

TODS Special Issues

Richard Snodgrass

rts@cs.arizona.edu

The March 2004 issue of *TODS* should be in your mailbox (as well as available in cyberspace, at the ACM Digital Library) around the time you receive this issue of *SIGMOD Record*. This issue is a first, a special issue dedicated to extended versions of SIGMOD and PODS papers from the 2002 conference. This special issue opens the topic of invited papers from conferences, and indeed, the more general topic of journal versions of conference papers.

A conference paper is the scholarly equivalent of fast food: quick to read (eat), sometimes healthy for you, and undoubtably convenient. The strict page limit (10–12 pages, depending on the conference) favors topics that can be introduced, developed, and evaluated in a short amount of space. One can read several conference papers, on a wide variety of topics, in the time that it takes to read one journal article.

A journal paper is the equivalent of a three-course dinner. It is allowed space (in terms of page count) to more fully examine related work, develop algorithmic or theoretical refinements to the proposed approach, and perform a more thorough evaluation of the central idea of the paper.

Some topics are perfect for a conference paper. The approach is simple and thus can be described fully in just a few pages, the applicability is routine, the related work minimal, and the requisite evaluation straightforward. Other topics are better suited for a journal paper: there may be a substantial prior literature that must be summarized and the positioning of the paper explained, the approach may be complex and require careful development, the evaluation may be quite involved, with many aspects to consider. In fact, some ideas require more space than even a journal article provides (as I have remarked in these pages three years ago¹). The impetus for the last two books that I wrote was exactly that: I knew that to fully discuss my ideas and the rationale behind them, a single journal article, or even two or three related articles, would not be sufficient, but the several hundred pages afforded by a book-length monograph was perfect.

Then there are those ideas that work well in both conference and journal form. The conference paper is a teaser, presenting just enough of the idea and its evaluation to be interesting and to get the idea out there. The journal paper then elaborates on the idea, expounding on exactly where and in what circumstances the idea applies, identifying exactly where the benefits reside and the magnitude of those benefits, and providing a full exposition of the idea, with all necessary detail.

It is this last category of papers that program chairs of SIGMOD, PODS, EDBT, and ICDT are encouraged to nominate for invitation. These nominations are evaluated by the Editor-in-Chief and a relevant Associate Editor to select those conference papers with the most potential for extension, for invitation. The invitation emphasizes the *TODS* prior publication policy², summarized here.

“A submission based on a paper appearing elsewhere must have major value-added extensions to the version that appears elsewhere. For conference papers, there is little scientific merit in simply sending the submitted version to a journal once the paper has been accepted for the conference. The authors learn little from this, and the scientific community gains little.

“The submitted manuscript should have at least 30% new material. The new material should be content material, not just the addition of obvious proofs or a few more straightforward performance figures. The submitted manuscript affords an opportunity to describe the novel approach in more depth, to consider the alternatives more comprehensively, and to delve into some of the issues listed in the other paper as future work. At the same time, it is not required that the

¹<http://www.acm.org/sigmod/record/issues/0103/chairsmg.html>

²<http://www.acm.org/tods/Authors.html>

submitted manuscript contain all of the material from the published paper. To the contrary, only enough material need be included from the published paper to set the context and render the new material comprehensible.”

The subsequently submitted manuscript is reviewed like all other *TODS* submissions, except that the reviewers include a subset of the reviewers of the original conference submission, in part to ensure that concerns raised during that review have been thoroughly addressed. Reviewers are told that the paper was invited and are asked to explicitly evaluate the submission for conformance with the above policy, informally referred to as the “30% rule.”

Invited submissions are not guaranteed to be accepted, even if they meet the 30% rule; in all cases, the paper must be up to *TODS* standards, as interpreted by the reviewers and the Associate Editor handling the submission.

A paper invited to *TODS* is not required to be submitted there; authors are certainly free to submit elsewhere (though few do). Just as importantly, authors of papers appearing in conference proceedings that would benefit from a more leisurely and thorough exposition are encouraged to submit those papers to *TODS*, with the proviso mentioned in the prior publication policy.

“The corresponding author of a *TODS* submission must inform the editor handling that submission about any paper by any author of the *TODS* submission that (a) is in submission, (b) has been accepted for publication, or (c) has been published, that overlaps significantly (more than a page or so) with the *TODS* submission. Such papers in categories (b) and (c) should be referenced by the *TODS* submission and discussed in the related work section, as appropriate. The corresponding author should also inform the editor about any overlaps that occur while the paper is under consideration by *TODS*. In all cases, the Editor will make the determination as to whether the overlap is acceptable.”

The six papers that appear in the March 2004 issue each resemble a wine maker’s dinner, with the courses carefully coordinated and complemented with the selected wine for that course, ending with a delicious dessert of future research questions. For one of the papers, a reviewer mentioned in confidential notes to the editor, “This paper is going to become a classic. Call me back in 10 years for a free drink on me if I am proved wrong.” One of the reviewers of the paper by Torsten Grust, Jens Teubner and Maurice van Keulen stated simply that “I expect [it] will be a widely cited paper in the area of query optimization for XPath/XQuery.” Similar enthusiasm has been expressed for every paper in this issue.

- “Archiving Scientific Data,” by Peter Buneman, Sanjeev Khanna, Keishi Tajima and Wang-Chiew Tan
- “Probabilistic Wavelet Synopses,” by Minos N. Garofalakis and Phillip B. Gibbons
- “Accelerating XPath Evaluation in Any RDBMS,” by Torsten Grust, Jens Teubner and Maurice van Keulen
- “Selection Conditions in Main Memory,” by Kenneth Ross
- “Characterizing Memory Requirements for Queries over Continuous Data Streams,” by Arvind Arasu, Brian Babcock, Shivnath Babu, Jon McAlister and Jennifer Widom
- “A Normal Form for XML Documents,” by Marcelo Arenas and Leonid Libkin

Bon appétit!