A Vision for the *ACM Transactions on Database Systems*

Richard T. Snodgrass

January 29, 2001

This document outlines my view of the current state of *ACM Transactions on Database Systems* (*TODS*), summarizes my experience to date with scientific publishing, identifies several concerns that I feel need to be addressed, and outlines my vision for evolving *TODS* to meet the demands of this internet age.

In brief, *TODS* is the premier journal in the database community. If named Editor-in-Chief, I plan a series of over a dozen initiatives that will further solidify its place, and will make *TODS* an innovator that will outpace all other journals in this community, and indeed, will put *TODS* in the forefront of scientific journals across all disciplines.

1 Current State of TODS

Simply put, *TODS* is in great shape, in terms of its prestige and impact. The Institute for Scientific Information (ISI, `www.isinet.com`) ranks *TODS* the second most impactful[1] ACM transactions, behind *ACM Transactions on Graphics*. It is the premier journal in the field of databases. ISI rates its impact factor as twice that of the *IEEE Transactions on Knowledge and Data Engineering* (*TKDE*). A separate citation analysis of database literature (`www.acm.org/sigmod/dblp/db/about/top.html`) considers over 100,000 citations in listing the 100 top-cited papers. 30 *TODS* papers appear on this list; 31 papers were from all other journals combined, the second largest being the *ACM Computing Surveys*, with 11 papers. This is a strong indicator of the superb quality of *TODS*.

Won Kim, the current *TODS* Editor-in-Chief, has been instrumental in maintaining the quality of this publication. He has reduced the turnaround time (before his tenure, it was between one and two *years*; now it is estimated at perhaps six months, though statistics are not available). Won also instituted the practice of having the program committee from the annual ACM SIGMOD and ACM KDD conferences refer one or two of the top papers to undergo an abbreviated review and accelerated publication in *TODS*.

2 My Experience with Scientific Publishing

I have had a long-standing and abiding interest in scientific publishing, first as a reader, then as an author, then as an editor, and now as a leader within ACM at large.

My first database paper was my TQuel paper in the June 1987 issue of *TODS* ; it is still my favorite. It appears in the above-mentioned list of most-cited papers. I continue to publish actively in this journal.

As an editor of several journals, I have handled an average of one paper a month for the last decade. I have been an Associate Editor for *TODS* for almost nine years, joining the editorial board when Won became Editor-in-Chief (EiC) in July 1992. In that time, I have handled 47 submissions. Four of those were desk rejects, which normally take only a few days. The average turnaround time, from receipt through

[1] ISI's impact factor is based on a citation analysis.

peer review to sending my editorial decision, for the remaining 43 papers (some of which were revisions) averaged 97 days.

I am on the editorial board of the *International Journal on Very Large Databases* (*VLDBJ*) since November 1996 and served on the editorial boards of the *TKDE* for two terms, from July 1995 to December 1999 (I'm still handling a few stragglers for that journal) and of the *International Journal of Computer and Software Engineering* for six years, from January 1991 to December, 1995. For these three journals I have handled a total of 56 papers, six of which were desk rejects. For the remainder, my turnaround time also averaged 97 days. In handling these 103 submissions in total, I have never exceeded four months in rendering a decision, and my average is just over three months[2].

I have been the sole Information Director for *TODS* , having originated the position in March 1996. The *TODS* web site (`www.acm.org/tods`) is one of the most actively maintained, with preprints generally available within a few weeks of acceptance.

I have chaired program committees for SIGMOD and VLDB and been a vice program chair for *ICDE*, and have formally reviewed several hundred conference and journals papers.

I originated the concept of the ACM Computing Portal (`www.cs.arizona.edu/people/rts/Portal/`), which will be a comprehensive collection of bibliographic information on *all* published computer science literature. This project is being overseen by the SIG Governing Board Portal Committee, which I formed and chair. As a first step, ACM is digitizing all of its journals, conferences, and newsletters, which is itself a massive task. The Portal Management Board, which I also chair, is responsible for collecting this material; it is now working to collect hard copy of all ACM conferences prior to 1985.

I was appointed to the ACM Publications Board in August 1999, and appointed Chair in May 2000. Thus far, my primary focus has been finalizing a comprehensive statement of reader/author/editor rights and responsibilities and a wide-ranging strategic plan, both of which can impact in a positive way the future of *TODS*, as I detail below.

As SIGMOD chair, I helped start the *SIGMOD Anthology* and *Digital Symposium Collection*, and have worked closely with the founding editors of these two innovative publications. As a direct result of that effort, a greater percentage of the published literature of the database community has been digitized and made available than for any other area of computer science.

It is my deeply held conviction that publications are among the most valuable benefits of ACM or SIG membership, that publications do much to bring together the scientific community, and that ACM and the SIGs should continue to innovate their publications in order to increase the service they provide to this community. The projects I have embarked on are consistent with that conviction.

3 Challenges

I see several challenges ahead for *TODS* in the next few years.

- *Declining subscription base*
  In July 1995 there were 3101 regular *TODS* subscribers. By March 1999, less than four years later, this had fallen fully by half, to 1686. At that rate, there will be *no* regular subscribers by the end of next year.

  This is somewhat less alarming that it initially appears, because some of the decrease is attributable to institutional Digital Library (DL) subscriptions, which have dramatically increased over that same period. Those at most research universities and major research labs already have access to the ACM DL, and many of those have dropped their individual subscriptions.

[2] My motto is, over 100 in under 100.

Nevertheless, this large drop is troubling, because it reduces the number of people who receive the print copy every three months (people are less likely to go the the library to look at recent copies, or periodically browse the DL), and it is the canary in the mine, indicating less interest in the publication.

- *Inadequate backlog*
  For the past few years, as ACM has changed its production process, *TODS* has been coming out later and later. The June 2000 issue was delivered just a few days ago, in mid-January 2001. The ACM Publications Board has under my direction been working very closely with the ACM publications staff to address late issues for *TODS* and for most other ACM journals. It is expected that the production process will be back to schedule within a few months.

  When that occurs, it will be evident that *TODS* has an inadequate backlog. The publications staff requires all the papers for a particular issue to be delivered by the EiC by the first day of the month four months before the issue date. So the papers for the June 2001 issue are due February 1, 2001, which is only few days from when this was written. But no papers have yet been delivered for that issue, so the backlog is at least three months behind. This was not particularly an issue when the production process was late, but it has now become a problem[3].

- *ACM as the preferred publisher*
  The ACM Publications Board is preparing, under my leadership, a wide-ranging strategic plan to establish ACM as the preferred high-quality computer science publisher. As an initial step, the pubs board has worked with the SIG Governing Board and the EiCs to state explicitly what Rights and Responsibilities it provides and expects. This document (`www.acm.org/pubs/rights.html`) is comprehensive and exacting in what it expects, of ACM and of its journals. Currently, neither ACM nor any of its publications fully implement the rights listed in this document.

- *Irrelevance in this internet age*
  A related concern is the perceived decrease in relevance of "old school" journals in an age of electronic journals, instant news, and a wide variety of (high- and low-quality) free information sources on the web. For the papers that appeared in *TODS* in the last year, the four issues that appeared from September 1999 to June 2000, the total processing time, from initial submission to appearance in print, was 25 months. It is difficult for such a slow and cumbersome medium to compete with conferences (there are five high quality database conferences a year, EDBT/ICDT, ICDE, PODS, SIGMOD and VLDB, and many other workshops and narrower and/or lesser quality conferences) and newsletters (*SIGMOD Record* and *IEEE Data Engineering Bulletin* being the two most prominent), which have a total processing time on the order of six months, and often are freely available on the web. I believe this is a driving factor in the declining subscription base, and implies that even with access to the ACM DL, *TODS* papers are being less frequently read.

How serious are these challenges? It is the impression of many on the ACM Publications Board that current journals will cease print publication in just a few years, and may become irrelevant in just a few more years. Business as usual is simply not acceptable[4].

[3] I note that several other ACM transactions also have inadequate backlogs.

[4] In this way, the situation is similar to that when I started as SIGMOD chair: the organization was highly regarded, yet membership was dropping precipitously, a situation many SIGs still find themselves in.

4 My Vision

I assert that it is a primary responsibility of the next *TODS* EiC to substantively address these challenges, for the journal to remain viable.

In the following, I present how I would address this situation, should I be named the next Editor-in-Chief of *TODS* . Just as the challenges are related (e.g., a declining subscription base reflects the perception of irrelevance), the solutions also interact with and mutually enhance each other.

It is important to emphasize that *all* initiatives must be undertaken in the context of maintaining the very high quality that *TODS* has achieved. Nothing should be done that detracts from the excellent reputation that *TODS* currently enjoys.

I strongly believe in soliciting a wide range of feedback, and so will frequently utilize the Editorial Board to gauge reaction to ideas. In fact, I will have yearly in-person meetings of the Editorial Board, probably at the SIGMOD conference, to sound out ideas and brainstorm on solutions to identified problems[5]. The EB is a valuable resource that has been under-utilized. In between in-person meetings, occasional conference calls of the EB will ensure that everyone is on the same page, and that problems get identified and addressed expeditiously.

4.1 Increasing the Subscription Base

As SIGMOD chair, I started an innovative program in which SIGMOD pays $5.00 per member ($4.00 per student member) for an electronic subscription to *TODS* . This arrangement will *double* the number of subscriptions, in one fell swoop; *TODS* will then have more subscriptions than any other ACM transactions. I then want to send, to each SIGMOD member as well as those with an electronic subscription, brief email messages, alerting them to newly printed issues, providing URLs to the DL for the new articles, and providing other information of interest.

While many EiCs and pub board members are sanguine about decreasing subscription levels[6], I find these decreases alarming. Decreasing subscriptions are not directly the problem, rather, this is a symptom of a more disturbing underlying concern: a lessening interest in the journal itself. Quite frankly, *TODS* has become a journal of a few (about a dozen per year), very long (average length: 40 pages), old (average age: over three years) papers, which are read by very few researchers. And the SIGMOD arrangement will only exacerbate this situation in the short term, because print subscribers who perused the publication when it arrived in the postal mail will no longer have this reminder. *TODS* needs to include more recent work, more and shorter papers, and other features to entice people to look at the journal, while retaining very high quality and material of lasting value.

I will track subscription levels closely, and will judge the efficacy of initiatives to increase reader interest at least partially through their impact on subscriptions. It is important to get *TODS* in front of more eyes. Improvements in the journal will be noticed only if the journal is being read.

4.2 Increasing the Backlog

I will consider a variety of initiatives to increase the number of papers accepted to *TODS*, to ensure a healthy backlog.

[5] Yearly meetings of the SIGMOD Executive Committee, instituted during my tenure, were very effective in developing initiatives for that organization.

[6] The same holds for SIGs. Many SIG chairs were resigned to falling membership rolls, until SIGMOD and SIGCHI showed that it was possible for a mature SIG to reverse the decline and to steadily grow.

- I will suggest an arrangement with the International Conference on Database Theory (ICDT) and the International Conference on Extending Database Technology (EDBT) (held in successive years) to have their program committees suggest one or two of their best papers for accelerated review and publication in *TODS*, similar to the arrangement already in place with the SIGMOD and KDD conferences. This will add one or two papers per year[7].

- A very popular feature I introduced for *SIGMOD Record* is the Influential Papers column, in which well-known authors provide a paragraph or two about a paper that influenced them. For this journal, I will invite authors of influential *TODS* papers from the past to write a 2–15 page retrospective on their original paper, giving their view of the context of that paper, the developments since the paper was written concerning the topic of the paper, and their current view of the material: how well it has held up and what has surprised them about more recent developments. These retrospectives would be shepherded by an editor to ensure appropriate content and style. I feel that this will add interest to the journal, that readers will first turn to the retrospective paper. I envision starting with two per year and increasing to one per issue, or four per year, adding perhaps 40 pages per year.

- *ACM Computing Surveys* publishes high quality tutorials, but looks for tutorials that are broad. I will invite submissions of more focussed tutorials, which are not relevant to *Computing Surveys* but would be relevant to the database community. I would prime the pump with invited tutorials, then transition into submitted tutorials, with the goal of achieving several papers per year.

- Some journals, such as *IEEE TKDE*, *VLDBJ* and *ACM TOMACS*, have special issues. I am reticent to devote an entire issue to a particular topic, but will consider special topic *sections*, of perhaps 50 pages in a 150-page issue.

In the last four issues, 13 papers appeared, averaging 40 pages each, for a total of 514 pages. If the average paper length is reduced to 35 pages (see Section 4.4), these changes in concert will increase a year's yield to perhaps 21 papers and perhaps 600 pages, while simultaneously addressing the backlog and increasing interest in the journal.

4.3 ACM as the Preferred Publisher

As mentioned above, currently no ACM publication fully implements the rights outlined in the Rights and Responsibilities document. *TODS* will lead this effort, by adopting specific approaches and initiatives to be one of the first journals to fully meet the criteria. In particular, the following steps will be taken.

- As mentioned in Section 4.2, the backlog must be increased so that *TODS* is always published on time.

- *TODS* must do a better job of handling submissions. As mentioned before, Won has reduced the turnaround time from something like eighteen months to an estimated six months in most cases. But much more improvement is possible, as indicated by my personal average of less than 100 days. I will prepare an Associate Editor handbook, providing a fool-proof strategy for achieving very fast turnarounds, with little time investment on the part of the AE.

- I feel that it is important that both the average and the maximum turnaround time be reduced. I will insist on a maximum turnaround time initially of six months; once this is achieved, which may take perhaps a year, I will announce a *guaranteed* turnaround time of six months. I will also announce

[7] I'm unsure if this arrangement already exists for PODS. If not, I will also suggest doing this with that conference.

a guaranteed total processing time of 30 months (six months for the first review, six months for a first revision, five months for the second review, two months for the final revision, four months for production, and seven months as a buffer)[8]. The following year, I will announce a guaranteed turnaround time of five months and a guaranteed total processing time of 24 months. At the end of the third year, should the mechanisms and procedures I will put in place be working efficiently, I will announce a guaranteed turnaround time of four months, and a guaranteed total processing time of 21 months. I know from personal experience that this is achievable (my maximal turnaround time was right at four months; my maximal total processing time, excepting a production delay for one paper, was seventeen months).

- I will publish backlog and turnaround time statistics, to let prospective authors know that (a) the guaranteed turnaround time is being met, with an average even shorter than the guarantee, so they can expect a fast review of their submission, (b) the guaranteed total processing time is being met, so they can estimate when their paper will appear, if accepted, and (c) the backlog is about right, so their paper, if accepted, will not reside long in the queue. This is in dramatic contrast to some other journals, such as *VLDBJ*, which claims an absurdly fast turnaround time of three months, but conveniently declines to provide any statistics. Anecdotally, my last *VLDBJ* submission took over *fifteen* months to review, that journal's claim notwithstanding. Authors will want to submit to a journal that makes promises and keeps them. If the number of submissions goes up, the number of acceptances will also go up, helping to address the insufficient backlog.

- Also on the subject of manuscript handling, I will implement with vigor the web-based manuscript tracking system being considered by ACM. I view this as critical to bringing *TODS* into this century. This will enable more frequent and appropriate interaction with authors during the reviewing process, and will reduce the infrastructure each AE needs to provide to keep track of the papers he or she is handling.

- Once a paper is accepted, it is out of the control of the Editorial Board. Nonetheless, I will monitor the production of accepted papers, both via statistics provided by the pubs staff and via interviews with authors, to ensure a fast, smooth and effective transformation from manuscript to printed article.

- Reviewers provide a critical service to any scientific journal. I will put in place a mechanism to ensure that reviewers are not overloaded with requests for *TODS* reviews, and will, when that mechanism is in place and working well, guarantee to reviewers that they will not have to review more than one *TODS* paper every x months (this time frame will be fixed once more experience is available; it will probably be between 8 and 12 months). I feel that reviewers will be more likely to provide fast turnaround if they are promised that there will be a substantial break before the next review is requested.

- The Associate Editor handbook will provide explicit guidance on achieving the other relevant objectives listed in the Rights and Responsibilities document.

My goal will be to fully implement the rights for *TODS* by the end of this calendar year.

4.4 Increasing Relevance in this Internet Age

Most of the initiatives outlined to this point will contribute to increasing the relevance and desirability of *TODS* to its readership and to its potential authors. In this section, I list further initiatives that go even further in this regard.

[8] While the *average* for the last year was 25 months total processing time, that time for individual papers varied from 11 months to 53 (!) months. Getting the *maximum* down to 30 months will require much diligence.

- The manuscript tracking system to be implemented will help in another objective: to move to an "electronic-everything" system. The *TODS* web site will be enhanced to feature four distinct areas, some with restricted access, for readers, potential authors, accepted authors, and editors. It will allow authors of submitted or accepted papers to determine exactly where their paper is in the process, and when the current step will be completed. As mentioned before, this web site will provide current statistics on how well the journal is doing in terms of turnaround time and total processing time, and whether the guarantees are being met. Submission, reviewing, and interactions between (a) the AEs and the EiC, (b) the AE and reviewers, and (c) the AE/EiC and authors will all be via email and the web. The goal is a totally paperless process.

- *TODS* papers are long, often 50 pages or more in final form. Won imposed an absolute limit of 50 pages for a submitted manuscript, a policy I support. But I would like to further reduce the average length of printed papers from 40 pages to 35 pages. I plan to do this by moving material, primarily appendices, proofs, and detailed code fragments, to the electronic version. This ancillary material will be fully reviewed, and will be considered part of the paper, but will not be printed. In this way, *TODS* can retain the detailed and comprehensive coverage typical of its papers, while reducing the length and thus increase the number of papers, and the breadth of concepts and insights, found in each issue.

- *TODS* must expand its repertoire beyond static text and figures. An aspect of the cooperation with SIGMOD discussed in Section 4.1 is including *TODS* on the *SIGMOD Anthology* and *DiSC*, where it will reach many others, e.g., through the library donation program. Here I think there is the opportunity to include with each article two types of ancillary material: the refereed appendices, proofs, and other material, and unrefereed addenda, such a source code, demonstrations, and sample data. The *TODS* material in the DL or on the *TODS* web site must become more dynamic, utilizing modalities other than prose and equations. As disk space is cheap, it is now practical to store a wide variety of ancillary material with papers, even if only the core part of the paper appears in the printed version.

In the longer term, there is the oft-expressed concern that journals will become the dinosaurs of the information age. My view is somewhat to the contrary. I feel that with the democratization of information on the web, more is available, but with a concomitant increase in variability. There is much more good information, but also much more poor or simply incorrect information, on the web. In such an environment, indicators of quality are even more critical. So I believe that the careful reviewing associated with *TODS* , along with its established imprimatur, will be of great value, and will be the basis for the journal's enduring legacy. That said, I think that the web maybe the downfall of lesser-known journals, and there will always be the need to innovate the strong journals for them to prosper.

In addition to the specific ideas outlined above, I plan to peruse other high-quality journals for additional ideas that might apply to *TODS*. As but one example, since 1987 the influential journal *Genetics* has prefaced each monthly issue with a (usually) short historical reminiscence or review under the heading "Perspectives: Anecdotal, Historical, and Critical Commentaries on Genetics." I think perspectives on databases would be an intriguing addition to *TODS*.

Let me emphasize again that I will involve the Editorial Board closely in the conception and implementation of these ideas, as I have done as Chair of SIGMOD, of the Publications Board, and of the SGB Portal Committee.

5 Summary

TODS is a highly regarded and influential journal. The Editor-in-Chief provides leadership and is responsible for running the journal efficiently while maintaining its quality. I regard the position of Editor-in-Chief as both an honor and a daunting responsibility.

I am heartened that the current Editor-in-Chief, Won Kim, has endorsed my nomination; this will help ensure a smooth transition. My close working relationship with Ron Boisvert, the ACM Information Director and Vice Chair of the Publications Board, and with Mark Mandelbaum, ACM Director of Publications and Jono Hardjowirogo, Publisher of ACM Journals, with whom I talk regularly, will also help in realizing these ambitious initiatives I propose.

My tenure as SIGMOD Chair ends June 30. With that obligation met, I pledge that if named *TODS* Editor-in-Chief, I will devote my substantial energy, skills, insight, and passion to this journal which I so respect, and will do my utmost to innovate and strengthen the journal, to meet the challenges and opportunities presented by our times.