

JACM's 50th Anniversary

This is Volume 50 of *JACM*; it marks *JACM*'s 50th year of operation. *JACM* was the first journal devoted to computer science, started when it was not clear to many that there was even a distinct field of computer science. Fifty years later, both the journal and the field continue to thrive. This 50th anniversary issue is meant to be a celebration of both the accomplishments of the past and the potential for the future.

Most of the issue deals with the future: I asked winners of the Turing Award and the Nevanlinna Prize to discuss up to three problems that they thought would be major problems for computer science in the next 50 years. (Jim Gray's Turing Award address was actually about twelve of what he thought were the most significant problems in computer science. Rather than asking him to pick three of them, he discusses all twelve in his article.) Their articles cover a wide range of topics. Although the articles were written independently, there is very little overlap between them. They address a wide range of problems, all fascinating. I hope that these articles will serve as a distributed version of Hilbert's Problems for our field. Unlike most of Hilbert's problems, most of these problems are relatively fuzzy. It's not always clear what counts as an answer. It will be interesting to see what the status of these problems will be 50 years from now, and how many will be viewed as "solved" (and how many others will turn out to be perhaps not such significant problems after all, in light of new developments).

To give a sense of the past, the issue also includes reminiscences of many of the former editors-in-chief of *JACM*. We are particularly fortunate that among those who were able to write are the first two editors-in-chief: Franz Alt (the founding editor) and Mario Juncosa. Their articles give a real perspective on what things were like at the beginning of *JACM*. The article by Mark Mandelbaum, who has been involved with the journal for 25 years, first as Managing Editor and now as ACM Director of Publications, gives a sense of things from ACM's point of view; it also includes some fascinating statistics.

Reading these articles inspired me to go the library and leaf through past volumes of *JACM* myself.¹ Volume 1, Issue 1 includes an article by John Backus on "The IBM 701 Speedcoding System". This system was a precursor to the system that included FORTRAN, and is clearly in the purview of what we view as computer science today. The issue also includes descriptions of application programs, such as an article entitled "Life Insurance Premium Billing and Combined Operations by Electronic Equipment".

By 1959, it is clear that the preponderance of articles in the journal are in numerical analysis, certainly the most mathematical aspect of computer science at the time. However, even then, there was a concern for balance. Indeed, an editorial by Walter Bauer, Mario Juncosa (who had just taken over as editor-in-chief), and

¹All the articles I mention are available in the ACM Digital Library (<http://portal.acm.org>), although looking at the journals also gave me a sense who the area editors were; this material is not yet available on the DL.

Alan Perlis, “ACM Publication Policies and Plans”, says

Every effort will be made in the future to publish good papers in computer programming and business applications. In addition, of course, papers on computer logic, system design, and hardware will continue to be most welcome Also, it is realized that important applications of the computer as an information handling device (real time control, language translation, etc.) are appearing with increasing frequency.

By 1970, the journal had much of the look and feel of today’s journal. Everyone on the editorial board had a specific area of coverage. The areas then were Artificial Intelligence, Combinatorial Theory, Computer Applications, Information Retrieval, Linear Algebra, Mathematics of Computation, Operating Systems, Operations Research, Programming Languages and Methodology, and Theory of Computation. Some of these areas still exist on today’s editorial board. There was also the same concern then about the coverage of the journal. Gerry Salton’s introductory editorial as editor-in-chief in 1969 says “A major effort will be made to attract important papers in all areas of computer science, including the applied areas. While maintaining the high standards to which everyone is accustomed, we shall do our best to produce a *Journal* which appeals to all classes of readers and which few will want to miss.” Salton returns to this issue in his 1970 editorial:

During the calendar year 1969, fifty articles were published in the *Journal of the ACM*,² falling approximately into the following categories:

automata and theory of computation	22
numerical applications	13
computer applications	13
computer systems	2

The list is indicative of the policies which have governed the *Journal* over the last 15 years: the *Journal* is relatively selective in what it publishes, and the concentration in the theoretical areas and in numerical mathematics is heavy.

. . .there is some concern that a great many worthwhile manuscripts never come to the attention of the editors or, alternatively, that manuscripts are not accepted when received because of editorial preferences for theoretically oriented material.

While such arguments may in fact apply in isolated cases, they surely do not accurately describe the general situation. Concerning first our orientation, we feel that it is essential that an outlet for important theoretical papers in computer science be maintained, and it is entirely proper that

²Articles have gotten longer over the years. In 1969, there were actually 52 articles (not counting an editorial, Richard Hamming’s Turing Award Lecture, and three errata) in 646 pages; in 2001, there were 38 articles and 1241 pages.

this outlet should remain the *ACM Journal*. Even so, we are making efforts to broaden our coverage, and almost half of our editors now handle various applied field

This goal of adding new area editors to ensure that *JACM* covers more of computer science seems to be a running theme in the history of the journal, and can be seen even in my first editorial as editor-in-chief (from 1997):

I would like *JACM* to cover “computer science, broadly construed”. *JACM* is already trying to cover large parts of computer science; just look at the list of areas covered by our editors. However, while *JACM* is viewed by many as the journal of record for the theoretical computer science community, it currently does not have the same central role in all communities within computer science. I would like to change that. As a first step, I have added three new areas to our coverage

Coverage has not been the only concern of editors-in-chief over the years. Salton wrote an 1971 editorial entitled “Some Thoughts on Scientific Information Dissemination” that discusses, among other things, the role of the journal. He argued that we “concentrate on that material which is really worth studying and preserving” and “de-emphasize the alleged requirement for speed in the publication and dissemination of information”. With regard to the first point, he said the following:

It may be unrealistic to suggest that we abandon the publication of marginal items—including, for example, most conference proceedings—and rely for our information on the 2 or 3 or 4 primary journals in each area, together with appropriate review publications, but it is a fact that everyone would be far better off (One may add that a drastic reduction in the number of new publications may become mandatory whenever the libraries decide to curtail standing orders for all newly published material.)

Salton did not foresee how conferences would come to dominate the field, nor the rise of the Internet and the new tools for information dissemination that have become standard. I would guess that far more people get papers now from CiteSeer or authors’ home pages than from journals or conference proceedings. Nevertheless, the issues that he raises regarding quality control and knowing which are the truly important papers in the field still remain. Computer science seems to be the only field where conferences play such a dominant role (and where conference publication is viewed as being, in most cases, on a par with journal publication). How this will affect the field remains to be seen.³

Although many of the themes at *JACM* have remained stable, there have been a number of changes in the journal during my tenure. I have already mentioned adding new subject areas. These include Machine Learning and Computational

³As a personal aside, I thought I was the first *JACM* editor-in-chief to write editorials. As should be clear from the comments above, Gerry Salton was actually the first and, in fact, the only other editor-in-chief to do so. Gerry unfortunately passed away in 1995, just before I joined the Cornell faculty. His administrative assistant at the time, Cindy Robinson, became my administrative assistant, and is the editorial assistant for *JACM*.

Learning Theory; Decisions, Uncertainty, and Computation; Networking; and Software Engineering. In addition, some areas were modified: Architecture became Parallel Computing and Architecture; and Database Theory became Database Theory and Systems. There will no doubt be further changes in the not-too-distant future, as the field evolves.

Another recent significant change for *JACM* is a new production path for papers: we now go directly from author-provided \LaTeX to PDF. This is good for authors, since it means that all the errors that were introduced in rekeying (or converting from \LaTeX to SGML) are eliminated. It is also good for ACM, since it gets to take more advantage of the work that authors have done in preparing the article. The \LaTeX to PDF path is not necessarily the way of the future. Other languages (such as MathML) may do a better job of providing an appropriate presentation for the web. Nevertheless, for now, I think it's the best we've got. By using \LaTeX directly and committing to send authors marked-up galley proofs showing all changes made, *JACM* has become one of the most author-friendly publications around.

The near-term future holds some other changes. I believe we need to be more proactive about getting the best papers from various conferences. We have already taken some preliminary steps. We have worked out an arrangement to get the winners of the Best Paper Award for FOCS (the IEEE Conference on Foundations of Computing), starting with the 2002 FOCS; in addition, the best papers of the 2002 PODS (the ACM Symposium on Principles of Database Systems) will be split between *JACM* and the *ACM Transactions on Database Systems* this year. We are actively pursuing the best papers of other conferences, while trying to tread the thin line between two slightly conflicting goals: (1) representing computer science, broadly construed, and (2) getting the best papers in each of the areas that it represents. Having a special issue of *JACM* completely devoted to one area seems to me to torque the journal a little too much. However, I am comfortable with having, say, three papers in one area in a particular issue of *JACM* (which is what I expect to happen in the case of PODS).⁴

Readers may have also noticed a classic paper by Stockmeyer and Meyer in the previous (November, 2002) issue of *JACM*. Although the result was well known, this paper had not been published in an archival journal. Results like this that stand the test of time belong in archival journals; I believe that *JACM* is an appropriate home for the best of them.

What about the longer-term future? It is always hard to predict the future of journals in general, let alone *JACM*, in light of all the changes happening in the world of academic publishing. Journal publication used to play a number of important roles:

- it made results widely available;
- it *registered* results, in the sense of giving them a timestamp (so that authors could claim priority);
- it *certified* results, in the sense of saying that the editor(s) of the journal believed

⁴The PODS issue will not be the first time that part of a *JACM* issue is devoted to papers from a particular conference. In 1977, for example, a selection of papers from POPL (the ACM Conference on Principles of Programming Languages) appeared in *JACM*.

them to be correct and significant;
 —it made papers look better, typesetting authors' typewritten manuscripts.

Authors can now make results widely available by posting them on the web; a web posting also serves to register the results. And, of course, with modern computer typesetting systems, authors can typeset papers as well as (indeed, in some cases, better than) publishers. Thus, I see the main role of the journal of the future as that of certification. Something will be needed to distinguish the important papers from all the rest. In most fields, perhaps the main way of distinguishing the important papers from the rest is by the journal in which they appear. In computer science conferences also play a significant certification role. There is relatively little pressure to publish a paper in a journal if a preliminary version has already appeared in a prestigious conference. I'm not sure how good this is for the field, but there's no question that it has had an impact.

We always say that journal publication is important; it's easy to explain why. Conference versions are severely page-limited, and generally have to scrimp on details of models, results, methodologies, and proofs. These details are usually in the journal versions (sometimes thanks to encouragement from referees). Feedback from referees (and others), together with the absence of strict conference deadlines, means that the presentation of a journal version of a result is typically far more polished and accessible, and less likely to contain errors, than the conference version. The data presented in Mark Mandelbaum's article provide further support for the greater impact of journal articles relative to conference articles. Moreover, there are regrettable examples where vague claims in conference proceedings have had the effect of preventing people from working on what may appear to be a solved problem, but in fact is not. Finally, the recent scandal in physics regarding Jan Hendrik Schön's irreproducible results emphasizes the need to have a careful journal review process to ensure scientific integrity and the importance that the community places on the certification process.

I am not saying anything that hasn't been said by many others before me. Despite the acknowledged importance of journals, we as a community have not supported journals perhaps as much as we should. Certainly the importance of journal publications for tenure has greatly diminished (again, this phenomenon seems to be unique to computer science), and a growing number of conferences have "Best Paper" awards, while only the Gödel Prize seems to emphasize journal publications. I still believe that journals like *JACM* will play an important role in the future. But, in the end, this is a judgment that the community has to make. If we want journals to play an important role, then there will have to be appropriate incentives put in place to ensure this. In any case, it will be interesting to see what the preferred method of communicating results will be 20 and 50 years in the future.

This will probably be my last editorial. I am coming to the end of my second three-year term; that's considered to be an upper bound for *JACM* editors-in-chief. It has been an honor to serve as editor-in-chief during this exciting time. *JACM* is largely run by volunteers, aided by a very professional staff at ACM headquarters. I would like to thank everyone over the past 50 years who contributed their time and talent to help make *JACM* the journal it is today. I was particularly fortunate to have an outstanding editorial board and an energetic and effective editorial

assistant, Cindy Robinson. I hope that my successor is equally fortunate. I look forward to the next 50 years!

Joe Halpern
(halpern@cs.cornell.edu)