Editorial: Early Career Researchers in Embedded Computing

As new professionals enter their independent careers as researchers and engineers in the industries, or as faculty at research universities, they are always confronted with a multitude of choices regarding their career path. Strategies for getting projects funded, count versus depth in publication, and the journal-versus-conference dilemma are quite common. This year at the ACM/IEEE Design Automation conference in San Francisco, Aviral Shrivastava and Andreas Gerstlauer, industry researcher Sherie Taylor, and IEEE CEDA representative Shishpal Rawat, who also spent his career in the industry, organized a very relevant 1-day workshop for new faculty and researchers from industry and government titled “Early Career Workshop.” Even though it was in the context of a design automation career, embedded computing is important in multiple areas, and I feel that it is relevant to dedicate this editorial on the lessons learnt from this workshop.

The first panel on “Academic Career Challenges,” chaired by Todd Austin, reflected on the specific career challenges he and four other faculty members faced while working towards establishing their careers. The most interesting point made concerned how much to write in the early days of one’s career. We all fall into the trap of trying to make all of the deadlines in our fields when we are new and eager to go to all the top conferences, and we often work ourselves and our students to death to make those deadlines. If he were to start afresh, he said that he would not do this. It was pointed out that the depth and impact of your papers are more important than your numbers. While numbers do play a small role for tenure and promotion committees, demonstrable impact in terms of follow up work by others, solving problems that have industrial or other impacts in terms of translational potential, and possibly solving a long-standing problem that takes time and patience to solve, can help one’s career far more. I think this is a very important piece of advice for our early career researchers. We have all been a victim of this numbers game, and even at advanced stage in our careers, we often have to do this so our students can graduate. But in the long run, unless the work makes a large impact, this strategy does not help anyone.

The second panel, “Making Publications Work for You,” was chaired by Vijaykrishnan Narayanan, and I was part of this panel. The panelists were all experienced editors or editors-in-chief of various ACM and IEEE journals or chairs of flagship conferences in the field. The issue here is whether to publish in journals or in conferences and whether publishing a conference paper with the extra “30%” material makes any sense. In this context, one of the most problematic issues is the exponential rise in the number of papers submitted versus the number of researchers available to provide in-depth reviews for these papers. As editors, we always face this conundrum. Everyone wants to submit papers and get the papers published as soon as possible, but the same people reject these review requests, citing their extreme engagement in research, teaching, or other deadline-driven activities. This is not a sustainable publication model.

Some time ago, I had suggested that the ACM start a point system. If someone reviews a paper, and does his/her job on time, he/she will be given some points. Unless a researcher has enough points, his/her submission to journals/conferences will not be reviewed and then sent back. For early career researchers, we have to certainly provide some initial startup points so that they can start publishing. We can create a blockchain across ACM/IEEE to keep track of these points, and as

a paper is submitted and reviewed, some points will be paid from a wallet. I do not know whether my proposal was taken seriously, but I do not see any other way to achieve a sustainable system.

The other problem is reviewer delinquency. We often have reviewers taking up six months or more, and then they do a quick and loose job in reviewing the paper, which affects the quality of the publication. I have had reviewers submitting after more than 12 months. In a rapidly developing field of technology, this dates the work and unnecessarily penalizes the author. In the early days when there was not such a deluge of journals and conferences, the system was in a kind of Nash equilibrium. Everyone approached their reviews as a sacrosanct duty towards the community, because if they did not, the system would halt, and they themselves would not be able to publish. Now, the equilibrium is no longer maintained, and we really need to solve this problem.

The ESWEEK conferences have already chosen a “journal first–journal only” publication model that helps reduce the number of reviews required, but there is still a problem throughout. In addition, Ph.D. defenses often get delayed because reviews are delayed, and the tenure clock starts ticking, making faculty nervous... we need to solve this problem. So, the lesson is that if you want your work published fast, you need to review as often as you write or more. I hope some game theory person takes up this problem seriously and come up with the correct strategy and balance.

The other question that generated a lot of discussion was from a young industry researcher about incremental improvement of an existing solution. His question was how to publish such work or whether to publish. If the work is not enough to warrant a publication, it might still be patentable. Patents are considered publications, so in my view researchers should consider that route. My fellow panelists were not sure that all that seems incremental might not be incremental, so it is a dilemma. I believe more social networking of researchers can solve this dilemma: One can discuss his or her work and seek out the views of their colleagues, especially those who are more experienced one, before submitting a paper to journals or conferences and engaging multiple reviewers on something that might not make it. Vijaykrishnan Narayanan and others also reiterated the earlier the point regarding– seeking depth and fewer but better publications.

The rest of the panels were on funding opportunities for young professionals and industry engagements, and I wish more early career professionals had attended. These panels are very relevant and provide an opportunity to clarify many doubts that their local support system might not provide.

Overall, I think more such workshops in conjunction with other conferences around the world could really create a support system for our early career professionals in their respective fields. It must be field specific, as different fields have different dynamics of publication, funding, and industry engagement. I congratulate the organizers for serving the community by organizing such an interesting workshop.

Volume 17, Issue 4, of ACM *Transactions on Embedded Computing Systems* includes articles on subjects ranging from cryptographic hardware, optimizing GPU computation, battery-aware control application in automobiles, application of RFID in trajectory detection, to architecture. Given that the ACM publications board has approved six issues per year, we are able to publish with eight papers, as there are two more issues coming out for Volume 17, and I am happy to say that we do not have any accepted papers being delayed in the publication pipeline. The review pipeline, however, needs improvement—and I urge all readers to kindly take review work seriously and be timely with submitting their reviews, as the reviewers are themselves authors as well.

Sandeep K. Shukla

Indian Institute of Technology Kanpur

July 14, 2018

Editor-in-Chief