

Chair's Message

One of the fundamental questions of philosophy, physics, and yes, even databases, is, what is the structure of time? Must it go in only one direction? Is it linear, branching, cyclic, even multi-dimensional? I'll invoke most of these aspects in this description of where SIGMOD has been and where it is going.

Imagine you are planning your wedding, perhaps a small affair, of say 50 people. Perhaps it will be an afternoon wedding, followed by a dinner and dancing. Still, there are lots of things to decide: the meal, entertainment, renting the hall, centerpieces, etc., etc.

Your dream darkens. The wedding grows and grows, to 500 of your closest friends. The afternoon and evening grow to a 5-day event. The cost grows to a quarter million dollars. You wake up sweating, with your heart pounding. Fortunately, it was just a nightmare.

Now you know what has been keeping Maggie Dunham busy this past 15 months, as she planned the SIGMOD'00 conference. Those of you who went know that Maggie did a fantastic job. Thanks Maggie, for your energy, your handling of a million details, and your desire that the week of the conference would be enjoyable and productive for everyone.

The registration fee in Seattle two years ago was \$525. At that meeting, Moshe Vardi asked that the registration fee be lowered. Maggie was able to get the registration down to \$400, which is the same as it was some 6 years ago, before high resolution projectors were needed in each room, before there was an email terminal room, and before the keynote addresses were video-taped. Getting the registration fee down while continuing to provide the facilities that are now expected was indeed an impressive accomplishment.

We are hoping that our example keeps the lid on conference registration fees for other database conferences. So my challenge to the other conferences is, do as we did, and drop your registration fees \$100. Let's have a race among database conferences to the lowest registration fee!

Almost 20 years ago, another nightmare: one of my papers was rejected by a conference, in fact, by SIGMOD. I was devastated.

I'm reminded of talking to a colleague who was on the graduate admissions committee during the Vietnam war. He found these decisions agonizing, because denying an applicant was in some cases tantamount to sending that person off to die.

Well, a SIGMOD rejection isn't that serious, but it definitely has quite negative repercussions. A few SIGMOD papers can do wonderful things for one's career; the converse also holds.

And rejections are common. For every paper that Jeff Naughton accepted, he had to reject 5 others. Being PC chair means almost always having to say *NO*. (Jeff reiterates this observation in his comments.)

So it was critical that the strongest papers be accepted. And Jeff did a superb job of choosing the right people for the program committee and of managing the process and the PC meeting, so that the very best papers were accepted.

But Jeff's job was even harder. When I asked him to be PC chair, I told him he was responsible for the entire technical program: the panels, the tutorials, the demos, everything. And as those who were there can attest that the program was very strong indeed this year.

Next year, SIGMOD will be in Santa Barbara. Jianwen Su will be General Chair, and Timos Sellis, Program Chair. Do plan on joining us in sunny California, May 21-24.

So how is SIGMOD doing? On the next few pages are reports from Meral Özsoyoğlu, the Vice Chair, Joachim Hammer, the new Treasurer, and others on how various projects are going.

But first, I wanted to highlight just a few activities of this organization. Our focus has been on building a worldwide database community, with four specific goals.

First, we want to emphasize the central place of databases within CS. To some, databases is administrative data processing, which is boring, boring. To others, databases is the Y2K problem, boring, boring. To me, databases represent absolutely core technology that has succeeded so well that everyone takes it for granted.

The SIGMOD awards celebrate the achievements of those in our community. This year's SIGMOD Innovations Award went to Rakesh Agrawal, the SIGMOD Contributions Award, to Michael Carey and Laura Haas, the SIGMOD "Test of Time" Award to two papers that appeared in the SIGMOD'90 conference a decade ago: "Encapsulation of Parallelism in the Volcano Query Processing System," by Goetz Graefe, and "Set-Oriented Production Rules in Relational Database Systems," by Jennifer Widom and Shel Finkelstein, and the SIGMOD Best Paper Award, to Hartmut Liefke and Dan Suciu, for "XMill: an Efficient Compressor for XML Data." My congratulations to all these deserving awardees.

And there are ACM-wide awards. As you all know, we celebrated last year Jim Gray winning the ACM Turing Award. This May, some 30 ACM Fellows were inducted. Over a fourth were from our community, which made me very proud: Peter Buneman, Mike Carey, Ron Fagin, Don Haderle, Hank Korth, Raymond Lorie, Moshe Vardi and David Warren.

Another goal that I have pushed throughout my tenure as chair is archiving and distributing technical material, which has always been a central role of SIGMOD, a role that has expanded in this age of digital libraries.

I thank Isabel Cruz for the amazing job she did on the DiSC'2000, which was released just a few months ago. Isabel is already hard at work on DiSC'2001, which will be even more impressive, in that it will include the proceedings of at least a dozen conferences held this year.

Michael Ley is also working very hard on the next volumes of the Anthology, which will be sent (free, of course) to all SIGMOD members in the fall, some 8 additional CDROMs. We've added some 20,000 pages to what was announced at the last business meeting:

- CIKM Conference on Information and Knowledge Management (Niki Pissinou)
- DBPL Databases and Programming Languages (Alberto Mendelzon and Sophie Cluet)
- PDIS Parallel and Distributed Information Systems (Sushil Jajodia)
- MFDDBS Mathematical Fundamentals of Database Systems (Jan Paredaens)
- MobiDE (Panos Chrysanthis)
- SIGIR Conference (Susan Dumais)
- TKDE (Betty Salzberg and Ahmed Elmagarmid)
- Foundations of Databases, a book by Serge Abiteboul, Rick Hull, and Victor Vianu

Thanks to all these people for their help in convincing publishers that the community values including their publications in the Anthology. Along with first volume, the 13 CDROMs will total over 100,000 pages of material.

Even better, next spring everyone will receive the *Silver Edition* of the Anthology, on DVD. This will comprise everything through 1999, about 70% of all the refereed material on databases that has appeared over the last 25 years. This includes about 250 volumes of material, a large bookcase totally full of journals, proceedings, and books, all on one or two DVD disks. Isn't storage technology amazing?!!

Much of this material turns out to be hard to find. We wanted to grab this stuff before it literally disappears. Even ACM and the IEEE keep only a few years worth of important material. So, one side benefit is that the Anthology will serve as an important historical record of our field—a record that we will now have for all time.

The third way we are building a worldwide database community is by encouraging other, related disciplines to also make their material available. Databases is a broad field: we use results from theory, operating systems, networking, hardware, and information retrieval. Hence, it is important that these related fields also capture their literature for easy access.

We have been encouraging other SIGs to follow SIGMOD's lead, with some success: ACM and all the SIG chairs have recently agreed to pay to digitize *everything* ACM has published in its 60 years of existence, back to 1940. By summer, 2001, the ACM Digital Library will include all journals, all conference proceedings, and all newsletters. This will cost in excess of \$500,000. It will cost SIGMOD very little, because most of our material has already been digitized, with the rest being completed within the next few months.

SIGMOD is also encouraging other publishers to “do the right thing”, in making their material as available as possible. We are working to get all of the Anthology onto the web, in many cases freely accessible.

Finally, the fourth way we are building a worldwide database community is to partner with other societies. The March 2000 issue of *SIGMOD Record* listed the many societies we cooperate with. Our regional partners include societies in Russia, Japan, and China. We are now working with those in Latin America to strengthen our ties there. Claudio Bauzer Medeiros chairs the Latin American Liaison Committee. We are also reaching out to the industrial community. Daniel Barbara has formed an Industrial Advisory Board to help there. See the first page of this issue for the members of these two committees.

I'm very appreciative of all the fine work that these people, and indeed, some 100 SIGMOD volunteers all over the world, are doing.

Rick Snodgrass
July, 2000

Vice Chair's Message

I am very glad to report that, two years in a row, SIGMOD is the only SIG (among those ACM SIGs with 2000+ membership) with an increased total number of members. There has been a continuous decline over the last decade in SIG membership, and SIGMOD has also been losing members. Last year, we were excited to see the reversal of this trend for SIGMOD, and the increase in membership for the first time in the last thirteen years. This year's increase is even more significant (more than 500!). In addition to SIGMOD Anthology and DiSC which have been very instrumental to boost membership, successful SIGMOD/PODS conferences, new conference registration fee structures that promote new members, and SIGMOD financial support program for undergraduate students to attend the conference, all helped in reversing the trend of decline in membership and resulted in a healthy increase. More than 100 new members joined SIGMOD with the SIGMOD/PODS 2000 conference at Dallas in May. As of May 2000, SIGMOD has 2939 members; it is the fifth largest among all 37 ACM SIG's and is still climbing up (it was the seventh largest last year). Another good sign for strong membership is that SIGMOD has the highest member retention rate between large SIG's (the top ten largest). SIGMOD member retention rate, 69%, is also one

of the highest of all ACM SIG's (see the full report at <http://www.acm.org/sigs/membership>). It is very likely that we will see SIGMOD to continue climbing up in the ladder in the list of largest SIG's in the coming years as well.

I also want to mention here that SIGMOD has recently adopted new guidelines for conference approvals emphasizing the quality of the conference, among other requirements such as non-conflicting time and venue with the SIGMOD conference. Conferences with SIGMOD in-cooperation support are also required to have their conference proceedings to be included in SIGMOD DiSC. These guidelines are posted in the SIGMOD web site <http://www.acm.org/sigmod/sigmodinfo>. If you have suggestions, or questions about the membership and conference approval, please contact me.

Meral Özsoyoglu
July, 2000

Treasurer's Message

SIGMOD's financial situation continues to be very healthy despite the heavy investments in our special member projects. Under the current fiscal year FY 2000 budget we have funded activities in excess of \$58,000 aimed at improving the value of the SIGMOD membership. For example, we funded the SIGMOD Anthology (total cost: \$40,000), DiSC (total cost: \$13,000), and gave out special awards and scholarship grants totaling \$5,000. We also froze membership dues and, for the second year in a row, the SIGMOD conference chair reduced registration fees for our conference.

Since our annual operating expenses (\$81,000) and revenues (\$80,000) are roughly balance, the above mentioned investments and services continue to have a direct impact on our fund balance. Our expected fund balance at the end of FY 2000 will be about \$275,330, which is down from \$336,433 in FY 1999, but still roughly \$63,000 above the required minimum fund balance mandated by ACM. However, given the fact that we are the only SIG with an increase in membership, this \$58,000 investment is money well spent.

In February, the executive committee passed the budget for FY 2001 which will go into effect in July 2000. As in previous years, we made every effort to reduce annual operating costs while continuing to invest portions of our funds to improve the value of the SIGMOD membership. Given the recent increase in members, we are also making a special effort to improve our retention rate. In all, we have allocated \$86,000 to fund special projects and expenses which include, for example, a membership drive (projected cost: \$10,000), non-US SIGMOD development in Asia, India, and South Africa (projected cost: \$8,000), continuation of anthology and DiSC (projected cost: \$54,000), awards and grants (projected cost: \$5,000), and undergraduate scholarships (projected cost: \$2,000).

Although our expected annual revenue in FY 2001 will be higher due to increased memberships (\$100,891) while annual operating expenses remain roughly the same (\$85,555), the additional income does not completely offset the cost of the special projects. Therefore, the expected fund balance at the end of FY 2001 will be around \$200,000. This is still about \$60,000 above the required minimum from ACM (recalculated annually) and considered a healthy fund balance for a SIG of our size. Given all the exciting and worthwhile projects we are undertaking, we are confident that we are investing our members' money in the best possible way. One of the challenges we must face when making the next budget is to continue providing our members with the highest possible membership value while at the same time preserving the fund balance.

Joachim Hammer
July, 2000

Information Director's Message

My previous message to the members, back in 1998, stated: "the first step in building the new SIGMOD Online will be to revamp the web site, with help from a professional design firm." We finally got this redesign done and are very pleased with the results. Please take a look at www.acm.org/sigmod and let us have your feedback on the new look-and-feel.

Content at SIGMOD Online continued to grow in the past two years, from 100 megabytes in June 1998 to 2.2 gigabytes in June of 1999 to 4.4 gigabytes in June of 2000. Some highlights of the content added from 1999 to 2000 include:

- full text of the VLDB Proceedings from 1982 to 1999, thanks to Michael Ley and the VLDB Foundation;
- full text of the SIGMOD 2000 and PODS 2000 Proceedings, available to members weeks before the conference, thanks to Nick Koudas and Bertram Ludascher respectively;
- Alex Labrinidis's **dbGrads** registry of graduating database students, with over 800 registered users and over 300 students listed as of last June;
- a collection of PODS-related pages created and maintained by Leonid Libkin;
- Dongwon Lee's **dbEvents** calendar of calls for papers;
- an XML version of the *SIGMOD Record*, thanks to Paolo Meriardo;
- the online version of DiSC 2000, thanks to Isabel Cruz.

Several exciting initiatives are in the works for the coming year, including a job posting repository, a directory of experts in database-related areas, a set of industry pages, and Volume 2 of the SIGMOD Anthology. Many thanks to the volunteers mentioned above and the many others who contributed to the expansion and development of SIGMOD Online.

Alberto Mendelzon
June, 2000

SIGMOD Anthology Editor's Message

1. As you read this report, the CDROMs of volumes 2, 3 and 4 of the ACM SIGMOD Anthology should be in production. For the most recent information check

<http://www.acm.org/sigmod/dblp/db/anthology.html>

or

<http://www.informatik.uni-trier.de/~ley/db/anthology.html>

2. The bibliographic records of the DBLP server are available in an XML-style format; you may download them from `ftp://ftp.informatik.uni-trier.de/pub/users/Ley/bib/records.tar.gz`

The records look like

```
<article key="AgrawalS93">
<author>Divyakant Agrawal</author>,
<author>Soumitra Sengupta</author>,
<title>Modular Synchronization in Distributed,
Multiversion Databases: Version Control and
Concurrency Control.</title>,
<pages>126-137</pages>,
<year>1993</year>,
<volume>5</volume>,
<journal>TKDE</journal>,
<number>1</number>,
<url>db/journals/tkde/tkde5.html#AgrawalS93</url>
<ee>db/journals/tkde/AgrawalS93.html</ee>
<cite>...</cite>
<cite>journals/dc/AgrawalBGS87</cite>
<cite>journals/csur/BernsteinG81</cite>
<cite>journals/tods/BernsteinG83</cite>
<cite>conf/podc/BernsteinG83</cite>
<cite>books/aw/BernsteinHG87</cite>
<cite>conf/sigmod/ChanFLNR82</cite>
<cite>journals/tse/ChanG85</cite>
<cite>journals/cacm/EswarranGLT76</cite>
<cite>journals/tods/Garcia-MolinaW82</cite>
<cite>conf/sosp/Gifford79</cite>
<cite>conf/ac/Gray78</cite>
<cite>journals/ins/HedayatHW89</cite>
<cite>journals/tocs/Herlihy86</cite>
<cite>journals/tods/KungR81</cite>
<cite>journals/jacm/Papadimitriou79b</cite>
<cite>...</cite>
<cite>...</cite>
<cite>conf/sigmod/StearnsR81</cite>
<cite>journals/tods/Thomas79</cite>
<cite>journals/tse/Weihl87</cite>
<cite>conf/podc/WuuB84</cite>
</article>
```

The outer tag is the record type: like in BibTeX the record types are “article” for a journal paper, “inproceedings” for a conference paper, “book”, “proceedings”, “incollection”, etc. The second level tags are the field types. Most fields have the same semantics as in BibTeX, new fields are:

`<url>` points to the position of the publication in a table of contents HTML page

`<ee>` points to an “electronic edition” (citation) page

`<cite>` is a citation link (DBLP-ID of the cited publication), or ‘...’ for references not (yet) covered by DBLP.

You may use the DBLP records as a test or example data set for your software as it was done in the (award winning!) SIGMOD 2000 paper by H. Liefke and D. Suciu or may create new services based on the information provided by DBLP.

3. DBLP now contains more than 100,000 citation links for database publications. The Web page `http://www.informatik.uni-trier.de/~ley/db/about/top.html` shows the most frequently cited DB publications.

The slides shown in Dallas are available on

`http://www.informatik.uni-trier.de/~ley/db/about/dallas.html`

Michael Ley
July, 2000

SIGMOD Digital Symposium Collection (DiSC) Editor's Message

This year we completed the second issue of the ACM SIGMOD Digital Symposium Collection, DiSC '00. We are especially proud to feature the video of the SIGMOD/PODS Plenary address by Dr. Jim Gray (Microsoft), the 1998 A. M. Turing Award Winner, entitled: "What Next? A Few Remaining Problems in Information Technology".

Last year, we mentioned that it was the aim of DiSC to become the "most comprehensive and timely source of information on database research." DiSC '00 has already obtained this objective, by including the full contents of the proceedings of SIGMOD '99, PODS '99, VLDB '99, ICDE '99, KDD '99, CoopIS '99, DOA '99, DASFAA '99, SSDBM '99, WebDB '99, and of the 1999 NSF Information and Data Management Workshop. The presentation slides and other material for various invited talks, tutorials, demonstrations, and panels of SIGMOD, PODS, VLDB, CoopIS, and the EDBT Summer School are also part of DiSC '00, as is the video of the SIGMOD/PODS Keynote Address by Dr. Theodore Berger (U. of Southern California).

Isabel Cruz
July, 2000

SIGMOD'2000 Program Chair's Message

Dear Fellow SIGMOD Members,

It is my pleasure to write a few brief remarks here summarizing comments I made at the business lunch in Dallas.

I would like to begin by expressing my sincere thanks to the people who made my job as PC chair possible, and who are responsible for the high quality program we enjoyed at SIGMOD.

The Program Committee Members had to review approximately 20 submissions each. Every one of the PC members completed every assigned review before the PC meeting (although some on the PC gave me some anxious moments!) I am truly grateful for the time, energy, and expertise each PC member put into this job.

Phil Bernstein and the members of his Industrial Program Committee did a great job soliciting and reviewing submissions for the Industrial Program. Similarly, Jiawei Han and his Demo Program Committee put together a first rate program of demos. Svein-Olaf Hvasshovd selected a great series of tutorials, while Alon Levy selected and assembled a great series of panels.

I would also like to thank Surajit Chaudhuri and the Microsoft Conference Management Toolkit staff for writing, maintaining, and running the software that managed the submission and review process.

Finally, I would like to thank each of the authors of the 248 submissions to the conference. As PC chair one truly begins to understand the number of active people in our community and the quantity and quality of work that they do. Obviously, the conference depends at its core on these authors.

Moving on, I will present a few statistics about the submissions to the SIGMOD 2000 research track. Overall, there were 248 submissions, of which 42 were accepted, for an acceptance ratio of 17%.

For a little history, we had:

1999:	205 submissions,	42 accepted
1998:	246 submissions,	42 accepted;
1997:	202 submissions,	42 accepted;
1996:	290 submissions,	47 accepted.

It will be interesting to see if this pattern holds; if so, Timos should receive approximately 200 submissions to SIGMOD 2001.

SIGMOD 2000 was truly an international conference. By region, our submissions break down was:

North America:	146
South America:	3
Europe:	57
Asia:	42

Over 40% of the submissions came from outside North America.

By submission count, the three most popular areas and their acceptance rates were:

Data Mining:	9/39	23%
XML/Web:	7/29	24%
Query Opt./Eval:	5/26	19%

Some areas turned out to be particularly tough:

OLAP:	0/10	0%
IR:	1/10	10%
Data Integration:	1/9	11%

Let me close with some changes I think our community should be considering with respect to the SIGMOD research program.

First, we reject too many papers! A fundamental question we should address is whether to increase the number of papers in the conference. Issues to consider are proceedings space, conference talk slots (although I am personally partial to having 100 parallel sessions in the first hour of the first day then letting everyone schmooze for the remainder of the conference), and a loss of prestige due to becoming less selective.

Second, the web has opened up the potential for new mechanisms in the reviewing process. Most interestingly, it would be possible to incorporate a round of author response in the loop. It could work something like this: 2 weeks before the PC meeting, all reviews are due. Authors could then read their reviews and prepare a one page response. Then this response could be included with the reviews during the PC meeting. The idea is that this can improve the reviewing process by giving the authors a chance to clarify misconceptions in their reviews. I have bounced this idea off a number of you already, and reactions have ranged from great enthusiasm to complete horror, with more reactions on the side of enthusiasm than horror.

Finally, I would like to thank those of the SIGMOD community who roped me into this job, especially Rick Snodgrass, who is the epitome of a great SIGMOD chair. It was an honor to serve as PC chair, and I look forward to seeing many of you also experience this honor in the future.

Jeff Naughton
July, 2000