increased level of competition has led to new problems with cheating on the refereeing side. These issues have led major conferences like ours to eliminate bidding for papers. Granted, it's a little hard to have bidding anyway when a conference has over a thousand submissions. But even if scale were not an issue, there are rings of PC members who arrange in advance to bid on each other's papers and give those papers positive reviews.

With so many submissions to major conferences, the program committees become enormous, e.g., 600 reviewers, and no one knows all the reviewers. Because many people have similar names, PC invitations sometimes go out to the wrong person, and sometimes that person accepts the invitation and submits questionable reviews. We had to be very careful about this issue at CIKM'17.

Have you seen the websites of fake conferences that imitate prestigious ones? When I searched on Google for the due date for ICDE not too long ago, the top hit was a site that looked at first like the real site for the [IEEE International Conference on Data Engineering](#). If I'd been a junior researcher, I might have been fooled into submitting. Maybe some people who aren't fooled are desperate enough to submit there just so they can list an ICDE 2017 paper on their resumes.

These things are happening because of the competitive pressure to get papers into top venues, which in turn is a side effect of the expansion of our community, which outpaces the growth in good research venues.

RS: What about the globalization of computer science research?

MW: When I first started out 30 years ago, almost all of the hot results were coming out of North America. Over the decades, Europe, Australia, and then Asia became places where you could do really good research and have it published in the top venues. Top conferences started to be held outside of North America, too. You even see best-paper-award-winning work coming out of countries that wouldn't even have had a paper at the conference 20 years ago.

So there's really been a diaspora of high quality research, making it much more global than before.

One big factor in these changes is the web and initiatives like the digital libraries at ACM and IEEE. Without the immediate broadcast of those papers on the internet, a researcher would have to be there in person at the conference to know about the latest research results.

Globalization of research also means that students don't have to come to a top university in the U.S. anymore to learn to do good research. The top U.S. universities won't be affected by this very soon, because each professor has so many students and the best students tend to be concentrated at the top universities, much more so than the best professors. (In fact, for a new PhD who wants to be a professor, I'd say that the number one reason to go to the highest-ranked department that you can is that the quality of students drops off much more quickly than the quality of professors.) But eventually, U.S. departments where most grad students are from overseas may see a major impact on student quality, if overseas students aren't motivated by other considerations like the ease of immigrating to Silicon Valley after graduation.

RS: What about computer science research in academia versus industry?

MW: In data oriented research areas, there's been a brain drain from academia to industry for the past 10 years or so, because industry has all the great data. If you don't have a great data set and some pressing issues associated with it, it's very hard to validate your ideas, or even to know what the important research problems are.

As an editor-in-chief of *ACM Transactions on the Web*, I sift through the best paper winners from a lot of major conferences to find papers that we might want to see in extended form in our journal. More and more of the best paper award winners are coming from industry. That was a surprise!

Many major conferences still have something called an "industry track." Industry tracks were originally intended as a way to allow people in industry to get a paper accepted and attend the conference, because those authors couldn't compete in the research track. Now I see many academic authors in the industry track, and many industry authors in the research track. In a world where industry authors routinely win best paper prizes, it's time to rethink those industry tracks! At CIKM'17, we're experimenting with turning ours into a case studies track, where practitioners can learn about the experiences of early adopters.

RS: You mentioned you joined the [Advanced Digital Sciences Center \(ADSC\)](#) in Singapore in 2009, shortly after it was created. How did that whole organization come about? Singapore is quite far from Illinois.

MW: I think it's fair to say that I was personally involved with all of our early hires. After a while, I was able to step back and be more focused on mentoring the young people we'd already hired, while they hired even younger people to work for them. I've been told I'm quite good at that. [Smiles]

The premise behind ADSC is that Illinois faculty members lead research projects going on 12,000 miles away in Singapore. We originally thought that the researchers in Singapore would be Illinois PhD students. But it didn't turn out that way, which is good because in hindsight I recognize that the original plan would not have worked well for research at a distance. Most Ph.D. students require frequent interaction. You can't really talk to a PhD student once a week on Skype and have them make speedy progress on a project.

We came up with a different model for research at a distance that works quite well. Instead of having Ph.D. students 13 time zones away from their professors, we hired people who already had PhDs, and created a ladder of postdoc and research scientist positions so that they could grow as researchers while still working for ADSC. We refer to those people as our PhDs. When professors hear that term, they think we are talking about PhD students, because that's the model that professors are used to. Instead, we hired people who already had a PhD, so that they had enough prior experience to be fairly independent and to rely on the project's professor back at UI (University of Illinois) mainly for strategic advice.

For research at a distance, it's not enough just to hire excellent recent PhD graduates, for several reasons. First, if you hire the best people you can find, you end up with lots of chiefs and no Indians, which is not a very effective model for getting things done. Second, if these recent PhDs really are going to develop as researchers while they work for ADSC, they need to practice hiring people and supervising them. Third, once they get a bit more senior, they need to practice the art of convincing people that their ideas are worth pursuing, by writing grant proposals. Eventually they will write successful proposals, and then the funds need to be used to hire additional people.

Sometimes ADSC's PhDs hire student interns who come for a summer or semester. But more often they hire full time, research oriented, software engineers. Usually these engineers recently finished an undergraduate degree and want to build a research record before applying to PhD programs at top schools. They know that they need to know how to do research before they even apply to grad school; to get admitted, they need to already have publications. It turns

out there's a job market for that type of engineer, and we were able to find those people and hire them. We found that a one-to-one ratio of PhDs to research engineers often works quite well, for a new PhD. As people get more senior, they tend to get grant money and hire additional engineers or interns.

The PhDs might talk to a UI faculty member over Skype once a week or so, to get advice in choosing directions or overcoming particular problems. Beyond that, the PhDs and their engineers are really quite independent in carrying out the research. And in hindsight, I can say that that approach would not work with most PhD students.

We found that the model has other advantages. The engineers are huge productivity multipliers for the PhDs they work for. Although the engineers might be the same age and have almost the same background as first-year graduate students, the engineers are devoted full time to the project, without the stresses and distractions of classes and exams. Compared to a new assistant professor with one grad student, probably our PhDs got a lot more research done. Certainly they were able to build bigger systems.

Another nice aspect is that our engineers tend to stay with us for several years, to build really a strong resume before applying to grad school or going off to a startup. Just like a grad student, each year the engineer knows more and is more productive than the year before. When they finally leave, it's both a sad and happy moment for their PhD, just like when a student graduates.

Most of our PhDs come from degree programs where the students are expected to do a postdoc before finding a permanent position. When they leave ADSC, some become professors, some join startups, and some join multinational companies. Most stay in Singapore, but some return to their native countries.

RS: Were your hires from Asia or from Illinois?

MW: Singapore is a nation of immigrants, so we hire people from everywhere. We draw from Asia the most, but also the Middle East, Europe, the U.S., and even Africa. So it's quite a mix. Some are attracted by the adventure of living in Asia, while others want to do Illinois-style research but need to be in Asia for family reasons.

RS: Did the professors at Illinois also travel to Singapore?

MW: Yes, typically twice a year for two or three weeks. In the beginning, it's helpful to spend more time there because you need to get to know the country. On the surface, Singapore's research system and research culture is deceptively similar to that of the U.S., which makes it easy to assume that the systems are the same all the way down. But there are a lot of differences below the surface.

Funding in Singapore works a bit differently than in the U.S. For many U.S. government agencies that don't have a research mission, their Singapore equivalents do have a research mission, as you can see by the facts that ADSC's original core funding came from the Ministry of Trade and Industry (equivalent to the U.S. Department of Commerce), and our group now has a grant from the Buildings and Construction Authority. For most grants from any agency, the overhead is a fixed 20 percent, the Principle Investigator (PI) can't pay his or her own salary out of the grant, and there may be salary caps for people hired under the grant. No-cost extensions are rare, which means that PIs may have to hire in a hurry or give back large sums at the end of the project. Singapore is such a small country that conflicts of interest are handled very differently than in the US. And so on.

Funding is just one example. At first blush, you don't see these differences from the U.S. way of doing things, but if you are really going to participate in the system, you have to understand them.

RS: Was there a commercial component to the lab?

MW: After a while, our technical staff caught start-up fever, and there certainly has been a lot of effort in that direction. I have my eye on four or so of the recent startups that I think are looking quite promising. We also do technology transfer by licensing, and by other companies hiring our people. It's easy to overlook that last method of tech transfer, but I think it's an important part of ADSC's educational mission. Young people build their research skills with us, and their next employer will benefit from that.

RS: What were the broad outlines of your role as director?

MW: You are who you hire, so the single most important thing I did was to help choose who to hire. That was quite tricky at first, because the applicants completed their PhDs in institutions

all over the world, in degree programs of different lengths. People graduate with very different resumes, depending on where they did their PhD. For example, back at Illinois, I was used to seeing resumes from fifth and sixth year grad students in strong U.S. programs. In contrast, a job-hunting PhD student from a top school in the U.K. with a three-year PhD program tends to have the resume of a third-year grad student at Illinois. I had to recalibrate the meaning of a PhD and realize that all these half-baked (by my standards) students in other countries were expected to finish baking in a postdoc position. They weren't going straight to permanent positions like the US students were. It's a different model for how to become an independent researcher—not a better or worse model, just different. We had to learn to sniff out the best technical talent, no matter what model they had been educated under, and convince them to join this unconventional research center.

When a new PhD joins ADSC, they face many of the same issues as a new assistant professor at Illinois. They need to understand the local opportunities and collaborations and industry that might want to partner with them and give them good data. They need to settle in and meet the right people and have mentoring so that they're moving along their chosen career path. We had a real diversity of career paths because quite a few of the PhDs were interested in becoming professors in their home countries, and different countries have different ways of deciding who to hire as a professor. That meant that different staff members had to build their resumes in different ways and in different stages. Some members of our staff were very interested in being entrepreneurs. Others wanted to join industrial research labs. And then, most of our engineers wanted to go off and get their PhDs in the U.S., Europe, or Asia.

As you can see, we had a real mix of career paths, and each person needed appropriate mentoring to ensure that they were meeting the right people, building the right resume, and applying to the right opportunities. I spent a lot of time doing that, and it's what I miss most about the job.

It also fell to me to communicate our successes to the outside world. I had no expertise in that area, because Illinois is famously bad at self-promotion. We are called the best-kept secret in the U.S., and ADSC eventually became the best-kept secret in Singapore, though I certainly tried to get the story out. ADSC's staff eventually received dozens of best-paper awards, and I never tire of mentioning that fact.

I was fortunate to have Professor Jesse Delia, who had been the provost at Illinois, taking care of everything at ADSC outside of the research program. That meant that I didn't have to worry about choosing health insurance providers or banks, or establishing policies and procedures, or figuring out how to get a Singapore grant proposal through the Illinois IRB or grants office.

Jesse dealt with the lawyers, export control, Office of Technology Management, financial matters, all of those bureaucracies. If I had had to do that, I would not have had the time needed to build the research side by helping people put together projects and get hired and oriented. Jesse was instrumental to our success.

RS: Were you given carte blanche, or were there certain directions that A*STAR or Illinois wanted the lab to go into?

MW: Our initial core funding was from A*STAR to work on the Human Sixth Sense Program, which included research on interactive digital media and its supporting infrastructure. The second round of core funding added a cybersecurity component. Our third round of core funding is from Singapore's National Research Foundation, for work on cybersecurity and analytics for the power grid. Illinois is a big university, so it is easy for us to be flexible.

While ADSC's focus has always included excellence in research, other aspects of our charter have changed over the years: how important external grants are, what kinds of collaborations are most desirable, the relative importance of startups and licensing, which populations to target in hiring, which aspects of workforce development matter most, and so forth. Singapore changes rapidly, so it is natural for these aspects to evolve also.

RS: How did the Singapore culture affect things at ADSC?

MW: Many parts of Asia, including Singapore, have a strong Confucian influence. My favorite thing about Confucian cultures is that if the person at the top wants something to happen, then it just happens. Once the person at the top blesses your idea, you don't have to work your way through layers and layers of bureaucracy to make it happen. That is glorious. And it's certainly not how Illinois operates. There, every professor is an independent-minded little chief, so Illinois is probably more bottom-up than top-down. In the west, managing professors is often compared to herding cats! But since Singapore is pervasively top-down, people there tend to assume that Illinois is the same way.

Another aspect of Confucian culture is great respect for seniority. In general, for important tasks, senior researchers are preferred over junior ones. White hair is ideal! The US isn't like that, and so I think our young ADSC staff from Asia probably felt quite empowered.

Singaporeans tend to be very risk averse, but in general, the young people who joined us are not; otherwise they wouldn't have joined a new research center with no track record. So, for example, they were much more willing to take advantage of startup opportunities than the average Singaporean would be.

RS: Research labs in the U.S. have evolved considerably over the past few decades. For example, Bell Labs was highly prominent, but then it was split up. Some major computing companies have no research activities at all. Given your experience, what models for labs outside of computer science university departments make sense for the near future?

MW: There's been a recent trend for large companies to buy innovation by purchasing startups. In many ways, that makes a lot of sense because it's really tough to let innovation flow from the bottom levels of a commercial organization. It's natural for any bureaucracy to resist change. Often resistance is fueled by the possibility that a new innovation might hurt the profits of an existing product line or line of service. So it could be quite reasonable for someone at the top to decide to buy innovation from outside, by acquiring a startup, getting the rights to their technology, and doing a good job of integrating it into the company. Often startups are founded in the hope of being acquired this way.

But it's not the only model out there. For example, Tesla has been very innovative, quick and nimble, and it has 33,000 employees already. Instead of buying innovation, Tesla is [giving it away](#). Many Google products emerged from the one-day a week that Google employees could work on their own projects. Until 2014, Facebook believed in moving fast and breaking things. Apple built a lot of innovative things in its time. But Google and Facebook don't work that way anymore, and the jury is out on Apple's fate without Steve Jobs. So perhaps when an organization reaches a certain size and age, it ossifies into something resistant to innovation, unless someone very special is at the top.

RS: Do you see that happening also in Singapore?

MW: Definitely yes. But remember that Singapore's culture is risk averse, and so there are relatively few entrepreneurs compared to the U.S. One factor is that in some parts of Asia, people believe that if you ever fail at anything, then you're going to be forever branded as a failure for the rest of your life. Ninety percent of high-tech startups are going to fail, but traditionally in some parts of Asia, having a company fail is considered a disgrace for the

founder's entire family. Compared to that, the U.S. is the land of second chances. You can always go back to college. You can always wipe out your debt with a bankruptcy and start over. You can always start a new company, and so on.

Singapore is not like that. My favorite example is that there is a national exam, the Primary School Leaving Examination, given to kids when they're around 12 years old. To some extent, the outcome of that exam determines the course of the rest of a person's life. The government realizes this is ridiculous and has been trying to move away from putting so much weight on this exam. Perhaps this system made sense back in the 1960s when Singapore was a new nation and needed to identify the cream of the crop and train them intensively to help with nation building. But it really doesn't make sense now that Singapore is well-established as a nation, suffers from an extremely low birth rate and needs to get the most it can out of each of its citizens. There's a reason why Singapore kids score so highly on international scholastic achievement tests, often coming in first or second among all nations, to a surprising degree. All that their parents think about from the time the kids are born is how they're going to do on that single test. So you can see how that culture would discourage risk taking.

RS: In contrast, the Silicon Valley culture is “fail fast and fail often.”

MW: Yes. There's a lot of money floating around in some parts of China because of all of the manufacturing successes that China has enjoyed, and both government and private individuals are looking for good ways to invest that money, whether it's in education or real estate or ongoing businesses or startups. In particular, in Beijing, a lot of creative energy has gone into startups in the last ten years or so. More recently there has been a lot of hot money for startups in Shanghai and even Shenzhen. But probably the largest concentration of them is still in Beijing.

RS: It's a joy to be able to interview you, in part because you are a very famous interviewer. You've interviewed, I think, almost 70 people accumulating more than 30 hours of interviews. We are talking over the phone, but you actually videotape those interviews. How much work is involved in doing each interview?

MW: There's a lot of work before and after the interview, but not so much during it. Beforehand, I have to figure out what questions to ask the person, which requires research. I ask their friends and close colleagues to suggest questions to ask. I compile those ideas, refine

them, and then share the resulting questions with the person being interviewed, because knowing what I'm planning to ask can be reassuring. This is not "60 Minutes," and I don't want to make people nervous.

We film using a video camera, but we also use bright lights and a more expensive microphone. Otherwise, the footage comes out very dark and with an astonishing amount of background noise. Sometimes I run the equipment, but whenever possible I have a volunteer cameraman, because it's hard to think about the conversation while also worrying about background noise, battery failure, whether the subject has squirmed out of the frame, and so forth.

During the actual interview, I sit off to the side and ask the questions. Afterwards we edit the interview down to about half an hour. We always record more than half an hour of material, and then snip and cut to make it the right length because there are always surprises during the interview, and some parts turn out to be more interesting than others. An expert volunteer edits the video and produces the podcast.

If a print version is also needed, that's a lot more work, because spoken and written English are almost two different languages. I edited the written versions for the first few years, and then expert volunteer Vanessa Braganholo took that over, for which I am extremely grateful.

RS: Are you going to continue your interview series?

MW: Back when we started, YouTube as we know it today didn't exist, so people who couldn't come to the major conferences never got to see the leading researchers in our field. They never got to hear famous people talk at length about where they thought the field was going and about important technical issues or the ups and downs of their careers. Back then, the interviews were the only way that people living so far away could build that connection.

But the world has changed, and more and more material on YouTube will let you see famous researchers giving conference keynotes and other presentations. Admittedly, I like to ask questions about an entire career rather than a single technical topic. But at least everyone can see and hear keynote speakers on YouTube, and get a sense of what they're like. We didn't have that at all when we began. Another change is that there are more and more industry interviews on YouTube, which I think is great.

Putting it all together, I think the interviews don't matter as much as when we first started. To decide whether to continue, a good test would be to see how thoroughly the people we consider to be the leaders of the field are represented on YouTube. For example, Pat Selinger is

all over YouTube, and people interview her all the time. But what about everybody else? I honestly don't know.

RS: Actually, I see the focus on the person's entire career as the continuing value of your interviews even in this YouTube culture. I personally don't know of other sources for the kind of in-depth, probing, and retrospective interviews that you do so well.

MW: I'm sure there is someone else in our community who could do a great job! Prospective volunteers should get in touch with Vanessa and me by email. Or if you are from a different research community, then reach out to your SIG. Our series is sponsored by SIGMOD, and I'm sure other SIGs would be interested in having similar series.

RS: Some people say we should get rid of PowerPoint, especially in the classroom. How can researchers and practitioners become better writers and presenters?

MW: I was very influenced by Michael Alley from Virginia Tech, who wrote *The Craft of Scientific Presentations*. I think his book should be required reading for anyone who inflicts PowerPoint presentations on an audience! (For a quick summary of the most important points from his book, check out his [talk](#) at MIT.)

Michael did a scientific study about how well students learn with different styles of PowerPoint. I thought that if he went to the trouble to do a controlled study, I should take his findings very seriously. Many of the things he found out are so obvious in hindsight that it is hard to understand why they aren't already standard practice. Let me give some examples.

Michael found that students do better if the main point of a slide is used as its title. For example, "Performance Results" conveys no takeaway message. A better slide title would be "Quicksort is the fastest sort method for big randomly ordered data sets," because everyone will get that message so much better!

Another of Michael's rules of thumb is not to put a lot of text on a slide, because people can't listen to you and read your words at the same time. Those two activities use overlapping parts of the brain, so students have to pick between one and the other. But if you put a picture on the slide as a visual aid, then people can look at the picture and listen to you at the same time.

Another of his tips is never to end a talk with a slide that says "Questions?", as so many people do. Instead the last slide should be the one that summarizes the takeaways of the talk you just

gave, so people can stare at it and absorb it while you answer questions. Otherwise, it's a wasted opportunity to reinforce your takeaway message.

The Craft of Scientific Presentation has many other great tips. After we read the book, my grad students and I switched to a much more visual focus in our slides. What in the olden days would have been the bullet points of each slide, that's what we say out loud now, and the slide has a helpful image for people to look at while we say it.

RS: I'm wondering if it might be effective to take a PowerPoint talk that doesn't follow any of those rules and make another version of the slides with pictures and the takeaway point at the top. And then use the new version during the talk, but also hand out the original bullet points as talk notes afterwards.

MW: That's how I put a new talk together. It's much quicker to do the traditional bullet point format because you can figure out what you want to say pretty quickly, but getting the right image to go with it often takes considerably longer. On the other hand, the effort you invest in coming up with the right image to convey an idea does pay off in the long run. You can reuse that image in all of your talks about that topic, and maybe in your papers too. Often you can use the same image to make more than one point, depending on what your presentation is about; so the slide title might change while the image stays the same, e.g., "RadixSort is great, when you can use it." There is a lot of payoff once you get the right image.

RS: That makes sense. The right image can capture the essence. Thank you so much for spending time with me this afternoon.

This interview has been condensed and edited for clarity.

Biographies

Richard Snodgrass is a member of the *Ubiquity* Editorial Board. He is also a Professor of Computer Science at the University of Arizona, working in the areas of ergalics and temporal databases.

Since 1987, Marianne Winslett has been a professor in the Department of Computer Science at the University of Illinois at Urbana-Champaign, where she founded the DAIS (Data and Information Systems) research group. After four years as the director of Illinois's research center in Singapore, the Advanced Digital Sciences Center, she returned to the U.S. in 2013. She is an ACM Fellow and the recipient of an ACM SIGMOD Contributions Award, as well as a Presidential Young Investigator Award from the U.S. National Science Foundation. Her research specialties include information security and management of scientific data. She holds a Ph.D. from Stanford University. Winslett is the former vice-chair of ACM SIGMOD and the co-editor-in-chief of ACM *Transactions on the Web*, and has served on the editorial boards of ACM *Transactions on Database Systems*, *IEEE Transactions on Knowledge and Data Engineering*, *ACM Transactions on Information and Systems Security*, the *Very Large Data Bases Journal*, and *ACM Transactions on the Web*. She has received two best paper awards for research on managing regulatory compliance data (VLDB, SSS), one best paper award for research on analyzing browser extensions to detect security vulnerabilities (USENIX Security), and one for keyword search (ICDE).

DOI: 10.1145/ 3158418