



BY MOSHE Y. VARDI

EIC YEARS

2008—

# CACM: PAST, PRESENT, AND FUTURE

The French adage “Plus ça change, plus c’est la même chose,” or, the more things change, the more they stay the same, still rings true today. Reading over the essays of my predecessors, one recognizes the thread that runs through all of them, which is the constant need of CACM to reinvent itself. In fact, I discovered an April 24, 1964 report from a Commission of Thoughtful Persons to the ACM Council that stated “It was felt that *Communications* was becoming too much of a journal and that a re-evaluation is in order.” I suspect this ongoing need to rethink CACM, a flagship publication for professionals working in a fast-moving and ever-changing field, will stay with us for the foreseeable future.

I came of age as a computer scientist in the late 1970s, during my formative years as a graduate student. I remember being highly influenced by some great research articles published during CACM’s “black-and-blue” years. E.F. Codd’s paper, “A Relational Model of Data for Large Shared Data Banks” published in the June 1970 issue was given on Mt. Sinai, from my perspective. (1970 was felt to be in the dim past for a graduate student in 1979.) It was clear upon its February 1978 publication that Rivest, Shamir, and Adleman’s paper, “A Method for Obtaining Digital Signature and Public-Key Cryptosystems,” was a seminal one. And in May 1979, De

Millo, Lipton, and Perlis’s “Social Processes and Proofs of Theorems and Programs” was instantly controversial. (Indeed, it still makes for interesting reading today, though the tremendous progress in formal-methods research has dulled its edge.)

The essays here by Denning and Cohen describe the changes that CACM underwent during the 1980s and 1990s. In 1996, the ACM Publications Board decided that CACM ought to be largely run by professional staff, guided by an advisory board. When Cohen’s tenure as EIC ended, the Board did not appoint a new EIC.

In retrospect, eliminating the position of CACM EIC and

THE FEBRUARY 1978  
 ISSUE OF CACM FEATURED  
 THE LANDMARK PAPER  
 “A METHOD FOR OBTAINING  
 DIGITAL SIGNATURE AND  
 PUBLIC-KEY CRYPTOSYSTEMS”  
 BY RONALD L. RIVEST,  
 ADI SHAMIR, AND LEONARD  
 ADLEMAN. THE NEW CACM  
 REVITALIZATION PLAN AIMS TO  
 DRAW RESEARCH WORK  
 OF GREAT INFLUENCE  
 INTO THE FOLD.

reducing the role of the editorial advisory board left a void. By the mid-2000s, it was increasingly clear that CACM needed attention. When Dave Patterson became ACM President in 2005, “fixing” CACM became one of the priorities of his presidency. In his January 2006 “President’s Letter,” he wrote “When I was running for ACM President, I asked people for feedback about the Association. The consistent advice I received was to do something ... about *Communications of the ACM*.” Plainly speaking, by 2005 dissatisfaction with CACM has become quite pervasive in broad segments of the computing community.

While the format envisioned in the early 1980s was that of a content-rich magazine, the reduced role of the editorial advisory board combined with a relatively small professional staff meant that most of the content came from submitted articles. Over the years those articles have evolved to be strongly slanted toward Management Information Systems. Over time, a significant segment of the ACM



membership lost interest in the publication.

Patterson argued that this state of affairs was unacceptable, concluding “ACM must present a compelling value proposition for individuals to join or stay members of ACM. This means our flagship publication must be the best it can possibly be.”

In mid-2005, Patterson commissioned the CACM Task Force to discuss how best to tackle the revitalization of CACM. Their conversations took into account the weaknesses within CACM coverage, the evolving composition of ACM membership, its outdated design, the emergence of the

Web and the ACM Digital Library, and the introduction of *Queue*, ACM’s magazine for young practitioners that debuted in 2003.

Patterson’s Task Force, like the 1983 counterparts, was heavily influenced by the content model of *Science*, one the most prestigious scientific publications (being published in *Science* is a career milestone for a scientist). *Science*’s content consists of an extensive news section, several columns (policy, education, and book reviews), and research articles. While the research articles are aimed at specialists, about one-third of the articles in each issue are accompanied by a short

Re-creating the prestige of the past for CACM seems to me to be an important step in restoring the image of computing. CACM is not just the flagship publication of ACM, it is also, or at least it ought to be, the flagship publication of the computing field. CACM is our “storefront window”; it ought to project an exciting image of a dynamic field.

“Perspective” piece aimed at a general scientific audience.

The final model proposed by the CACM Task Force in late 2005 is similar to the model envisioned by the 1983 redesign effort noted in Denning’s essay. It consisted of news, columns, computing practices, and research articles to be drawn from ACM conferences and journals and accompanied by brief but broadly aimed perspectives.

I was asked by Dave Patterson to take over the task of revamping CACM in early 2006. After some hesitation, due to other commitments and the magnitude of the task, I accepted the job.

The computing field went through a perfect storm in the early 2000s: the dot-com crash, the telecom crash, the offshoring scare, and a research-funding crisis. After its phase of glamour in the late 1990s, the field seems to have lost its luster. For those of us in academia, the plunging undergraduate enrollment is evidence for the “image crisis.” At the same time, I fervently believe that our field has a glorious future. Re-creating the prestige of the past for CACM seems to me to be an important step in restoring the image of comput-

ing. CACM is not just the flagship publication of ACM, it is also, or at least it ought to be, the flagship publication of the computing field. CACM is our “storefront window”; it ought to project an exciting image of a dynamic field.

I love the idea of getting research articles from computing research conferences. The reliance of our field on conferences is unique among the science and engineering disciplines. The fast publication cycles and the sharp focus of the conferences move our field forward very quickly. At the same time, our conference-based culture has resulted in a severe fragmentation of computing research. A reader of *Science* can keep track of progress across all of science. In contrast, it seems impossible to keep track of progress across computing research, even at a high level of abstraction. A strong Research section in CACM might be able to help us re-create the unity that our field had in its early days, as reflected in the early CACM.

I embarked on the CACM renewal project in late 2006. With all due respect to the CACM Task Force, I was not yet ready to adopt their conclusions.

The Task Force consisted of a group of ACM insiders. I thought it best to carry out a much broader conversation before committing to a particular model.

During the winter and spring of 2007, I participated in four major conversations about the “new CACM.” The first conversation was with the SIG Board, which consists of the chairs of all ACM’s Special Interest Groups. I then organized three focus groups—in New York, San Francisco, and London—each consisting of about 25 computing professionals from industry and academia. In each conversation I described the history of CACM and presented the recommendations of the Task Force. This was followed by very lively conversations typically lasting several hours. It turns out that people, though unhappy with CACM, do care passionately about it and do have many ideas on how to revitalize it.

There were several main points that echoed throughout these conversations. First and foremost, there was almost unanimity that ACM must have a print flagship publication. With all the advances in online publishing, there is yet no substitute for paper

publishing. The idea of bringing select research articles back into CACM was quite popular, but the participants felt strongly that such articles should be rewritten and well edited to address CACM's broad readership. At the same time, there was strong consensus that CACM should continue to publish submitted, peer-reviewed articles, although with other new content to be included in the magazine, the space available for submitted articles would be more limited and thus the articles published would need to be of the highest quality.

As far as content addressing the interests of practitioners, the general feeling was that *Queue* provides high-quality content, focusing on software practice, and CACM should not attempt to compete with *Queue*. In addition, the participants uniformly expressed a desire to see a strong news section and an "edgier" opinions section (as put by one focus group attendee: "Let us see some blood on the opinion pages of CACM"). There was a general agreement that CACM does not have the look and feel of a high-tech publication and that a graphic redesign is long overdue. Finally, there was a lively debate, and no agreement, on the issue of possibly renaming the publication. Many felt that 50-year-old name "*Communications of the ACM*" is a valuable brand; others felt a new name is needed to send a strong message that CACM is being "rebooted."

One of my main conclusions from these conversations was that the practical content offered by *Queue* is of the highest quality and likely of interest to a broad cross-section of ACM members,

not just "young practitioners" for which it was originally intended. Thus, if we want CACM to serve the broad interests of the ACM membership, then it should offer the core content available in *Queue*. Thus, the Practice section of CACM should become the print outlet for *Queue*. The *Queue* editorial board should continue to produce high-quality content and develop further the *Queue* brand through an enhanced presence on the ACM Web site to serve practitioners.

With the decision to bring *Queue* articles into CACM to address the needs of practitioners, the content model for the new CACM fell into place. The focus will be on cutting-edge material, organized in five sections: News, Practice, Breakthrough Research, Refereed Articles, and Opinions and Views.

**News:** The news section will have a very distinct "voice," covering research and practice in computing on a global scale. There will be an increased emphasis on research from a news perspective. As a monthly publication, CACM will not compete with more timely news services, but rather analyze the latest news in greater depth for computing professionals and interpret its potential impact. ACM news of broad interests will also be covered.

**Practice:** CACM will offer coverage of cutting-edge, emerging technology issues. Such material would be of interest to both academics and practitioners. To obtain this material, the *Queue* editorial board, led by Stephen Bourne, will continue to function as it has in the past—

providing articles, columns, and interviews.

**Breakthrough Research:** The goal is to bring research articles, covering the broad spectrum of computing research, back to CACM. This will be a combined effort of conference and program chairs (ACM as well as non-ACM) and the CACM editorial board to select the best research material coming out of conferences and share that information with the wide readership of CACM. Each selected research paper will be adapted to the broad CACM readership and will be preceded by a short "Technical Perspective," giving readers an overview of the important points and significance of the research. These perspectives will be written by noted experts in the field. (This issue of CACM contains two sample research articles beginning on page 104 that span the spectrum of computing research: an article on MapReduce, with perspective by David Patterson, and an article on locally sensitive hashing, with perspective by Bernard Chazelle.)

**Refereed Articles:** CACM will continue to publish, as it has since its early days, peer-reviewed articles. Such articles must be of the highest quality and of interest to a broad sector of the ACM membership. We will seek both solicited and unsolicited submissions. In particular, the CACM editorial board will solicit survey and review articles in topics of broad interest.

**Opinions and Views:** The final component in the content model of the new CACM is a section dedicated to opinions

and views, typically of a non-technical nature. Feature columns from respected voices in computing have always been a very popular element in CACM. The plan is to maintain the most popular columns, as well as add new dynamic ones. In addition, the editorial board will continually solicit opinion pieces on topics of broad interest to the computing community. Controversial issues will not be avoided, but be dealt with fairly, by representing both sides of the issue. To maintain the freshness of this section, the editorial board will seek, on an ongoing basis, new material in areas such as global issues, ethics, history, education, policy, interviews with luminaries, and the like.

“Healthy soul in a healthy body,” said the ancient Greeks. To that end, CACM will embark on a complete graphic overhaul; from cover logo to career opportunities and everything in between. All these very different elements must have a cohesive, professional, high-tech look and feel. There is also a need for a complete redesign of CACM’s Web site, with the goal of creating a dynamic, rich Web site that is much more than an online store of the magazine’s content. The aim is to think of CACM as consisting of a print publication, a rich Web site, and email channel to readers.

The reader who compares the content model proposed here with the content model described in Denning’s essay would be struck by the similarity. A natural question is whether this model would be successful

now, when it has not been so successful in the past. In my opinion, previous redesigns of CACM tended to put more attention on the content and look and feel, but less attention to developing the organizational structure that can support the production of a high-quality monthly publication. By way of comparison, CACM is produced by a staff of five professionals, while *Science* is produced by a staff of 100 professionals! By harnessing the current staffs of CACM and *Queue*, and by hiring additional professionals with critical skills, we will have the professional staff that ACM’s flagship publication deserves.

For CACM to stay attuned to the interests of the computing community, it needs an active and authoritative editorial board. One must be mindful that editorial board members are volunteers, who typically shoulder the commitments of a full-time job. While it is tempting to conceive of a small cohesive board, the burden of producing a monthly publication would overwhelm a small volunteer board. The solution is to constitute the board as a collection of semi-autonomous sub-boards, corresponding to the various sections of the publication. The full editorial board is expected to have around 50 members.

This plan for a revitalized CACM was presented to and approved by ACM Publications Board in May 2007 and ACM Council in June 2007. Staffing and recruiting for the editorial board are currently under way. We have recruited an impressive

panel of sub-board chairs to date: Breakthrough Research—David Patterson and Stuart Russell; Refereed Articles—Al Aho and George Gottlob; Opinions and Views—William Aspray and Nigel Shadbolt; News—Marc Najork and Prabhakar Raghavan; and for the CACM Web site—Marti Hearst and James Landay. The launch of the redesigned CACM is planned for later this year. Watch for an announcement soon!

In conclusion, I’d like to add one final point. We live in a consumer society, so it is easy to evaluate products from a consumer perspective: “Is CACM a satisfactory product?” “Am I getting my money’s worth for my ACM membership?” ACM, however, is not a consumer-product vendor, it is a professional society. We are not ACM customers, we are ACM members. CACM is not a product, it is a project. For this project to succeed, the membership of ACM must collectively undertake it. Let us—together—make CACM the exciting publication it should be. Please write to me at [cacm-eic@acm.org](mailto:cacm-eic@acm.org). **C**

---

MOSHE Y. VARDI ([cacm-eic@acm.org](mailto:cacm-eic@acm.org)) is Karen Ostrum George Professor in Computational Engineering and Director of the Computer and Information Technology Institute at Rice University, Houston, TX.

---

© 2008 ACM 0001-0782/08/0100 \$5.00

**THE RESEARCH PREVIEW  
BEGINS ON PAGE 104**