

Viewpoint

Technology, Conferences and Community

Jonathan Grudin
Microsoft Research

Commentaries in the January, May, and December 2009 issues of *Communications* address a perceived crisis in Systems and other Computer Science fields [1, 2, 3, 4]. The discussion covers our shift from a traditional emphasis on journals to the current focus on conferences, and the challenges of conference reviewing, but at its heart is our sense of community. One commentary speaks of a “death spiral” of diminishing participation [3]. This Viewpoint considers how and why these changes occurred.

In a nutshell, the commentaries note that a focus on conference publication has led to deadline-driven short-term research at the expense of journal publication, a reviewing burden that can drive off prominent researchers, and high rejection rates that favor cautious incremental results over innovative work. Some of the commentators suggest that we return to emphasizing journal publication. To understand whether this is possible, and I doubt it is, we must understand why Computer Science in the United States shifted to conference publication in the first place. It was not simply that Computer Science requires quick dissemination of results: Conferences did not become prominent in Europe or Asia, or in other competitive, quickly-evolving fields such as Neuroscience or Physics.

Technology and a Professional Organization Drove the Move to Conferences

By the early 1980s, the availability of text editing or word processing among computer scientists enabled the relatively inexpensive production of decent-looking proceedings prior to a conference. This was something new. Anticipating that libraries might shelve proceedings, ACM printed many more copies than conferences needed, at a low incremental cost. ACM also made them available by mail-order after a conference at a very low price. Papers in ACM conferences were thus widely distributed and effectively archival. *These are the two features that motivated the creation of journals centuries earlier.*

Proceedings in Europe and Asia rarely had after-conference distribution, so to be archived, work there had to progress to journal publication. The shift to a conference focus did not occur. In 2004, a prominent UK researcher wrote about the CHI Conference:

HCI's love of conferences is a fluke of history. We all know this. CS in general has suffered from it, but is steadily moving away. CHI however digs in, with more and more death rattles such as CHI Letters. Being conference centred is bad for any field: bad for its archival material, bad for its conferences, and worst of all, really bad for the respect that we command with other communities. SIGCHI needs to move away from bolstering up conference publications. It needs to use journals for journal stuff and conferences for conference stuff. [5]

He was wrong about the direction of Computer Science, and at least premature in diagnosing CHI's expiration. The point, though, is that he saw the problem as an American problem, affecting CHI but not European HCI.

Knock-on Effects

This change in the complex ecology of scholarly communication was followed by a slow sequence of adjustments. ACM and IEEE had considered conference papers to be ephemeral, and expressly allowed

verbatim or minimally-revised republication in journals and transactions. With proceedings effectively archived even before digital libraries arrived, this policy was formally ended early in the 1990s.

A significant consequence is that it is increasingly difficult to evolve conference papers into journal articles. Publishers, editors, and reviewers expect considerable new work, even new data, to avoid a charge of “self-plagiarism.” Republishing the same work is undesirable, but we have inhibited the use of review and revision cycles to clean up conference papers, expand their literature reviews, and engage in the deeper discussions that some feel are being lost.

The pattern extends beyond Systems. I edited *ACM Transactions on Computer-Human Interaction* and serve on the editorial boards of *Human-Computer Interaction*, *Interacting with Computers*, and *ACM Computing Surveys*. By my estimation, no more than 15% of the work published in highly selective HCI conferences later appears in journals. Journal publication is not a prerequisite for being hired into leading research universities. Today, the major U.S. HCI journals mostly publish work from Europe and Asia, where conferences are less central.

Now let’s consider reviewing, a primary focus of discussion, before turning to the impact of these changes on our sense of community.

Conference Selectivity and Effects on Reviewing

In other fields, journals focus on identifying and improving research quality; large conferences focus on community-building and community maintenance; and workshops or small conferences focus on member support through specialist discussions of work in progress. This reflects Joseph McGrath’s division of group activities into those focused on production, team health, and member support [6].

When conferences became archival, it was natural to focus on quality and selectivity. Even with authors preparing camera-ready copy, the expense of producing a proceedings was proportional to its page count. Libraries sales were a goal prior to the emergence of digital libraries in the late 1990s. Libraries were more likely to shelve thinner proceedings, and needed to be convinced that the work had lasting value. These pressures drove down conference acceptance rates. In my field they dropped from almost 50% to 15% before settling in a range, 20-25%, that is acceptably selective to academic colleagues yet not brutally discouraging to authors, we hope.

But it *is* discouraging to have submissions rejected. I know few if any people who submit with no hope of acceptance. In most fields, conferences accept work in progress. It is also discouraging when we see a paper presented and immortalized in the digital library that seems less worthy than a paper that was rejected. Review processes are noisy, and more so as the reviewer pool expands to include graduate students and others. Multidisciplinary fields, with diverse methodologies and priorities, deliver even more random outcomes.

Previous commentaries emphasized that caution and incrementalism fare better than innovation and significance in conference assessments. An incremental advance has a methodology, a literature review, and a rationale for publication that were bullet-proofed in the papers it builds on. We try to channel papers to the most expert reviewers in an area, but to them incremental advances loom larger than they will to others. With pressure to reject ~75% and differing views of what constitutes significant work, the minor flaws or literature omissions that inevitably accompany novel work become grounds for exclusion. And in a zero-sum game where conference publication leads to academic advancement, a novel paper can be a competitive threat to people and paradigms [1], echoing concerns about journal conservatism in other fields.

Birman and Schneider [3] describe the risk of a “death spiral” when senior people cease to review. Although engaging young researchers as reviewers is great, many feel more comfortable identifying minor flaws and less comfortable in declaring that work is more or less important.

Every year, conferences in my area adjust the review process. But significant change is elusive, given the forces described above.

| SIG | 1990 | 2000 | 2010 |
|-----------|---------|--------|--------|
| PLAN | 12,335 | 4,362 | 2,323 |
| GRAPH | 11,811 | 6,298 | 7,216 |
| SOFT | 10,824 | 3,313 | 2,489 |
| ART | 8,955 | 1,917 | 1,123 |
| OPS | 6,801 | 2,356 | 1,828 |
| CHI | 5,023 | 4,950 | 4,415 |
| ARCH | 4,685 | 1,730 | 1,384 |
| ADA | 4,089 | 685 | 292 |
| MOD | 3,952 | 2,941 | 1,922 |
| MIS | 3,082 | 755 | 497 |
| (all 30+) | 103,489 | 47,042 | 41,008 |

Table 1. Membership in the top 10 Special Interest Groups of 1990. Today only two of 34 SIGs have 3000 members.

Impact on Community

A leading neuroscientist friend described their annual meeting as a “must-attend” event that is “where people find out what is going on.” There are 15,000 presentations and 30,000 attendees. The quality bar is low. It is a community-building effort in a journal-oriented field.

In contrast, despite tremendous growth in many CS specializations, attendance at many of our conferences peaked or plateaued long ago. So has SIG membership. Conferences proliferate, dispersing organizational effort and the literature, reducing a sense of larger community.

In my field, CHI once had many vibrant communication channels—a highly regarded newsletter, an interactive email discussion list, passionate debates in the halls and business meetings of conferences,

discussants for paper sessions, and in the late 1990s an active Web forum. All of them disappeared. The CHI conference gets more submissions, but attendance peaked years ago. When a small, relatively polished subset of work is accepted, what is there to confer about?

High rejection rates undermine community in several ways. People don't retain quite the same warm feeling when their work is rejected. Without a paper to give, some do not get funding to attend. Rejected work is revised and submitted to other conferences, feeding conference proliferation, diverting travel funding, and dispersing volunteer efforts in conference management and reviewing. In addition, high selectivity makes it difficult for people in related fields to break in—especially researchers from journal-oriented fields who are not used to polishing conference submissions to our level.

A further consequence is that computer scientists do not develop the skills needed to navigate large, community-building conferences. At our conferences, paper quality is relatively uniform and the number of parallel sessions small, so we can quickly choose what to attend. In contrast, randomly sampling sessions at a huge conference with 80% acceptance leads us to conclude that it is a junk conference. Yet with a couple hours' preparation, combing the many parallel sessions for topics of particular interest, speakers of recognized esteem, and best paper nominations, and then planning meetings during some sessions, one can easily have as good an experience as at a selective conference. But it took me a few tries to discover this.

Courtesy of Moore's Law, our field enjoys a constant flow of novelty. If existing venues cannot rapidly shift to accommodate new directions, other outlets will appear. But there can be abrupt discontinuities. Our premiere conference for many years, the National Computer Conference, collapsed suddenly two decades ago, bringing down AFIPS, then the parent organization of ACM and IEEE. Over half of all ACM Turing Award winners published in the AFIPS conferences. Most of those published single-authored papers. Yet the proceedings disappeared, until they were recently added to the ACM Digital Library. The field moved on—renewal is part of our heritage. But perhaps we can smooth the process.

Possible Directions

Having turned our conferences into journals, we must find new ways to promote community-building. Rolling back the clock to the good old heyday of journals, ignoring changes wrought by technology and time, is hardly plausible. But technology, having created problems, may help solve them.

With storage costs evaporating, we could separate quality-determination from participation by accepting most conference submissions for presentation and online access, while distinguishing ~25% as "Best Paper Nominations." Making a major conference more inclusive could pull participation back from spin-off conferences.

A more radical possibility is inspired by the revision history and discussion pages of Wikipedia articles. Authors could maintain the history of a project as it progresses through workshop, conference, and journal or other higher-level accreditation processes. Challenges would have to be overcome, but such an approach might ameliorate reviewer load and multiple publication burdens—or might not.

We are probably not approaching the bottom of a "death spiral." But when AFIPS and the National Computer Conference collapsed, the transition from profitable success to catastrophe was remarkably rapid. Let's continue this discussion and keep history from repeating.

1 Crowcroft, J., Keshav, S. and McKeown, N. Scaling the academic publication process to Internet scale. *Commun. ACM* 52, 1 (Jan. 2009), 27–30.

- 2 Vardi, M. Y. Conferences vs. journals in computing research. *Commun. ACM* 52, 5 (May 2009), 5.
- 3 Birman, K. and Schneider, F. B. Program committee overload in systems. *Commun. ACM* 52, 5 (May 2009), 34--37.
- 4 Wing, J.M. and Guzdial, M. CS woes: deadline-driven research, academic inequality. *Commun. ACM* 52, 12 (Dec. 2009), 8--9.
- 5 Gilbert Cockton, email communication, 22 Jan. 2004.
- 6 McGrath, J. E. (1991). Time, interaction, and performance (TIP): A theory of groups. *Small group research*, 22, 2, 1991, 147—174.

| SIG | 1990 | 2000 | 2009 |
|-----------|---------|-------|--------|
| PLAN | 12,335 | 4362 | 2,557 |
| GRAPH | 11,811 | 6298 | 8,474 |
| SOFT | 10,824 | 3313 | 2,637 |
| ART | 8,955 | 1917 | 1,202 |
| OPS | 6,801 | 2356 | 1,600 |
| CHI | 5,023 | 4950 | 4,845 |
| ARCH | 4,685 | 1730 | 1,363 |
| ADA | 4,089 | 685 | 317 |
| MOD | 3,952 | 2941 | 1,952 |
| MIS | 3,082 | 755 | 543 |
| (all 30+) | 103,489 | 47042 | 42,969 |

Table 1. Membership in the top 10
Special Interest Groups of 1990.
Today only two of 34 SIGs have
3000 members.