Status of *ACM Transactions on Database Systems*

Richard T. Snodgrass

June 7, 2004

This document outlines my view of the current state of *ACM Transactions on Database Systems* (*TODS*), summarizes my experience to date with the journal, identifies several concerns that I feel need to be addressed, and outlines my vision for evolving *TODS* over the next three years.

In brief, *TODS* continues to be the premier journal in the database community. Several pressing problems present when I became Editor-in-Chief have been resolved, following the initiatives I proposed to the ACM Publications Board three years ago. In particular, the inadequate backlog has been resolved, the turnaround and end-to-end times have been significantly reduced, and the number of papers appearing in the journal is growing. I feel that *TODS* is well on its way to achieving my vision of becoming an innovator that will outpace all other journals in this community, and indeed, will be in the forefront of scientific journals across all disciplines.

In my original proposal to the Publications Board[1], I listed four challenges.

- Declining subscription base
- Inadequate backlog
- ACM as the preferred publisher
- Irrelevance in this Internet age

In this status report I reprise those challenges and detail how I have confronted each one.

It is important to emphasize that all of these initiatives were undertaken in the context of maintaining the very high quality that *TODS* has achieved. Nothing was or should be done that detracts from the excellent reputation that *TODS* currently enjoys. I also emphasize that I have frequently utilized the Editorial Board in developing approaches to these problems and implementing the nascent policies and procedures.

In this analysis I present a collection of metrics that in concert characterize important components of the health of the journal. The historical record of these metrics, presented here for the first time, captures the ebb and flow of this publication. It is my belief that the best way to improve any process is to determine adequate means of quantifiably measuring the performance of that process, then put into place initiatives and refinements to improve that performance, regularly revisiting the metrics to determine what works, and to what degree. I welcome recommendations and advice from the Publications Board of additional metrics.

I also welcome any feedback from the Board on ways that I can better fulfill the role of Editor-in-Chief.

[1] "A Vision for the *ACM Transactions on Database Systems*," January 29, 2001.

1 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.

As of April 30 (thanks to Mark Mandelbaum for this information) there are 1194 *TODS* subscribers (688 member subscribers and 506 nonmember—mostly library—subscribers). There were 577 Institutional Members (libraries that get *TODS* as part of getting all the print journals in a package). This is down from 750 in 2001. Of the 577 IMs, 337 also subscribe to the DL. This is down from about 400 in 2001.

These decreases in institutional subscriptions have been more than offset by the spectacular growth of consortia and corporate site-license revenue. As of April 30, 2004 ACM has 43 consortia subscribers, compared to 18 in 2001. At an average of $50K per consortia (and 20 institutions per consortia), that's quite a substantial increase in revenue and in the number of institutions ACM is reaching. Also, in 2001, there were 10 corporate site-license deals; now there are 23.

Given this increase in total number of institutions with a DL subscription, the drop in individual subscriptions is understandable. Those at most research universities and major research labs now have access to the ACM DL, and many of those have dropped their individual subscriptions. It is also comforting that the decrease seems to be leveling off: in the first period (1995–1999), the decrease was 32 (1%) subscribers a month, and only 8 (0.4%) subscribers per month in the second period (1999–2004).

What this means to the journal is that individual subscriptions can no longer be used to gauge interest in the publication. Neither can the number of consortia subscribers, because those represent only highly aggregate data, in terms of institutions, individual readers, and even journals (consortia subscribers subscribe to the entire ACM Digital Library).

It also means that the declining *individual* subscription base is not worrisome, as long as it continues to be accompanied by increases in total number of institutions subscribing to the ACM DL.

Hence, I am now focusing more on submission rates and backlog than on subscription rates as indicators of the interest in the field of the journal.

2 Inadequate Backlog

In the period prior to my taking over as Editor-in-Chief (EiC), *TODS* had been coming out later and later. The June 2000 issue was delivered over six months late, in mid-January 2001. This was the result of two factors: late delivery of issues to HQ, and slow production of issues by the ACM publications staff. However, the staff through changes to the production process quickly addressed the production delays for *TODS* and for most other ACM journals, highlighting the former problem, that of an inadequate backlog for *TODS*.

This problem with the backlog was reflected in the size of each issue. Basically, as soon as the minimum number of papers, three, was available, the issue was produced. While *TODS* was certainly viewed as being very high quality, there was the sense that this was a tired journal that was just holding on, becoming more and more irrelevant in this fast-moving Internet age. Indeed, several in the database community declared journals dead, with all interesting research activity presented at conferences.

So the backlog problem was an indicator of weakness of the journal. This was of great concern to me. Addressing the inadequate backlog has been a central focus for me during my first term as EiC.

The publications staff requires all the papers for a particular issue to be delivered by the EiC by the first day of the month three months before the issue date. So an adequate backlog can be defined as having a sufficient number of papers available for an issue the day the issue is due.

I'll shortly go into detail on how this was done, but let me summarize my performance relative to the three-month requirement for delivery of issues to the publications staff. (I only started keeping detailed records in mid-2002.) The third column indicates the difference between the delivery date (when I accepted the last article for the issue) and the three-month required advance.

Issue	*Delivery Date*	*Delivery Status*	*Backlog Status*
June'02	—	4 months late	—
Sept'02	—	2 months late	—
Dec'02	8/12/02	12 days late	—
March'03	2/11/03	71 days late	87 days late
June'03	3/25/03	25 days late	59 days late
Sept'03	6/17/03	16 days late	47 days late
Dec'03	8/2/03	29 days early	40 days late
March'04	11/10/03	21 days early	39 days late
June'04	1/20/04	40 days early	—
Sept'04	5/7/04	23 days early	—

I am proud of the way that I have gone from a woefully inadequate backlog, in which as issue of three papers was filled months after the issue date, to a healthy backlog of up to six papers an issue delivered in advance.

The fourth column indicates the official backlog status[2]. What is the reason for the disparity between these two columns? I recently was informed that EiC's were responsible for ensuring that the authors sent their final versions to ACM HQ on time (I had the erroneous impression that the pubs staff took over once the paper was accepted). It turns out that some of the authors delayed in getting their final versions in, with the result that the official issue delivery date was about one month later. I have taken steps to rectify this. For example, I expect to have confirmation that all the authors have delivered their papers to ACM for the September'04 issue by this week, making that issue officially about one week late. Future issues will be officially on time.

How was this turnaround accomplished? It is very difficult to change the backlog of an established journal such as *TODS*. What was needed was a sea change in the discipline's perception of the journal. I believe I have effected such a change, through a number of specific initiatives.

2.1 Invited Submissions

I have worked closely with the executive committees of the International Conference on Database Theory (ICDT) and the International Conference on Extending Database Technology (EDBT) (held in successive years) and with the International Conference on Principles of Database Systems (PODS) 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 conference. I should emphasize that all invited papers undergo a thorough review, with most of them requiring a major revision and second round of review.

This effort has resulted in the first SIGMOD/PODS special issue (March 2004), with six papers invited from the SIGMOD'02 and PODS'02 conferences. The December 2004 issue will contain two papers invited from ICDE'03; the March 2005 issue will be a special issue with papers from SIGMOD'03 and PODS'03, and the September 2005 issue will likely contain three papers invited from EDBT'04.

[2]From `http://www.acm.org/pubs/eic/backlog.html`

2.2 Surveys

ACM Computing Surveys publishes high quality tutorials, but looks for tutorials that are broad. *TODS* now has a policy of encouraging focused tutorials surveys, which are not relevant to *Computing Surveys* but would be relevant to the database community. We have started receiving submissions of these papers. The first survey was published in the September 2003 issue, with another one in the following issue.

2.3 Promotion

I have tried to increase the visibility of *TODS* through a quarterly column in *SIGMOD Record*, started two years ago.

- "*TODS* Perceptions and Misconceptions," September 2002
- "Rights of *TODS* Readers, Authors and Reviewers," December 2002
- "ACM *TODS* in this Internet Age," March 2003
- "*TODS* Reviewers," June 2003
- "Journal Relevance," September 2003
- "Developments at *TODS*," December 2003
- "*TODS* Special Issues," March 2004
- "Developments at *TODS*," June 2004

The consistent message of these columns, which range from one to five pages, is that *TODS* offers a useful service to the community and is an attractive place to publish one's papers. I feel that this column has made a difference in the perception of *TODS*.

2.4 Reduced Turnaround and End-to-End Times

As mentioned above, the turnaround time and end-to-end time have both been reduced significantly, which I think has increased submissions considerably.

2.5 Increased Submissions

The result of these efforts is that submissions have roughly doubled over the past three years (see Table 1). Aggregating by year, the last half of 2001 (when statistics were first kept) showed 28 submissions, all of 2002 had 59 submissions, all of 2003 had 91 submissions, and the first five months of 2004 has 41 submissions (note that the second quarter of 2004 is not yet complete).

An increase in submission rate results in an increase in number of papers accepted and appeared, as shown in Figure 1. The average issue of *TODS* in the eighties contained six articles, while in the last seven years the journal has averaged only three articles per quarterly issue, or a paper a month.

In 2003, 15 articles appeared, the most in a decade. And 2004 is looking even better; the first three issues alone comprise 15 papers.

Quarter	*Submissions*
Third 2001	12
Fourth 2001	16
First 2002	11
Second 2002	20
Third 2002	19
Fourth 2002	19
First 2003	24
Second 2003	15
Third 2003	23
Fourth 2003	29
First 2004	15
Second 2004	26

Table 1: Number of Submissions per Quarter

Figure 1: Number of articles per volume in ACM *TODS*

3 ACM as the Preferred Publisher

The ACM Publications Board approved in 2001 a sweeping policy, termed the "Rights and Responsibilities in ACM Publishing" [3], summarizing the rights and responsibilities of readers, authors, reviewers, editors, program chairs and committees, and libraries vis-a-vis journals, transactions, magazines, conference proceedings, and SIG newsletters published by ACM.

This policy document contains a mixture of current practice and goals for the near future. I and the *TODS* Editorial Board have taken these stated rights very seriously and have developed mechanisms, policies, and procedures to ensure each of these rights. As a result of these efforts, *TODS* was the first ACM Transactions to fully implement all 42(!) rights listed in that policy.

I summarize some of the specific steps *TODS* has taken to ensure the stated rights and responsibilities.

3.1 Issue Timely Review and Clear Feedback

Reducing the review time for submitted manuscripts has been a priority.

Response time of submitted manuscripts concerns three related measures. *Turnaround time* is the interval between the submission, usually electronic, of a manuscript or a revision and the sending of the editorial decision, now almost entirely by electronic mail. The *queue time* is how long accepted papers remain in the backlog, waiting to appear. And the *end-to-end time* is the interval between the original submission of a manuscript and the appearance in print of (generally a revision of) that manuscript. (An alternate definition uses the time the paper appears electronically, though that time is much harder to determine for papers in the past, and so is not reported here.) I consider each in turn.

Won Kim, the previous EiC, reduced the turnaround time from something like eighteen months to an estimated six months in most cases. But I felt that more improvement was possible. I prepared 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, queue time, and end-to-end time be reduced. In my original proposal, I stated the following goals.

- A guaranteed maximum turnaround time of six months.
- A guaranteed maximum end-to-end 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). Further reductions were mentioned as later goals.
- Published backlog and turnaround time statistics.

How have I done with respect to these goals?

Figure 2 shows various aspects of turnaround time[4]. The turnaround time has been slowly decreasing over the time I have been EiC. This figure shows four sets of data. The bottom line is the *average turnaround time*, a moving average of the turnaround time for papers submitted in the indicated month. To smooth monthly variations, the moving average includes all of the submissions for the previous year. Each data point represents dozens of papers. The value for January 2004 (the latest date for which statistics are available), 3.0 months, is the average turnaround time, that is, the average time required to process a manuscript from submission to editorial decision, for all of the papers submitted between (inclusive) February 2003 and January 2004.

[3] `http://www.acm.org/pubs/rights.html`

[4] This figure is from `http://www.acm.org/tods/TurnaroundTime.html` where it is updated monthly.

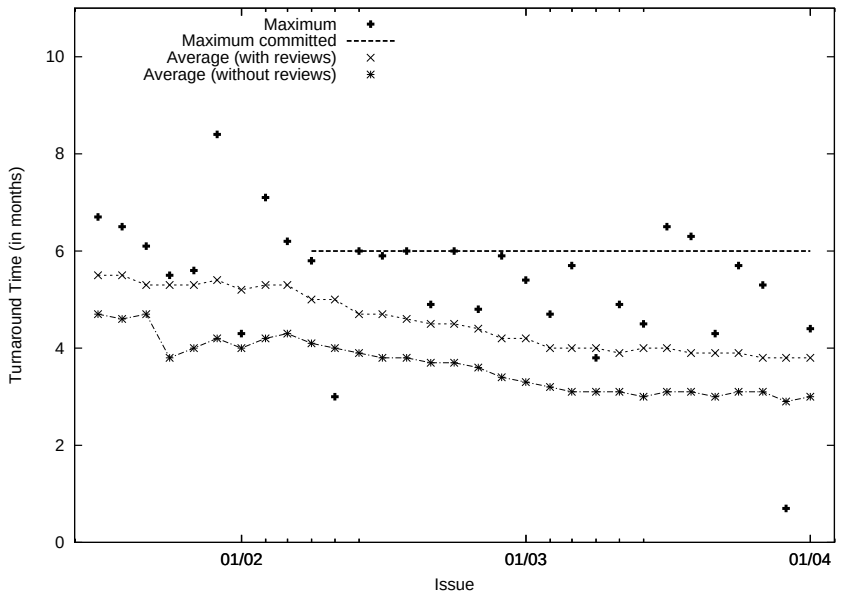

Figure 2: ACM *TODS* Turnaround Time

The next line up is the average turnaround time for external reviews only, a moving average of the turnaround time for papers submitted in the indicated month. This includes only submissions that went out to external reviewers and specifically excludes desk accepts and rejects. The value for January 2004, 3.8 months, is the average turnaround time for external reviews of all the papers submitted during the year up through January 2004.

The straight line is the *committed maximum turnaround time*, the boundary that the Editorial Board has committed to not exceed, for any submission. The Editorial Board recently established a formal policy stating its commitment to providing an editorial decision within 6 months. I don't know of another journal that has been willing to publicly announce such a commitment and then provide statistics indicating compliance.

The individual points, one per month, denote the maximum or peak turnaround time for submissions in the indicated month. Each point represents a single, unusually slow paper submitted during the indicated month. For all the papers submitted in January 2004, the longest turnaround time was 4.4 months. For the year this commitment has been in place, only two papers have required more than six months (and then, only a few days more).

I now turn to queue time, a metric that has been discussed in the December 2003 meeting of the Publications Board, specifically for *TODS*.

Given that production of an issue requires three months, the absolute minimum queue time is three months. This occurs only when all the papers in the issue are accepted on the last possible day, and enjoy instantaneous preparation of the final version by authors. A more realistic minimum assumes that papers are accepted uniformly across the three months before the deadline, and require two weeks each for authors to prepare the final version. These assumptions imply a goal of an average queue time of five months.

Listed below is the acceptance date for the first paper accepted for each issue, and the length of time between acceptance and publication of that paper. This represents the maximum (not average) queue time for papers in that issue.

Issue	*Acceptance Date*	*Maximum Queue Time*
June'02	4/8/02	2 months
Sept'02	2/7/02	7 months
Dec'02	7/23/02	5 months
Mar'03	9/12/02	6 months
June'03	2/11/03	4 months
Sept'03	4/6/03	5 months
Dec'03	6/27/03	6 months
March'04	3/31/03	11 months
June'04	8/21/03	10 months
Sept'04	2/3/04	7 months

For most issues, the maximum queue time is around the goal of five months. (Those issues with a queue time less than about five months were finalized after the production deadline, and so represent anomalies.)

The maximum queue time has always been less than twelve months. Equivalently, no paper was accepted a year ahead of its publication. However, the the March'04 and June'04 issues have maximum queue times which are significantly longer than the other issues.

The March'04 issue was a SIGMOD/PODS special issue, with all the papers coordinated. That one paper mentioned here was accepted after one review; it had to wait for the other invited papers to have a second review (the next earliest paper for that issue was accepted on 8/6/03, 8 months before the issue came out).

The June'04 issue followed this special issue. Since the special issue was full, papers that were accepted too late for the Dec'03 issue had to be delayed for the June'04 issue. The next earliest paper for that issue was accepted on 10/10/03, 8 months before the issue came out. Indeed, three of the papers in the June issue were delayed by this special issue.

I now turn to end-end time, which comprises the turnaround time for each cycle, the time for the author(s) to prepare zero, one, or several revisions, the time the paper sits in the queue waiting for a slot in an issue, and the time for the publisher to copy edit, typeset, proof and print the paper. Figure 3 shows the data for *TODS*, calculated from the submission date as indicated on the last page of the article in the journal and from the cover month of the issue[5]. This data does not take into account that some issues over the past few years were printed late, nor does it include data for the first volume, as papers in that volume do not have a "submitted on" date.

The end-to-end time started at 11 months in 1977, grew to almost 42 months in 1991 (can you imagine waiting over three years for your paper to wind its way through the reviewing and production process?!), then fell in spurts, to 27 months in 2000.

I have been proactive at reducing the revising time (a maximum of six months for major revisions, down to one month for simple changes); of course the reduction in turnaround time also helps. The end-to-end time is now 18.3 months, the lowest in over twenty years.

As with turnaround time, maximum end-to-end time, also shown as individual points in Figure 3, is equally important. After all, the individual author cares about how long his or her paper *might* take.

The longest end-to-end time for *TODS* was an incredible 73 months, or over six years. Even the past five years (1999–2003) has been disturbing: 53, 57, 55, 55, and 37 months. I have been responsible for part of 2001 and for 2002 and 2003. However, all of the papers with very long end-to-end times were submitted before I was appointed as EiC.

[5]This figure is also on `http://www.acm.org/tods/TurnaroundTime.html`.

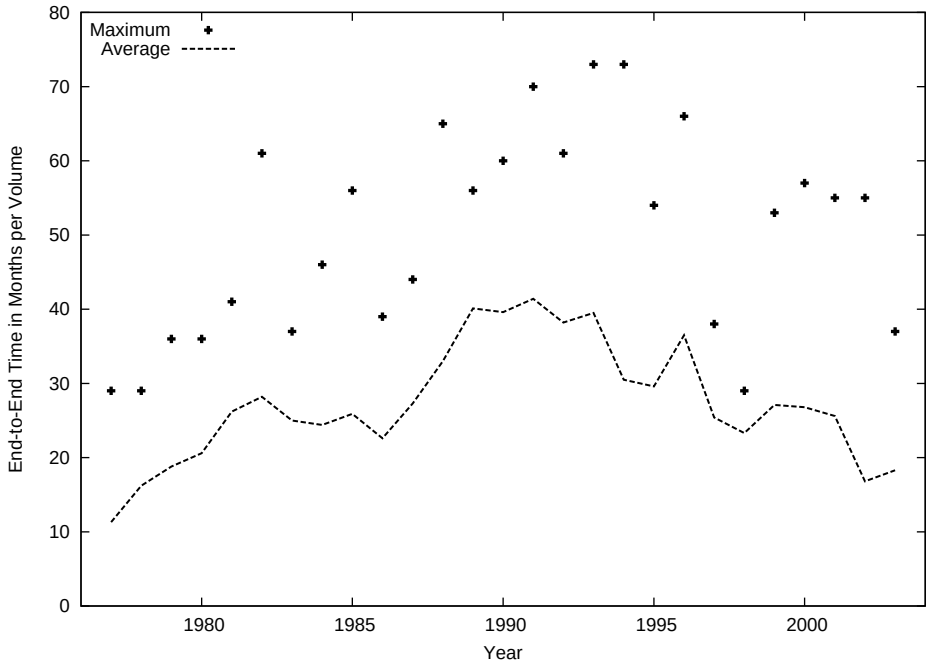

Figure 3: ACM *TODS* End-to-End Time

For papers initially submitted when I was EiC, the longest end-to-end time for 2002 was 13 months and for 2003 was 23 months.

3.2 Provide Statistics

Another item in the Rights and Responsibilities policy is to provide statistics for each journal, transaction, and newsletter on its average turn-around time and its current backlog of articles.

TODS is the first database journal, the first ACM journal, and indeed the first journal of any discipline that I am aware of, that publishes its turn-around performance[6] for all to see.

The current backlog of articles is also maintained[7]. Authors can judge for themselves how responsive the journal is.

3.3 Use Referees Sparingly

Another item is to request referees to review only submissions for which the editor feels they have expertise, and request only a limited number of reviews over the course of a year. *TODS* strives to not overload referees. Specifically, *TODS* now has an explicit policy that referees will be expected to review at most one *TODS* paper in any twelve-month period.

Yet another item it to give a reasonable length of time for a review, where the particular length of time depends on the publication. *TODS* now has an explicit policy to allow at least two months for an initial formal review.

[6] `http://www.acm.org/tods/TurnaroundTime.html`
[7] `http://www.acm.org/tods/Upcoming.html`

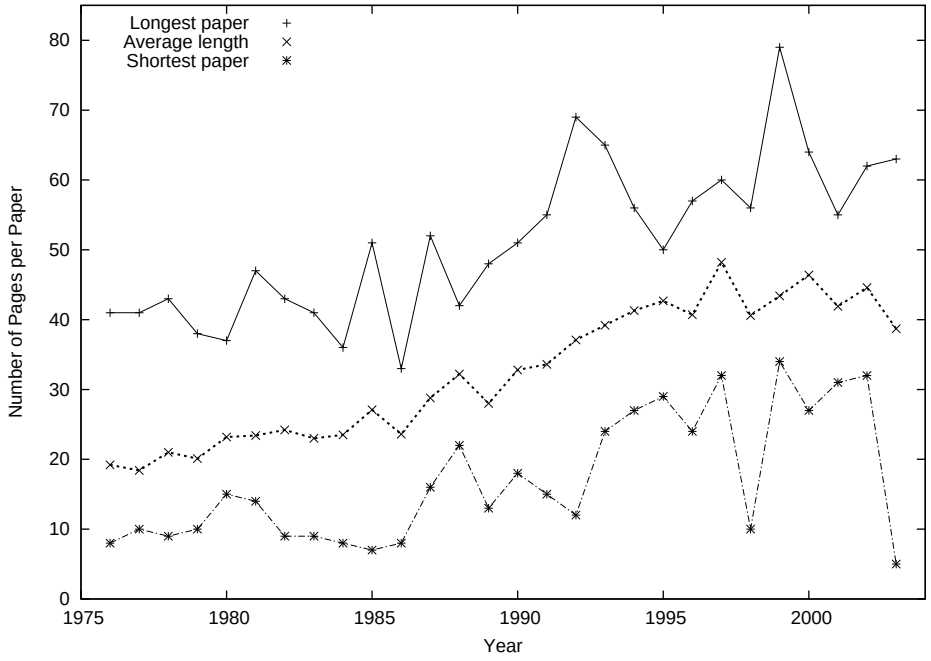

Figure 4: Article length per volume in ACM *TODS*

3.4 Have Clear, Written Policies

The Editors-only portion of the *TODS* web site has a description of the fifteen significant policy discussions of the Editorial Board over the last three years, along with a comprehensive 19-page Associate Editor Manual, with procedures and more detailed policies.

A full analysis of how *TODS* ensures all the rights and responsibilities put forth in the ACM policy may be found in the December 2002 issue of *SIGMOD Record*.

4 Irrelevance in this Internet Age

The fourth concern raised in my original proposal was 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.

This concern was one of the drivers for my efforts to reduce average and maximum turnaround time, queue time, and end-to-end time. I also endeavored to increase the relevance and desirability of *TODS* to its readership and to its potential authors, through my column in *ACM SIGMOD Record*.

A related concern was that of article length. Readers nowadays are reticent to wade through long articles.

Figure 4 provides the historical record for article length in pages for *TODS*. The top line states the length in pages of the longest article in each yearly volume, the middle line indicates the average length, and the bottom line states the length of the shortest article.

Quite frankly, all three trends are disturbing. The average article length has more than doubled, from 19 pages in 1976 to a peak of 48 pages in 2000. The average article that year was longer than the longest article in 1976 (at 41 pages). The shortest article that year, at 27 pages, was longer than the average article for the

entire first decade of *TODS*' existence. In seven separate years an article of at least 60 pages appeared (one weighed in at a whopping 79 pages).

As can also be seen in these graphs, we've made headway in the average paper length (down to 38.7 pages for 2003, the lowest in a decade) and shortest paper (down to five pages, the shortest in the history of *TODS*, due to the new policy we adopted to encourage short papers that nevertheless contribute to the state of the art).

The longest paper (63 pages in 2003) is still too long. Most readers are simply not willing to slog through a paper that is in reality a short monograph.

Another aspect I've emphasized is to augment articles with two types of ancillary material: the refereed appendices, proofs, and other material, and unrefereed addenda, such a source code, demonstrations, and sample data. 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. The *TODS* material in the DL or on the *TODS* web site is becoming more dynamic, utilizing modalities other than prose and equations.

Curtis Dyreson, the *TODS* Information Director, has been proactive at collecting these materials. Examples can be found at the following links (some of these papers were processed before my term as EiC).

- `http://doi.acm.org/10.1145/502030.502031` (a technical report)
- `http://doi.acm.org/10.1145/503099.503102` (a URL to a project)
- `http://doi.acm.org/10.1145/352958.352963` (an extended version of the paper as well as citations for two additional papers)
- `http://doi.acm.org/10.1145/383734.383737` (another URL with related technical reports and data)
- `http://doi.acm.org/10.1145/288086.288087` (additional documentation and code)

5 Goals

I have set the following goals for my requested second term as EiC.

Backlog I plan to deliver all of the issues during my second term on time, by the official metric.

Turnaround time My goal is to keep the average turnaround time consistently under three months and the average turnaround time for papers going to reviewers under four months. I would also like to announce a guarantee maximum turnaround time of *five* months, retroactive to January 2004.

End-to-end time My goal is to keep the average end-to-end time consistently below eighteen months (not seen since 1978). This will be accomplished in part by reducing the maximum revising time to five months, shadowing the maximum guaranteed reviewing time. I would also like to announce a guaranteed maximum end-to-end time of 30 months, and later reduce that further to 27 months, which has never been achieved with *TODS* before.

Number of papers per issue My goal is to increase the number of papers per volume back to that in the heyday of *TODS*, of 24 papers per year (not seen since 1985).

Maximum article length My goal is to bring this down from about 65 pages to 50 pages (not seen since 1995). However, I do not want to negatively impact quality. So I will work to adopt this policy, with the mechanism being to encourage more use of electronic appendices. Thus the paper version will be shorter, but the DL version will have all the details, should the reader desire that.

Average article length My goal is to bring this down to 33 pages (not seen since 1990), by reducing the maximum article length and encouraging electronic appendices for many papers.

Expanding the DL content My goal is to increase the content even more, to include more software, data, and other resources.

My original proposal mentioned three initiatives that I haven't yet gotten to.

Retrospective papers I plan to pursue this within the following year; doing so will require close interaction with the Editorial Board.

Special topic sections I'm now less enamored of the special topic section possibilities, especially since the issues are now of healthy length.

Other journals 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*.

6 Page Budget

I end with a request.

With the average paper length going down (cf. Figure 4) but the number of papers per volume going up (cf. Figure 1), what is the impact on the total page count per volume? Figure 5 depicts the page count of each volume, with the current page budget of 530 pages also indicated.

Several points can be made. First, either the page budget in the past has been larger, or past EiCs have been notorious in not following their page budget. Second, despite the increase in the number of papers for 2003, the 2003 volume was still under 600 pages, because of the reduced average article length. Third, the average number of pages for the volumes I have produced (2001–2003) is 525 pages, so I have (thus far!) kept within my page budget.

The planned increase to 24 papers a year at 33 pages each (on average) will result in a yearly output of 792 pages. Given this projected doubling of the number of papers per volume, I request an increase of the page budget of 260 pages, to 790 pages per year.

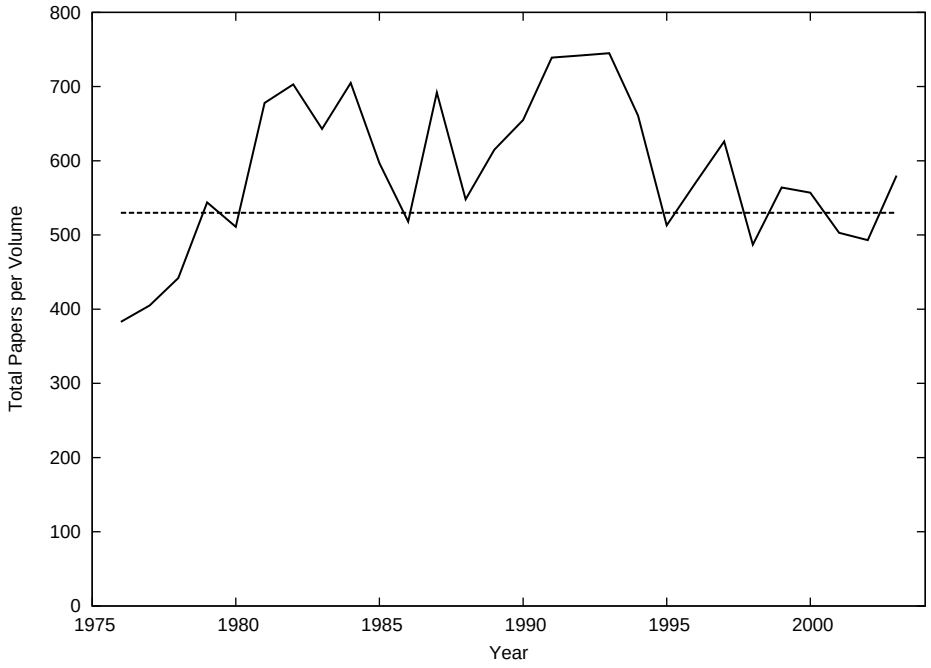

Figure 5: Number of pages per volume in ACM *TODS*

7 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 have worked hard to convert *TODS* from a high-quality but tired journal with an inadequate backlog and thin, late issues into a vibrant journal with a doubled submission rate, full issues delivered on time, and a responsiveness and transparency that is matched by few if any computer science journals. I'm especially proud of the *TODS* Editorial Board for reducing the average turnaround time to three months, and for upholding their commitment to not exceed six months, while simultaneously maintaining the imprimatur that has always been a hallmark of *TODS*, and indeed, of all the ACM Transactions.

I request that I be appointed to a second three-year term to stabilize the improvements thus made to the journal and to make progress towards the goals I list in Section 5. I also request that the page budget for *TODS* be increased to 790 pages per year.