

Data Mining: 100 Feet of SIAM Journals

Thirty-five three-foot shelves. That's 100 feet of SIAM journals, going right back to the beginning in 1953. This is what you find in a little room at SIAM headquarters, and out the window are the skyscrapers of Philadelphia.

I recently spent a happy hour looking over the collection, noting the transition from thin volumes in the early years to thicker, more professional ones nowadays. The first paper published by SIAM, in September 1953, was "Automatic digital computers in industrial research," by R.F. Clippinger, B. Dimsdale, and J.H. Levin of Raytheon Co. Here is their first sentence, the very first sentence ever published in a SIAM journal:

FROM THE SIAM PRESIDENT

By Nick Trefethen

"Seven years have passed since the first high speed automatic digital computer, the Eniac, first startled the world with its 1,000 arithmetic operations per second."

You may marvel at the change from kiloflops back then to petaflops, and soon to be exaflops, today. (Actually, they may have been kilo, but they weren't quite flops!) Personally, I marvel at the fact that even back in 1953, before I was born, machines capable of a thousand operations per second were already years old. Computers were there on the ground floor as part of the original motivation for the founding of SIAM. In fact, the ENIAC was built just a few blocks from the SIAM offices, at the University of Pennsylvania.

SIAM does many things, including publishing books and organizing conferences. But for permanence of intellectual impact, it's hard to beat those 100 feet of journal articles. You can look in any of those old volumes and see classic contributions on some of your favorite topics, by some of your heroes.

Now there was a reason why I went to look at those journals, and this is the question I'd like to ask SIAM members. Is our journal publication process too slow?

I decided to extract some data. I looked at the first four papers in the most recent issue of each of SIAM's ten regular research journals, all lined up on those shelves. (This didn't include the three all-electronic journals, the newest of which is *SIAM Journal on Financial Mathematics*.) I noted when each paper was submitted, when it was accepted for publication, and when it was published electronically.

Here's what I found: The mathematically average paper in this sample of 40 papers has 2.5 authors and is 23.1 pages long. It was submitted on February 13, 2008, revised at least once after the first round of referee reports, accepted for publication on February 27, 2009, and published electronically on June 18, 2009.

That's 12.5 months from submission to final acceptance, 3.5 months from acceptance to electronic publication, 16.0 months in total. And this is my question. Is 16 months from submission to publication OK, or is it a problem? If it's a problem, is there something we can do about it without jeopardizing SIAM's high standards?

Some people tell me that long refereeing times are more or less universal in mathematics, and our culture is just different from that of the engineers and physicists. Some think referees are overworked already and the last thing we would want to do is pressure them still further. Some say the referees would speed up soon enough if we offered a bit of incentive!

What do you think? Does it bother you to wait a few months for referee reports? Is it a matter of culture? Cultures can be changed; should SIAM try to take a lead on this? How? I would be glad to hear your views at trefethen@maths.ox.ac.uk.

"Automatic digital computers in industrial research," by the way, had 3.0 authors and was 15.0 pages long. It was submitted for publication on February 1, 1953, and appeared in print in September. The electronic version was published online about 52 years later.



Photo by Lois Sellers