

## Chair's Message

The Jargon File (<http://www.science.uva.nl/~mes/jargon/t/top.html>) reports that in 1989, hacker Paul Boutin was asked “What do you think will be the biggest problem in computing in the 90s?” Paul’s laconic answer was “There are only 17,000 three-letter acronyms.” (Those hackers among us will immediately counter that, to be exact, there are  $26^3 = 17,576$ .) Do you know that over half (<http://www.atomiser.demon.co.uk/abbrev/>) of the possible TLAs (three-letter acronyms) have been allocated?! I wish to talk about a particular TLA: LPU (and I don’t mean the airport code of Longapung, Indonesia).

I’ve been corresponding with an august group of over a dozen conference and program chairs, trying to solidify the conference schedule for 2002. It turns out to be very difficult to meet the following three objectives we set for ourselves: (1) the notification date for one conference should precede the submission date of the next conference by at least three weeks, to allow for adequate revising, (2) the notification date should follow the submission date by at least ten weeks, to allow for adequate reviewing, and (3) no deadline (submission, notification, camera-ready copy) should be simultaneous with another deadline, with a conference, or with a major holiday. We think we’ve finally gotten it, but it has taken several iterations, and flexibility on the part of everyone.

This negotiation brings to mind the commonly observed phenomenon of repeatedly submitting a paper to conference after conference until it gets accepted. Those who have served on multiple program committees quickly notice papers going through three, four, even five conferences before being accepted, or, much more frequently, the author finally giving up. And our effort to arrange for non-overlapping reviewing periods, where there is the opportunity for rejected papers to be submitted to the next conference, may in fact encourage this behavior.

Why are there so many poor papers? I feel that one cause is a prominent paper-writing strategy called “slicing and dicing”. The objective of this exercise is to get the most papers from the research results of one’s program of study. If the papers are cut too thin, they don’t get accepted. If too much is put in each paper, fewer papers result. The solution is to identify the LPU (*Least Publishable Unit*): that kernel of results that is broad enough to just barely be accepted by a conference or journal. If one is lucky, a research program might have within it four or five LPU’s, which will look very good on one’s vita.

You can peruse most any conference proceedings and identify papers written with this strategy. It is also possible to compare the LPU granularity of various conferences. An LPU, say, of SIGMOD, is easily two or three LPU’s in some other database conferences, not to be named here.

The problem is that this strategy also yields papers that have the *lowest potential utility* (another form of LPU). The world is awash in write-only papers. As an indicator of this, the Anthology ([www.acm.org/sigmod/dblp/db/about/top.html](http://www.acm.org/sigmod/dblp/db/about/top.html)) lists the top 122 referenced papers, out of 100K citations that were analyzed. These papers alone totaled 16% of the analyzed citations. If we assume that every paper references on average 20 other papers, then this portion of the skew alone relegates some 7000 papers to not being referenced by a single other paper. A paper that is never referenced is likely to have little or no impact. What is the use of many papers on a vita if they are never read by anyone else?

I recommend the opposite: that authors try to achieve the MPU (*Maximal Publishable Unit*, or *Most Potential Utility*) with each paper. Try to write a paper that may be the *best* in the chosen conference or journal. This may not maximize one’s vita (at least the quantity in one’s vita), but will help to maximize the impact.

However, life is not so simple (it never is). It is easy to arrive at a paper that has too much content, which reduces the chance of acceptance. My first conference submission was to SOSOP, the top operating system conference (yes, I started out as an O/S guy; it took me a while to see the light). It was rejected, and in retrospect, rightly so. I submitted it to *TOCS*, where it was accepted. Quite simply, the story I was trying to convey with that paper was not tellable in 20 pages, but was possible in 40 pages. So, some results are not appropriate for conferences: they require too much set-up, the proofs are too complicated, whatever. So one must first determine whether a result is appropriate for a conference, for a journal, or even for a book (the story of my last book could not have been told even in a journal paper). Then, if the paper is destined for conference, determine the conference for which the paper represents that venue's MPU, and submit it there. The chances are quite good that the paper will be accepted. Ditto for journals. Ditto for book publishers. There's my unsolicited advice, free for the taking.

Speaking of conferences, I hope that you have registered for the SIGMOD'01 conference, to be held in Santa Barbara on May 21–24, with the co-convened PODS'01 conference and several workshops. When you come, don't forget to bring a spare database textbook to donate, helping disseminate the ideas our community has originated and nurtured. This year, books will be donated to Egypt.

I want to thank Dave Maier, who rotates off the SIGMOD Awards Committee. Dave has served three years on this committee, the last as chair, always providing cogent input and advice. I thank him for his efforts, and also thank Phil Bernstein, who has taken over as chair, and David DeWitt and Jeff Ullman, who join as new members, for agreeing to serve on this important committee. The perceived worth of the SIGMOD awards is directly correlated with the stature of those on the committee, and this year's committee lends great prestige to the awards it names.

The SIGMOD Elections Committee, under the able direction of Meral Özsoyoğlu, has assembled an impressive slate of candidates. It is such a comfort to know that SIGMOD will be in good hands. My term ends in June, and so this is my last chair's message.

My tenure as chair has been truly a wonderful, incredible experience. SIGMOD is full of amazing people, willing to invest generous amounts of time, skill, and passion in activities that benefit others in our community. The first page of this *SIGMOD Record* lists many of these people, but there are also many more working with the conference and other activities. My joy over the past four years has been to interact with such people, and to get to know them as colleagues and as friends. I am deeply grateful for the opportunity you have given me to serve you, and I thank you from the bottom of my heart.

See you in Santa Barbara!

Rick Snodgrass  
January, 2001