Conferences, Community, and Technology: Avoiding a Crisis

Jonathan Grudin
Microsoft Research
One Microsoft Way
Redmond, WA USA
+1 425 706 0784
jgrudin@microsoft.com

ABSTRACT
Computer Science in North America has embarked on a course unique in academic scholarship. It has turned conferences into repositories of polished work, little of which ever evolves into journal articles. Senior researchers feel that the conferences are in crisis. I consider the origins and consequences of the shift to conferences, concluding that it has led to an evolutionary cul-de-sac that the Information field would do well to avoid. The crisis is described as centered on reviewing, but it is at heart a crisis of community.

Categories and Subject Descriptors
K.2. History of Computing
K.7.2. The Computing Profession: Organizations.

General Terms
Management, Design, Human Factors.

Keywords
Community, Disciplines, Conferences.

1. THE ECOLOGY OF SCIENTIFIC AND TECHNICAL COMMUNICATION
It is often observed that new technologies are inserted into existing processes to replace older technologies. "Design a word processor with the typewriter as a model," we were advised. Then, over time, the processes and the technologies are restructured. The goal is to improve upon the status quo, but long-term consequences can be unpredictable.

My topic is the effects of new technologies on the processes of research dissemination. We inserted transformational technologies into a complex ecology of books, journals, and conferences with minimal reflection on what would eventually replace the "iron horse" stage. Email replaces informal conversations and phone calls, word processing is used to prepare articles, PowerPoint replaces slides, authors put articles on the Web or blog their findings. We know this is just the beginning. What are the destinations, and how will we reach them?

Between 1997 and 2003, when I was Editor in Chief of *ACM Transactions on Computer-Human Interaction*, the Internet and the Web were perceived to threaten the existing business models of publishers. My 2004 essay *Crossing the Divide* outlined goals of scholarly communication shown in Table 1 [5]. The goals frequently conflict. Careful reviewing trades off against rapid dissemination. The goal of archiving all useful results can run up against page and cost constraints. A submission that an author considers to be an original contribution, an editor may declare to be out of scope. This creates a complex force field that drives books, journals, conferences, and workshops to different niches, each niche representing a different weighting of priorities. The resulting landscape varies across disciplines, and within a discipline can vary by country or continent, marked by differences in the nature of the scholarly activities, the approaches to assessing contributions, legacy practices and traditions, and the state of digital technology integration.

Digital technologies have affected scholarship in diverse fields. Physics and Mathematics are frequently-cited examples. My focus is on lessons for Information from the Computer Science experience

2. A CRISIS ENGULFING COMPUTER SCIENCE
In 2009, four essays in three issues of *Communications of the ACM* [1, 2, 9, 10] addressed "a growing crisis" in the computer science community. They argue that a focus on conference publication has led to deadline-driven, short-term research at the expense of journal publication, a reviewing burden that drives off prominent researchers, and high rejection rates that favor cautious incremental results over innovative work. In one essay, Ken Birman and Fred B. Schneider observe that in Computer Science, "in the past, journal publications were mandatory for promotions at leading departments. Today, promotions can be justified with publications in top conferences." [1] The resulting deluge of conference submissions creates reviewing challenges.

Although Birman and Schneider focused on reviewing, other of the interlocking Table 1 goals arose. Conference deadlines insure timely dissemination of results, but undermine originality, as novel papers are "time consuming to read and understand, so they are the most likely to be either completely misunderstood or underappreciated." More submissions lead to fewer broad program committee discussions and lower quality reviews. Birman and Schneider describe a "death spiral" in which senior researchers cease participating in review panels.

A *CACM* reader eagerly anticipating solutions may be disappointed. Birman and Schneider recommend (i) returning journals to prominence, a plea echoed by other commentators, and (ii) giving authors of conference submission no feedback, either to discourage premature submission or to reduce reviewer workload. In short, a plaintive call for an unlikely return to the past, lacking analysis of why Computer Science in the United States shifted to conference publication in the first place. It did not happen to Computer Science in Europe or Asia, or in other competitive, quickly-evolving fields such as Neuroscience or Physics.

Function	Time	Goal
Production	Venue creation	Defining scope.
		Defining quality or soundness.
		Defining originality.
	Reviewing	Measuring the value of submissions.
		Helping authors improve submissions.
	Publishing	Disseminating results quickly.
		Distributing results broadly.
		Archiving and providing access to all useful results.
		Publishing on schedule / maintaining content flow.
		Adhering to page count constraints.
		Making or not losing money.
Group well-being	Long-term effects	Growing and maintaining a research community.
Member support		Helping individual community members succeed.

Table 1. Goals of journals, conferences, and workshops. (Based on [5].)

3. A DISRUPTIVE TECHNOLOGY AND PROCESS: WORD PROCESSING AND ARCHIVED PROCEEDINGS

Major changes generally have multiple causal factors. In this case a new technology and a related decision seemed to transform the situation like a key opening a lock. The use of text editors and word processors by computer scientists enabled timely, inexpensive production of presentable conference proceedings. This perturbed the complex ecology, setting in motion a series of adjustments that are still being worked out.

Prior to the 1980s, the rare proceedings available at a conference required expensive editing and typesetting or typewritten pages with figures and tables pasted in. By the early 1980s, most computer science researchers had access to text editors or word processors, graphics packages, and printers that supported standard font sets. Conferences published formatting instructions for final versions that were intended to yield a consistent look. Text processors of the 1980s had limited formatting capability, so proceedings lacked today's uniformity, but they looked decent enough. Costs were contained by having authors do most of the work and by shifting from hardcover to trade paperback format.

CHI conferences had inexpensive proceedings prepared in advance from the beginning in 1983. The first international HCI conference, INTERACT 1984 in London, first produced a two-volume provisional paperback proceedings available on site, then a single-volume hardcopy proceedings with a more uniform look. From 1985 on, few if any major CS conferences produced proceedings after the event.

This technology change was not disruptive by itself. The second factor was the existence in North America of a non-profit professional organization that served computer scientists and organized many major conferences. The Association for Computing Machinery (ACM) saw an opportunity in proceedings of trade paperback quality appearance, low production cost, and very low per-copy incremental printing cost. ACM printed many more copies than there were conference attendees and set about marketing the surplus to libraries, the lifeblood of technical and scientific publishers.

In addition, some conference-cosponsors, such as the Special Interest Group on Computer-Human Interaction (SIGCHI), sent copies as a benefit to their thousands of members. Finally, and perhaps of greatest significance, mail-order copies could be purchased very inexpensively by anyone, years later. With library uptake slow, there was effectively an inexhaustible supply.

As a result, fifteen years before the digital library, ACM conference proceedings were archived and widely accessible. These were the two original purposes of journals! The ecological balance of technical communication in was disrupted, with effects that are still being worked out a quarter century later

4. UNINTENDED CONSEQUENCES (1): JOURNAL DECLINE

As editor of *ACM TOCHI*, I was frequently told by senior colleagues that they considered journals irrelevant. This is a development that the commentators lament. Why did it happen?

The short answer is that the major players had incentives to sharply drive up the quality of conference papers. This reduced the incentive for continuing to improve the work and raised the bar substantially for those who tried to do so.

To sell proceedings to libraries, ACM had a stake in papers being of the highest possible quality. In addition, libraries were more likely to acquire and shelve thinner volumes. Many authors, when they realized that conference papers would be immortal, desired to make a good impression. Also, when Digital Libraries, site licenses, and Internet access arrived, academic Computer Scientists realized that their conference product could be easily viewed and judged by colleagues evaluating job candidates, tenure cases, or promotions. A self-policing function arose: If we

let the quality waver, we could lose hard-won respect from peers in other disciplines.

Consider CHI as an example. A 1982 conference led to its formation. Proceedings for that conference were not available after the event. It had a 45% acceptance rate. Over CHI's first 13 years, the median acceptance rate was 27%, the maximum 39%. For the next thirteen years, the median was 22% and never exceeded 25%. For three successive years it was 15%-16%. The 25% ceiling that has held since 1995 coincides with the rise of HCI within academic Computer Science. At many U.S. universities, computer scientists convinced colleagues from other fields to weight papers from very selective conferences highly in appointments and promotions; 25% was a good demarcation.

Assume that the authors of a CHI paper would like to improve it, by responding to reviewer suggestions for which there had not been time or space in the final version, by extending the literature review to actually discuss some of the papers cited, by expanding their own discussion, or by including additional analyses or details. In the 1980s, ACM and IEEE policy was that authors of conference papers, which were not archival, could republish them as journal articles, which were. Usually journals expected more, but many excellent conference papers were republished verbatim or close to it. However, as conference papers developed an afterlife by being effectively archived (and later unequivocally archived in Digital Libraries), IEEE and ACM shifted policy to discourage republication, now considered self-plagiarism.

How much must be changed to merit journal publication? That bar has been steadily raised by publishers, editors, and reviewers. Some even consider the merger of two related conference papers into a single journal article to be unacceptable: A new publication requires previously unpublished data. The result, which I have not seen discussed despite its centrality to the decline of journal publication, is that authors of selective conference papers often find it prohibitively difficult to publish in journals.

Correlational data exists that bears on the above hypotheses, but given space constraints, I will conclude this section with a supportive logical argument. In Europe and Asia, professional organizations did not provide low-cost post-conference access to proceedings. Authors who wished their work to be accessible had to progress it to journal publication. Journals remained the major or only academic measure of achievement. Articles in the leading U.S. HCI journals shifted from being mostly authored by Americans to being mostly authored abroad. Interest in journal impact factors has been high among European computer scientists, but not among Americans -- impact factors were generally not calculated for conferences and did not measure citations appearing in papers in selective conferences.

European and Asian conference acceptance rates generally stayed higher, although some rose under competitive pressure from U.S. conferences. Only recently have I seen growing acceptance of selective archival Computer Science conference papers in some European countries.

In 2004, a prominent UK researcher wrote about CHI:

> 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. [3]

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.

Birman and Schneider decry the erosion of journals and describe a "death spiral" in which people overburdened by deadline-driven conference reviewing cease reviewing for journals. Perhaps in our dynamic field the shelf-life of some results is short or a conference paper captures the essence of the research, but I agree that the additional reflection afforded by an iterative review and revision process is valuable. However, considering the forces that led to the present state, a return to journal preeminence seems unlikely.

Information Schools, comprising computer scientists and researchers from other disciplines, wrestle with the assessment of publication venues—but so do many other schools with CS departments. The iCaucus and Information Conference will have to decide whether a new journal should be formed and whether the proceedings should be archived. However, the crisis that the Information field should work to avoid is not this fifteen-year-old dilemma. It is an emerging second-order effect of the shift to conference publication. The U.S. Computer Science crisis is a crisis of community.

5. UNINTENDED CONSEQUENCES (2): COMMUNITY DECLINE

The core problem confronting Computer Science is that their major conferences focus on assessing and showcasing the field's quality work, a role formerly filled by journals, and have largely abandoned the community-building and community-maintenance function that conferences traditionally fill. In the absence of an effective replacement, there has been a gradual but cumulatively significant decline in the sense of community in major Computer Science sub-disciplines, with no bottom in sight. Diverse factors may be at work, but let's step back to consider a framework from social psychology (Table 2).

	Production	**Group Well-Being**	**Member Support**
Inception	Production demand and opportunity	Interaction demand and opportunity	Inclusion demand and opportunity
Problem-Solving	Technical problem-solving	Role network definition	Position and status attainments
Conflict Resolution	Policy resolution	Power and payoff distribution	Contribution and payoff distribution
Execution	Performance	Interaction	Participation

Table 2. McGrath's Group Functions and Modes [8].

Joseph McGrath identified functions and modes of activities in teams or groups. At different times, groups take on new tasks (inception), work on them (execution), solve problems that arise, and resolve conflicts [8]. Of significance to us are the columns. Groups continuously engage in activities that address production (their *raison d'être*), team health, and member support. We may address the second and third without conscious consideration, but we ignore them at our peril.

Studies of group support technologies tend to focus on the lower left cell, *performance*—effects on productivity, return on investment. Technologies that have positive effects on performance in experiments may founder in practice due to negative effects in other cells. Conversely, technologies that show no short-term performance benefits in studies may have positive effects in other cells that could benefit performance over longer periods [4, 6, 7].

Table 1 is further evidence of a bias toward the production function. It seemed reasonable, yet it focuses overwhelmingly on production. Contrast it with Table 3.

This assignment of function to venue omits considerable nuance. For example, doctoral symposia or full-day workshops held in conjunction with a conference provide member support. But the broad picture is clear, as is the contrast with other fields. A friend described the annual Neuroscience Conference as a must-attend event: "It is where you find out what is happening!" It has 15,000 presentations and 30,000 participants. Quality is not the point, community is. Journals are where he finds quality; workshops are a source of information and feedback for work in progress.

Highly selective conferences work against many of the group well-being goals in Table 3. When 75% of paper submissions are rejected, it is difficult for researchers from allied fields or new researchers who do not know the conventions to break in. Setting aside the fact that being rejected is generally an off-putting experience, many people must present to get travel funds, so engagement is curtailed. The rejected material becomes fodder for spin-off or sub-group activities, which proliferate, scattering people, their energy investments, and the relevant literature. Community identity declines. For a typical topic, only one in four submitted papers is presented and much work in progress is not even submitted, so the conference is not a place to find out what is happening in one's specialty area. This further opens the door for new or competitive venues.

Membership data for Computer Science special interest groups since 1990 can be found at www.acm.org/sigs. SIGCHI membership peaked in 1992. It fluctuates but is currently down about 20% from that level. Conference attendance peaked in 2001. This is true despite an unquestioned increase in faculty, students, and practitioners focused on HCI. Graduate students are a steadily rising fraction of conference attendees and presenters. Practitioners disappeared from the program and then from the audience. Some member support functions are served—students get visibility and speaking experience, professors get their names on papers whether or not they attend. But for most people, rewards for attending have diminished. The papers can be read in the proceedings. In the early years, papers were assigned discussants, but the polished papers that make it through today's competitive review processes leave less room for comment. This is especially true given the Birman and Schneider observation that original or controversial papers are unlikely to survive the review process. One frequently hears statements such as, "I submitted two papers. The original and interesting one didn't make it. The more boring, incremental paper did."

In ACM SIGCHI, once-active community forums are gone. The newsletter, the *SIGCHI Bulletin*, was vibrant through the 1980s. A market research study in the early 1990s found it was avidly read. The past decade it withered and died. The CHI email distribution list used solely for event announcements was once a lively

Function	CS Venue (then / now)	Goal
Production	**Journal / Conference**	Collecting and distributing research results.
Group well-being	**Conference / Not Clear**	Establishing community identity.
		Developing members, maintaining engagement
		Recruiting new members.
		Interacting with parent and sibling organizations.
		Interacting with competitive or rival organizations.
		Managing subgroups and spin-offs.
Member support	**Workshop/ Workshop**	Helping students get visibility and jobs.
		Helping assistant professors get tenure.
		Helping associate professors get promoted.
		Helping full professors get honors.
		Helping practitioners prosper.
		Recognizing research and service contributions.

Table 3. Goals in U.S. Computer Science and venues before / after the shift.

discussion forum. The web-based CHIPlace forum was a focus of community discussion a decade ago; use trailed off and it was taken down. Business meetings held at the conference were once heavily attended and a source of passionate argument—a petition circulated at one conference forced an election of officers. Today gatherings are poorly attended; complaints over reviewing and heavy-handed program committee members are a major focus. A sense of community is found at the program committee meeting, restricted to a small number of mostly senior people. If pressures to save time and travel expense lead to distributed program committee meetings, social interaction will decline further.

The bottom line is that a conference that rejects 75% of submissions may not fill a community-building role unless it has some other irresistible draw—ICIS and SIGGRAPH can be must-attend venues despite high paper rejection rates due to their links to job interviews and exhibitions, respectively. Otherwise, this path seems problematic.

6. ALTERNATIVE PATHS

Academic Computer Science benefits from the status quo, which is intricately woven into its accreditation process. Change will not be easy. It may not be necessary, although as noted in the CACM commentaries, reviewers are more difficult to enlist. Conferences once run entirely by volunteers now contract out much of the work, but reviewing cannot be outsourced. People whose papers are rejected shift their efforts to the conferences that subsequently accept the papers. Birman and Schneider's solution, 'stop writing useful reviews," does not seem viable given the claims for quality in which we take such pride.

Computer Science in Europe has recently shifted toward greater recognition of selective conference papers. Other fields have not. The role of our strong, non-profit professional organizations was significant. How other fields react as online preservation becomes ever easier is something to watch.

Information Schools have some time to explore options. It seems appealing to stress the community-building focus of most major conferences. I attended an American Anthropological Association meeting—7000 anthropologists! Presentation quality varied, the high energy level did not. But it will be a challenge, particularly given that Information School faculty from the Computer Science community may be unaware that other ways of life are possible.

Those who frequent highly selective conferences expect polished work. They often decide which session to attend on short notice. Attending a larger, less selective conference, many complain bitterly about presentation quality. They do not realize that with an hour or so preparation based on the program and other materials at hand, one can have as positive an experience at a large inclusive conference as at a selective conference. But reeducating the Computer Scientists among us is not the only challenge.

The accreditation process must be considered. A large, inclusive conference could accept 80% of submissions for presentation in parallel tracks and identify 20% as Best Paper Nominations. This could provide a quality measure for those who need one, enable more people to present and learn what is going on in their areas, and help people plan their attendance around strong papers.

Another concern is so-called self-plagiarism. It is not expensive to host 80% of submissions online, but is it all archival? One can attach labels, but people put everything on their CVs. This difficulty is inherent in our increasingly visible world. With no cost to putting drafts online, the issue of multiple versions being published in some form is unavoidable.

Perhaps technology, having helped create the problem, can help fix it. Wikipedia's articles with complete version history and discussion pages are a possible model. The mediawiki software has weaknesses, not least of which is that references are not an object type. But perhaps a researcher at some point registers a draft with a system, controlling access, after which all versions, comments, and reviews are recorded. The work may initially be private, then opened to friends or colleagues, later submitted to a workshop, then conference, and maybe a journal or journal-level accrediting process. At each stage the version history is there, review comments are accumulated, the work develops. Self-plagiarism isn't an issue; anyone can inspect the history.

There are issues, technical and otherwise. What happens when an author combines two or three works into one larger work, or when co-authorship changes? How is copyright managed if I submit to an ASIST workshop, then want to submit a version to a CHI conference, and later to a for-profit publisher's journal?

7. CONCLUSION

Information Schools wrestle with issues of identity, direction, and quality measurement. Computer Science offers one model, which appeals to some because they come from the field or because Computer Science has been successful. But in the view of many, Computer Science is in trouble. Therefore, it makes sense to look closely at the current state and understand how it came to be. It was reached not through planning and consideration of alternatives, but because the field was pushed unwittingly down a path by using technology in an obvious and beneficial way, which nevertheless had unintended consequences.

There is more to say and more questions to ask. Let's continue the discussion.

8. REFERENCES

[1] Birman, K. & Schneider, F.B. 2009. Program committee overload in Systems. *Communications ACM, 52,* 5, 34-37.

[2] Crowcroft, J., Keshav, S. and McKeown, N. 2009. Scaling the academic publication process to Internet scale. *Commun. ACM, 52*, 1, 27-30.

[3] Cockton, G. 2004. Email communication, 22 January.

[4] Dennis, A.R. & Reinicke, B.A. 2004. Beta versus VHS and the acceptance of electronic brainstorming technology. *MIS Quarterly, 28*, 1, 1-20.

[5] Grudin, J. 2004a. Crossing the divide. *ACM Transactions on CHI, 11*, 1, 1-25.

[6] Grudin, J. 2004b. Return on investment and organizational adoption. *Proc. CSCW 2004*, 274-277.

[7] Grudin, J. 2008. McGrath and the Behaviors of Groups (BOGs). In T. Erickson and D.W. McDonald (Eds.), *HCI Remixed: Reflections on Works That Have Influenced the HCI Community.* MIT Press.

[8] McGrath, J. 1991. Time, interaction, and performance (TIP): A theory of groups. *Small group research, 22*, 2, 147-174.

[9] Vardi, M. Y. 2009. Conferences vs. journals in computing research. *Commun. ACM, 52*, 5, 5.

[10] Wing, J.M. and Guzdial, M. 2009. CS woes: deadline-driven research, academic inequality. *Commun. ACM, 52*, 12, 8.