



Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2008-01

Deja Vu All Over Again

Denning, Peter J.

Deja Vu All Over Again (January 2008) PJD reviews his years as Editor of the CACM. He concludes that the current "revitalization" effort is based on the same model as the 1982 effort. The big difference is that ACM has allocated 100% of the required budget, whereas in 1982 ACM allocated less than 10%.

<http://hdl.handle.net/10945/35487>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>



BY PETER J. DENNING

EIC YEARS FEBRUARY 1983—SEPTEMBER 1992

DÉJÀ VU ALL OVER AGAIN

After a 10-year struggle within ACM to define a Journal for All Members (JAM), a “new” *Communications* was launched in the cold of February 1983. CACM was to leave behind its pure research past and transform into a professionally useful, interesting, monthly magazine for all members. The CACM that evolved in the decade following 1983 is substantially the form you find today. I was the EIC who managed the transition.

To understand the “new” CACM, you need to understand the “old” CACM that preceded it. Stu Lynn has reported that a simple disagreement over the covers led to the formation of the ACM Publications Board in the mid-1970s and to a major restructuring of the ACM publications in the late 1970s. The 1970s were a major growth phase for ACM and the computing field, with a continuous stream of amazing new discoveries and inventions. ACM offered its authors two research publishing venues: *Journal of the ACM* (JACM) and CACM. CACM was the preference for papers about systems, architectures, and applications; JACM for theoretical papers. But these two journals could not accommo-

date the growth of the computing field.

THE PUBLICATIONS STRUGGLE

By the 1970s, the publications budget, which covered JACM, *Computing Reviews*, *Computing Surveys*, and CACM, was about half the ACM budget. The member cost of CACM alone was about half the annual dues. ACM revenues were very tight and everyone was sensitive about returns on investment.

CACM and JACM could not keep up with the explosive growth of scientific discoveries and technology inventions. By the mid-1970s there were major queues—and delays averaging three years—in both publications. Authors and readers alike complained bitterly to the ACM leadership and Council. Presi-

dential campaigns turned on proposals for improving publications, especially CACM. Unfortunately, there was not enough money to pay for the additional pages that would eliminate the CACM backlog. And even if there were, a typical issue would be over half an inch thick! Eventually there was a consensus favoring a major restructuring of publications to allow for more research publications, each self-supporting with its own subscriber base.

At the same time, an increasing number of SIGs wanted to start Transactions in their disciplines. The most active promoters were programming languages, computing systems, databases, graphics, and office automation. The SIGs had surplus funds to put into these publications.

Under the leadership of President Tony Ralston, ACM backed a project in AFIPS (American Federation for Information Processing Societies, now defunct) to launch a *Scientific American*-style magazine for computing. It was called *Abacus*. The prototypes were slick and compelling. Around 1975, AFIPS declined to launch *Abacus* for financial reasons. Ralston tried to persuade ACM to launch a scaled-down version of *Abacus*, but it was too costly. The *Abacus* concept, however, established a beachhead in the minds of everyone thinking about the form of an improved CACM.

In 1978, the Publications Board, under the leadership of Stu Lynn, forged a consensus around a long-range publications plan. The plan called for the establishment of a line of self-supporting research Transactions in areas of established need. New Transactions in the areas of greatest backlog in CACM were of highest priority. By 1983, six Transactions were launched and the corresponding departments discontinued in CACM. Today, there are 32 Transactions and five more are on the way.

The long-range plan also called for CACM to transform into a concept called “Journal for All Members” that included aspects of *Abacus*. However, it took until 1982 for enough of a consensus to form around this idea that it could be incorporated into CACM.

COMING TO A HEAD: 1982

When I was president of ACM (1980–1982) I heard numerous complaints about CACM. At that time, six Transactions had been launched or were about to debut, and CACM’s corresponding research departments were eliminated. Although the backlogs were gone, so was the technical content. Now the readership had no news whatsoever about research advances in computer systems, databases, graphics, programming languages, or computer architecture. At least with the backlogs

they saw three-year-old material. Now they saw nothing.

I spent a lot of time working with ACM leadership to forge a consensus around the JAM ideas as a way to transform CACM and respond to the members. The Council asked me to serve as EIC when the new CACM launched in early 1983. With the active participation of ACM Council, we put together a plan for CACM with these elements:

1. *News*. Refocus from ACM to industry. Eventually spin off all ACM news and calendars into a separate newsletter. (Done in 1990 with the debut of ACM MemberNet.)

2. *Computing Practices*. Expand coverage of technology topics, case studies, and how-to articles for practitioners, especially software developers. Hire new editors and writers to work proactively with practitioners to develop articles. (Ed Sibley was the chief editor for this.)

3. *Research*. Continue the existing research departments in emerging areas. Work with SIG conferences to get best papers in all other areas, especially in the departments that had been spun off to Transactions. Rewrite these articles so they can be appreciated by ACM professionals outside the immediate research area of the author. Where necessary, get experts to write opening perspectives to help readers appreciate the context and significance of a research paper.

To understand the “new” CACM, you need to understand the “old” CACM that has preceded it.

VOLUME 31 NUMBER 5

MAY 1988

COMMUNICATIONS OF THE acm



IN THIS ISSUE **Stalking the Wily Hacker**

THE MAY 1988 ISSUE FEATURED A FIRST-PERSON ACCOUNT OF CLIFFORD STOLL'S YEAR-LONG ODYSSEY OF SILENTLY TRACKING A GERMAN COMPUTER PROGRAMMER WHO BROKE INTO THE COMPUTER SYSTEM AT LAWRENCE BERKELEY NATIONAL LABORATORY (AMONG 40 OTHERS WORLDWIDE) TO DISCOVER THE INTRUDER WAS A SPY SELLING SOFTWARE AND MILITARY DATA TO THE KGB. THE STORY, "STALKING THE WILY HACKER," RECEIVED WORLDWIDE MEDIA COVERAGE.

4. *Articles.* Establish a new line of contributed and professionally written articles in the *Abacus* style.

5. *Columns.* Commission regular columns from excellent writers. (The first was "Programming Pearls" by Jon Bentley, beginning August 1983.)

6. *Design.* Hire a professional design company to create a new look and feel for CACM that integrated all the elements noted here. Consult with them on

every issue.

This plan drew on the many ideas from the JAM proposals, reader surveys, and comments. We believed it would establish a new balance among these elements that would prove to be much more satisfactory than the CACM of the day.

Council endorsed the final design and editorial plan in 1982. The new CACM was launched in February 1983 after a special issue in January to com-

memorate the best of CACM in its first 25 years.

But there was one problem: ACM Council wanted us to implement the plan but did not have the funds to hire all the staff required to execute the plan. We were able to hire two new editors and one journalist, but not the five editors and three journalists we thought we needed.

Therefore much of my time as EIC was spent on finding creative ways to implement as much

of the plan as possible within a meager budget.

WHAT WORKED AND WHAT DIDN'T

The news section took several years to find its footing. The biggest problem was finding news items that would still be fresh by the time the issue was published.

"Computing Practices" was our biggest challenge. Many practitioners are not inclined to publish and so it is necessary for the editorial staff to visit many conferences as well as solicit and write articles. We hired journalist Karen Frenkel, who wrote many articles and conducted many interviews; but these articles were quite labor-intensive. We needed three more Karens, but we did not have the budget. Her works were a big hit with readers. Once, Karen and I visited Apple Computer to interview Steve Jobs (published April 1989). When we asked if he thought the Internet would be crippled by hackers, he buried his head in his hands for a full minute; then looked up and said, "No, they see it as a critical infrastructure for their own work."

Another major success was the case studies conducted by Alfred Spector and David Gifford of MIT, who visited project managers and engineers at major companies and interviewed them about their projects, producing no-holds-barred pieces. This section was wildly popular among the readers. Unfortunately, the labor-intensive demands of the post got the best of them after three years, and we were not able to replace them. Also by that time, companies were getting

more circumspect about discussing failures and lessons learned in public forums.

I would say we improved CACM's coverage of computing practices, but not to the degree we envisioned. In 2002, former ACM president Stephen Bourne persuaded Council to undertake a major initiative in the computing practices area by founding *Queue* magazine. *Queue* got the budget needed to do this right and ACM finally learned how to do it well.

Readable research. We found that many of the articles submitted to the remaining research departments were much less technical than articles submitted to the old departments. It was much easier to edit them into the article format. We also found that making arrangements with SIG conferences for best papers was much more difficult than we thought; they were not a fruitful source for CACM.

When we saw this approach to research was not viable, we seriously investigated imitating *Science* magazine's approach. The idea would be to invite research papers from all sectors of computing, edit the acceptable ones heavily to make them accessible to our audience, and have a rapid review process. We envisioned a day when the *New York Times* would cite a scientific breakthrough in a forthcoming article in the CACM—just like in *Science*. We visited *Science* magazine to find out how they do it. To our dismay, we discovered that the number of staff required to handle the rapid review and editing process was well beyond our means. We abandoned this idea.

Eventually we decided to dis-

continue the research category altogether and concentrate on doing the articles category well.

Articles. It was quickly apparent that our resources would not allow us to realize our dream of giving articles the full *Abacus* treatment. We would need 10 articles editors and we only had two. Moreover, we knew that many *Scientific American* readers found the articles shallow, and many authors felt their work was so rewritten it was no longer theirs. By 1985 we had abandoned the *Scientific American* model and settled instead on Sigma X's *American Scientist* model. Their editors solicit papers from leading researchers, asking them to write articles specifically for their publication. Editors work with authors to improve sentence and article structure for the best connection with the reader; the objective is to improve readability while retaining the author's own voice. *American Scientist* readers felt its articles had good depth, and authors felt it was still their own work. We could provide the editing and scouting needed to run this model from within our existing resources.

We established regular special sections to concentrate on emerging areas discovered by our editors. One of our first was a compendium of the best computing humor of all time (Apr. 1984, with Peter Neumann as editor). Our first outreach section—Computing in the Frontiers of Science—was published as a joint venture with the IEEE Computer Society (Nov. 1985).

Columns. We cultivated a stable of regular columnists to comment on a variety of issues. The

In the grand traditions of ACM, there are always people who think we can do a better job.

first was Jon Bentley's "Programming Pearls," (1983), which proved to be the CACM's most popular column of all time. After five years, Jon retired from the job, saying he was burned out from the schedule. "Literate Programming" in 1988 (Chris van Wyck), "Legally Speaking" in 1990 (Pamela Samuelson), "Inside Risks" in 1990 (Peter Neumann), and "Viewpoint" in 1983. Reader surveys told us this was the most popular feature in CACM; the majority of readers turned first to the columns section.

Design. The redesign was a complete overhaul: new typography, stylistic opening pages to articles, illustrations, and professionally designed covers. Our Fifth Generation Computing Systems cover won an award (Sept. 1983). In 1990, we moved all graphic design and layout in-house.

MISSION ACCOMPLISHED

We launched in 1983 with the mission given us by the ACM Council: Transform CACM to a magazine style, embodying the JAM concepts that would be interesting and useful to members every month.

We conducted regular reader surveys and focus groups to help us assess how well we were doing; and we made many adjustments. We continued to be very creative

because the budget was not there to hire the personnel needed to fully realize the mission.

A number of our issues and covers received industry awards.

A recent survey of scientific journals confirmed that CACM is now highly ranked. It has the third-highest citation count across four key computing categories: Software Engineering, Information Systems, Hardware and Architecture, and Theory and Methods. As a result of this increased reputation, the submission rate for good articles has been rising.

We believe we achieved our mission and helped CACM achieve a high stature in the community.

I stepped down in 1992 to chair the Publications Board and lead the Digital Library Project.

WISDOM OF THE AGES

In the grand traditions of ACM, there are always people who think we can do a better job. When David Patterson was president of ACM, many researchers told him they thought CACM had never regained its vaunted glory of the 1970s. Patterson set up a committee to review the current model and propose ways to recharge its content and scope.

When I first talked with the committee, they were not aware that the reason many research

departments had left CACM was the Publications Plan approved by Council in 1978. It was not the work of capricious editors, but of top ACM and SIG leadership.

Moshe Vardi was tapped to spirit this revitalization effort. He spent months gathering feedback from focus groups, studying reader surveys, talking with many individuals, and reviewing every aspect of CACM from bottom to top. A new CACM plan was proposed (see page 44).

It's the same plan we submitted in 1982! Right down to the models envisioned for each section. We thought our plan then—developed through a consensus process—was sound and I am delighted the consensus today is much the same.

There is one major difference. The current ACM leadership has agreed to fully fund the plan. They will be able to hire all the editors they need. No cutting corners. CACM can now become truly great. 

PETER J. DENNING (pjd@nps.edu) is the director of the Cebrowski Institute for Innovation and Information Superiority in the Naval Postgraduate School in Monterey, CA.



ACM AUTHORS TURNED A.M. TURING WINNERS AND ACM PRESIDENTS*

THE GOTLIEB YEARS 1962–1964

Future Turing Recipients

Alan Perlis (1966)
Maurice Wilkes (1967)
John McCarthy (1971)
E. W. Dijkstra (1972)
Donald Knuth (1974)
Allen Newell (1975)
John Backus (1977)
William Kahan (1989)
Fernando Corbato (1990)
Peter Naur (2005)

Future ACM Presidents

Harry Huskey (1960–1962)
George Forsythe (1964–1966)
Anthony Oettinger (1966–1968)
Bernard Galler (1968–1970)
John White (1990–1992)

THE LYNN YEARS 1969–March 1973

Future Turing Recipients

Charles Bachman (1973)
Frederick P. Brooks (1999)
E. F. Codd (1981)
C. A. R. Hoare (1980)
Donald E. Knuth (1974)
Niklaus Wirth (1984)

THE ASHENHURST YEARS April 1973–1983

Future Turing Recipients

Butler Lampson (1992)
Dennis M. Ritchie (1983)
Ken Thompson (1983)
C. A. R. Hoare (1980)
Ronald Rivest (2002)
John Cocke (1987)

Francis Allen (2006)
Nicklaus Wirth (1984)
Adi Shamir (2002)
Leonard Adleman (2002)
Peter Naur (2005)
Robert W. Floyd (1978)
Richard M. Karp (1985)
Ed Feigenbaum (1994)

Future ACM Presidents

Stuart Zweben (1994–1996)
David Patterson (2004–2006)

THE DENNING YEARS 1983–1992

Future Turing Recipients

Robert E. Tarjan (1986)
Jim Gray (1998)
Ronald L. Rivest (2002)
Adi Shamir (2002)

Future ACM Presidents

Stuart Zweben (1994–1996)
Barbara Simons (1998–2000)
David Patterson (2004–2006)

THE COHEN YEARS 1992–1996

Future Turing Recipients

Juris Hartmanis (1993)
Frederick P. Brooks (1999)
Robert E. Kahn (2004)

Future ACM Presidents

Barbara Simons (1998–2000)
Maria Klawe (2002–2004)

*All works were published prior to award/office announcements.