PUBLICATION CULTURE IN COMPUTER SCIENCE
A White Paper for the UNC-Chapel Hill Scholarly Communications Convocation
January 2005

by
Jack Snoeyink, Department of Computer Science

Computers facilitate the changes in modern scholarly publication, since computers make it possible to communicate in ways beyond the printed page. Although Computer Science has been dealing with this fact for decades, its own publication culture is confused by current changes. I'll briefly survey this culture, and how current technology conveys authority on digital communication.

Conferences vs. Journals in Computer Science

Computer Science is a discipline whose main goal is to create and distribute synthetic artifacts and whose technology for doing so changes rapidly. As our department's founder, Fred Brooks, has said[1], "When one discovers a fact about nature, it is a contribution per se, no matter how small. Since anyone can create something new [in a synthetic field], that alone does not establish a contribution. Rather, one must show that the creation is better." This has resulted in some quirks for Computer Science publication culture.

The academic conference is the preferred vehicle for dissemination, and some conferences require more effort than journals. The extreme example is the computer graphics conference, ACM SIGGRAPH, that is sponsored by the professional society, the Association for Computing Machinery: an 8 page submitted paper, usually accompanied by videos or demonstration software, is commented on by five referees, and authors write a brief rebuttal before the conference committee decides which 15-20% of the submissions to accept. Accepted papers are presented before an audience of 3,000 at a conference of 25,000, and the proceedings go to many libraries. (My other conferences have between 100-300 attendees, but still accept only 25-40% of submissions.)

Journal review and printing times have grown to be measured in years, rather than months, so journals are used only for archival purposes. Many journal papers are expanded (15-20 page) versions of 8-10 page conference papers. In fact, many journals publish special issues devoted to invited papers from specific conferences: for the 2004 ACM Symposium on Computational Geometry, which I co-chaired, all three key journals in my field asked us to edit special issues. At one time ACM TOG (Transactions on Computer Graphics) refused to accept full (15-20 page) versions of papers presented at SIGGRAPH, but there the reason given was that the conference proceedings had three times the readership of this prestigious journal.

The sharing of software and artifacts is less structured: Research software is usually designed to be used only in the project that created it, since the effort to maintain software that is portable and user-friendly is an order of magnitude greater than that

---

[1] Fred Brooks, Academic Careers, p. 35, 1994.

required to make it work. Many packages are made freely available, with the author retaining copyright (or copyleft). Some are patented, but software patents are far more prevalent in industry than in academia. Although some people include download or page-hits in their CV, the tangible reward of distributing software is primarily its mention in tenure and promotion letters. (Some academics made lucrative jumps to industry, especially during the dot-com boom.)

These quirks are recognized by leaders within Computer Science, whose main concern is that they be recognized as part of the discipline's publication culture by deans and by tenure and promotion committees. The Computing Research Association has written a best-practices memo[2], which summarizes: "Though standard publication is one indicator of academic achievement, other forms of publication, specifically conference publication, and the dissemination of artifacts also transmit ideas. Conference publication is both rigorous and prestigious. Assessing artifacts requires evaluation from knowledgeable peers. Quantitative measures of impact are possible, but they may not tell the implied story."

On the other hand, these quirks of publication culture can pose problems for collaborative work. One of my current concerns is getting publications for my Computer Science students who are working with biochemists. In Biochemistry, it is considered double publishing to take an idea that has been presented at a conference and submit a paper on it to a journal, and algorithms are relegated to a short "Methods" section in a paper and distributed freely on the web or buried in expensive commercial packages.

Changing technology, changing cultures

Technology changes at a more rapid pace than publication culture. This remains true in computer science, despite the many relevant technology changes that from the discipline (browsers such as Mosaic and Netscape, search engines from OpenText, Yahoo, and Inktomi to Google, infrastructure from Akamai to Blogger). Computer science has become comfortable with its emphasis on conferences, especially as the professional societies make the conference publications and their journals available online through library or personal subscriptions. (I've memorized my UNC PIN because of how often I use it for off-campus access to the ACM Digital Library, and the portion of IEEExplor that we get through their Computer Society – the entire package from IEEE is too expensive for a school with no engineering department.) Thus, other disciplines are pushing the envelope in some areas: Physics has developed (and aggressively uses) online preprint archives, and Biology curates many more data repositories.

Rapid increases in journal costs have lead to a few small revolts. The most dramatic to me was that in January 2004, the entire board of the Journal of Algorithms (created in 1979 under Academic Press, just publishing Volume 50 under current owner Elsevier) resigned and created a new Transactions on Algorithms published by the professional society, ACM. Scott Aaronson[3] commented in a blog: "As for why this sort of thing doesn't happen more often, I think part of the answer is that scientists are the most

---

[2] Computing Research Association, Evaluating Computer Scientists and Engineers for Promotion and Tenure, http://www.cra.org/reports/tenure_review.html, 2000.
[3] Scott Aaronson at February 13, 2004 03:28 AM <http://www.freedom-to-tinker.com/archives/000514.html>

reluctant revolutionaries on Earth. Look at the web sites that ***favor*** boycotting companies like Elsevier – for the most part you'll find detailed price data, guarded prose, and skeptical counterarguments, not indignance [sic] over being charged a fortune to buy back a community's own intellectual work." But why should academics tinker when tenure and promotion may be at stake?

Authority in digital communication

One of the questions that must be faced when changing publication culture is how any replacement will gain prestige and authority. Elsevier's Author Guide suggests that this is a service that they provide when they explain why they ask for exclusive rights. "The research community needs certainty with respect to the validity, normally obtained through the peer review process, of scientific papers. The scientific record must be clear and unambiguous. Elsevier believes that by obtaining the exclusive distribution right it will always be clear to researchers that, when they access an Elsevier site to review a paper, they are reading a final version of the paper which has been edited, peer-reviewed, and accepted for publication in an appropriate journal."[4] At the other extreme is the blog, where Rebecca Blood[5] comments "When I began blogging I imagined that someday there might be hundreds of Weblogs, with tens of thousands of readers. Instead, the availability of often free and easy-to-use tools overturned that broadcast model. Instead of dozens of Weblogs with a million readers, there are now well over four million Weblogs worldwide – most with only a few dozen readers, according to studies by Blogcensus (www.blogcensus.net) and Perseus Developent Corp. (www.perseus.com/blogsurvey)"

It is easy to say that a community gives authority, but blogs show that communities can be small. I believe that the professional societies are best positioned to enact changes, if they can avoid acting too much like commercial publishers.

With digital technology, we could also measure authority much sooner than in the past: Google came to dominate the search engine market because of its "PageRank" technology explicitly conveyed authority on web sites from studying their links. Some bright young computer scientists should come up with a PaperRank system, but then they would be tinkering with their own tenure and promotion.

---

[4] Elsevier's Author Guide: http://authors.elsevier.com/getting_published.html?dc=CI

[5] Rebecca Blood, How Blogging Software Reshapes the Online Community, Comm ACM, 47(12), 53-55, Dec 2004