

Chair's Message

As I write this, many of you are busy preparing your submissions to next year's SIGMOD/PODS Conference. As you read this, the Program Committees of these conferences will be busy reviewing them. I thought the timing was right to publicly discuss with all of you an issue that I have individually discussed with some of you. The issue is the place of conferences and journals in disseminating our research results. I should warn you that my intention is not to write a well-reasoned essay, but to engage you in some discussion on the subject.

There seems to be an emerging attitude among some that conference publications are "it" and there is no value in journal publications (a slight exaggeration to make a point). I have been told that "the community has spoken" about the choice of conference publications over journals. I have heard that young researchers are discouraged from expanding their conference papers for journal submission in the belief that their time would be better spent on preparing new conference publications. I even heard from some senior members of the community that they don't read journal publications anymore. I fully agree that first class conferences such as SIGMOD, PODS, VLDB, ICDE, etc, are excellent outlets for disseminating new ideas. What I will argue, however, is that they are not always sufficient for a complete and comprehensive treatment of a topic.

So, I'll play the devil's advocate and point out where conference publications fall short and why journal publications are important. In making my points, I will occasionally put on my other hat as an Editor-in-Chief of the *VLDB Journal* and also draw heavily from the *TODS* reports that Rick Snodgrass is publishing in *SIGMOD Record*.

I have two major arguments against relying very heavily on conference publications as the major dissemination outlet. The first is that I do not believe that a every topic can be fully treated within the page constraints of a conference publication. Even though we are now up to 12 pages for SIGMOD (other major database conferences have similar length limits), this is not always sufficient for an in-depth and complete treatment of a topic. How many times have you seen in a conference publication that some experimental results, or some proofs, or descriptions of data structures, or fundamental algorithms are left out due to space considerations? Far too often, I'd say. The result is that readers are expected to make inferences about the workings of the techniques (or have back channel communication with the authors to fill in the gaps). The replicability of experimental results is almost never attempted at least partly because the methodology or the experimental setup is not fully exposed. This can't be good.

One can argue that we should increase the page limit to a reasonable level so that authors can fully expose their ideas. This brings me to my second major objection to relying too heavily on conference publications – the review process. Having served on numerous PCs and having chaired some of them, I know how hard the PC members work to select the best papers for a conference. Their load is almost unrealistically high, and increasing the page limit will only add to that workload. So, that is not a realistic option. If anything, we need to seek ways to control the workload in order to improve the evaluation process (for example, I strongly believe that the load of each PC member should be limited to no higher than 12 papers). The workload also limits the amount of time a PC member can spend on each paper, undoubtedly impacting the quality of the review. There are occasions when the reviewers misunderstand (or don't understand) a paper, but the conference reviewing process does not allow for the authors to be engaged in an iterative discussion with the reviewers. The result is the introduction of an unnecessarily high "chance factor" into the evaluation process.

Journal publications address these issues, but the common claim is that they introduce their own problems. Perhaps the most common complaints are (1) journal publications may not be the same quality as the top conference publications, because it is easier to get papers accepted to journals; (2) the impact of conference publications is not very high since they are mostly “write-only” media; and (3) the review process takes too long. Let me respond to these complaints using my own data about the VLDB Journal and what Rick reported for TODS.

It is not true that it is easier to get papers published in journals. One has to compare apples to apples. In that sense, it only makes sense to compare top journals with top conferences. The acceptance rate for the VLDB Journal in 2001 was around 16% – about the same as the SIGMOD and VLDB Conferences. So, the selectivity level of the VLDB Journal is comparable to the major conferences. Where there is a difference is that the journal review process takes multiple rounds and engages the authors to arrive at a more acceptable resolution (generally).

The impact argument is also suspect. Rick reports that according to the Citeseer impact evaluation, TODS is ranked in the top 4% of the 800 or so journals and conferences. VLDB Journal is ranked in the top 5.6%. The rankings of both journals are higher than many conferences (including top ranking ones). So, there is significant interest in the community in the papers that are published in these journals.

The final argument is about the length of the review process. Unfortunately, this is the Achilles heel of journal publications. Indeed, the review process can be far too long. However, a number of journals are working hard to reduce the review times. Rick reports that the average turn-around time for TODS (from submission to an editorial decision) is 4.3 months. For VLDB Journal, the times are longer: in 2001 it took about 4 months for us to reach the first round decision. However, we are working to reduce this time period to 2.5-3 months. There is no reason why a review should take any longer. (I should also note that we are responsible for the length of these review times and we are individually and collectively empowered to reduce them.) Both journals make the accepted papers available on the Web shortly after they are typeset by the authors. The upshot is that publications in these journals takes only marginally longer than publication in conferences.

There are many things related to journals that need to be improved. Perhaps I’ll touch on those some other time. However, even with their shortcomings, they serve an important function. It is not appropriate, in my view, for the community to discount their role or importance. Conferences are excellent venues for the presentation and discussion of new ideas. However, the papers that finally get published in the top journals are generally more thorough, more complete, and deeper. There is a role for both of them. We have been very successful over the years in convincing tenure and promotion committees and university bodies about the value of the conferences (rightfully so), we now have to convince ourselves that journals are equally valuable and important venues to publish fuller research results.

M. Tamer Özsu
November, 2002