

Editor's Notes

This issue is especially important to me as it marks the completion of my first year as the *Record's* Editor. So far so good, but I thought it would be a good time to think about the following question: "what is the *Record's* role within ACM SIGMOD?" After some thinking (and three versions of these Notes!), I thought that this can be translated into the following questions: should the *Record* become a "real" fully refereed journal? Should it be more of a venue for articles, maybe not as "researchy" (I know this is not a word, but English is a forgiving language), but of more interest to the SIGMOD community at large?

I think the latter should be it. Actually that is what I think it has mostly been, and the regular columns are a good example of that. Winslett's Distinguished DB Profiles and Ross' Influential Papers columns are examples of what one could not find elsewhere but on the *Record*, and I have received good feedback about those in general. Çetintemel's Research Centers column helps most of us to place our colleagues' work and environment within a wider framework besides their papers, and Cooper's Articles, Reports and Notes column allows one to have a better idea about other events, in a world where there are already too many events to follow. This being said, it does not mean that there is no room for more research-oriented articles. There is, but the key is on their coverage. Libkin's Database Principles column, for instance, is a good example. It always features a research-oriented article but typically of a broader coverage than a typical, say, *PODS* paper. I dare to say that those papers would not make into *PODS* or *ICDT* (due to their broader scope, not quality!) but they are perfectly fine for the *Record*. (Similar reasoning applies to all other columns, the one above are only used as examples to motivate my idea.) In fact, this very issue is an excellent example of articles that I consider far-reaching contributions, more on that later. On this front, I believe, and have heard from many others, that the *Record* is doing its job well.

Now, let me venture into a more delicate terrain, what I think the *Record* is not. It is not to be viewed as a full-fledged archival journal. Yes, it does publish original research articles, which are reviewed but not refereed, i.e., not as rigorous as in a typical journal process. And at this point a fair question that you may have is: what is the difference between reviewing and refereeing? For that I refer to ACM's "Policy on Pre-Publication Evaluation"¹ which I took the liberty to summarize below:

- *Refereed material is subjected to a detailed peer review ... Refereeing is generally directed to scholarly material for purposes of ascertaining originality, correctness, novelty, importance, and clarity of exposition ...*
- *Formally reviewed material is subjected to a structured evaluation and critique procedure following a defined process uniformly applied as with refereeing, only without requiring that the tests of scholarly originality, novelty and importance be applied in the previous sense ... Evaluation for technical accuracy is still required, and the criterion of clarity of exposition may be interpreted as readability by a certain intended audience.*
- *Reviewed material is subjected to a more informal and not necessarily uniform process of volunteer review, with standards dependent upon the publication and the type of material ...*

Furthermore, as per ACM's policy "Publications consisting only of reviewed or unreviewed material are considered 'informal'" (which bears no relationship to unimportant!) There is no positive or negative spin on this, this is ACM's policy, and the *Record* follows it.

¹ http://www.acm.org/pubs/prepub_eval.html (Thanks to Rick Snodgrass for pointing me out to this documentation).

This brings me to a side note, that may be important for those less experienced and/or perhaps more “aggressive”. Once I was asked by a Department Chair if the *Record* was a *typical* archival journal. Clearly, my view is that it is not, and that conflicted with that of a faculty member in his/her department, incidentally seeking promotion. Since I am the Editor, I believe my definition prevailed, and the faculty member’s credibility suffered with the exaggerated claim of having one too many journal papers. A word of advice, if I may, be careful with your claims with respect your publications on the *Record*!

So, what is the conclusion I am coming to? The *Record* should strive to be broader in its coverage, without losing sight of correctness and novelty. Articles of the survey type, or presenting state-of-the-art and research directions (or non-directions for that matter), position and visionary (not delusional!) papers –these should form the core of the *Record*. Special issues can also be accommodated, and in this case they should be more research-oriented towards a common topic of general interest. This being said, I think that perhaps I should avoid having in the *Record* papers that could very well fit in conferences or workshops. We all know that there ought to be a venue for every technical paper one can write. On the one hand, this is not to say that I will not accept research articles, at least not yet, but I will be stricter about the “coverage” aspect. On the other hand, that is to say that good papers may be rejected, because in my opinion (or of the reviewers) there is a better dissemination venue than the *Record* and/or its contribution is too narrow, despite being correct and original.

The above is mainly what I have in mind, and I would most certainly like to hear your opinions on the ideas above. The *Record* has to be what you want it to be and these ideas have to be fully debated among the *Record*’s associate editors and within SIGMOD’s Executive Committee as well before any action can be taken. Your views would be important input to that discussion, so feel free to email me at record@sigmod.acm.org (I would appreciate if you use “SIGMOD Record Feedback” or something similar in the subject header).

Let me now give you some statistics about the *Record* since I took over its editorship. All four issues of 2005 consumed 396 pages –not bad considering that I was told the *Record* has budget of 400 pages/year! As of the time of this writing, I have received 30 article submissions. Those do not include invited or solicited contributions, articles I considered of general interest to SIGMOD’s community or the special section in Sept./2005. Of those, 9 have been accepted for publication, 12 have been rejected and the remainder is under review. Given that the average submission has 6 pages, one can clearly see that the bulk of published material is made of the regular columns or those articles I would classify as of general interest, and which are not counted within the submitted articles above. While at it, let me thank all those who have helped review papers for the *Record*! They have helped a lot with their own time for the good of the community, and have also been quite forgiving with my requests and “reminders.”

How about this issue though? Besides the usual columns and some submitted research articles, it contains a few articles that illustrate well the spirit of what I mean above. The report by Amer-Yahia, which I placed under the Event Reports column, is about one of the panels at SIGMOD 2005. I welcome very much Amer-Yahia’s initiative and would like to encourage other panel organizers, not only of the SIGMOD conference, but of other events as well, to write such reports. Still on the spirit of broader coverage this issue also includes a paper by Gray and colleagues on scientific data management. This article has been published earlier on Cyber Technology Watch (February 2005), and is re-published here (with permission of the original publishers) because the authors felt that the *Record* provides the broad reachability and the right audience for this paper. The papers by Stonebraker, Çetintemel and Zdonik, as well as the one by Franklin, Halevy and Maier are also very much in line with what I mean by a broader article that is not as focused towards one contribution (as the vast majority of journal and conference papers) but is just as important as it can help shape one’s research program.

The other papers I want to mention are likely to spark some discussion among many of us (somehow the word “heated” comes to mind). The first is by Rahm and Thor where they present an analysis of the citations within our community in the last 10 years using some of the well-known venues we have. I was particularly proud to see that the top-three most cited articles in an archival publication² have appeared in the *Record*. It is also important to note that those papers are survey papers, which again sends me back to the view that the *Record* is indeed a good venue for such broader papers. Note that certainly there is still place for more authoritative venues, though some of those are typically much slower in their turnaround time (I recently learned from an author whose paper is still under the first round of reviews after over two years!) The second paper is by Sidiropoulos and Manolopoulos, where they also use citation analysis to rank publications. Although the authors contextualize their discussion about “prize awarding” I suggest you to read it beyond that narrow goal. I should also say that the opinions and facts stated in both papers are those of the authors and not necessarily endorsed by ACM SIGMOD. Nevertheless, in my opinion, they help illustrate very well my view of the *Record*’s role within our community: to provide news and articles of interest to the community as whole. A third related paper is the one by Bernstein and his colleagues reflecting the VLDB 2005 panel on “database publication practices”. It provides some interesting statistics about the paper submission process to some of our best known archival publications.

It is also time for my first “errata”: in June’s issue the paper by Hull and Su (“Tools for Composite Web Services: A Short Overview”) was missing some references that were mentioned in the text. That escaped both Libkin’s and my notice, and we apologize for the inconvenience. At the time you read this it should have been replaced online.

Finally, I have three thank-you notes, in no particular order: (1) to the ACM staff who handles the *Record*’s actual publication, in particular to Julie Goetz, (2) to Alex Labrinidis to helping me putting the *Record* available via SIGMOD Online as soon as it is ready, and (3) to Ken Ross; this issue also marks his last issue as the editor for the Influential Papers column (he is deservedly moving on to be associate editor on *ACM TODS*, and I hope to soon have someone taking over as associate editor of that column). Oh, there is a fourth thank-you. This may well be the longest “Editor Notes” the *Record* ever had, and if you managed to read these notes up to this point, you certainly deserve a big thank-you for your interest!

Mario Nascimento, Editor.
October 2005

² Rahm and Thor refer to “journal” publications, however since I want to distinguish the *Record* from a “typical” journal, I chose to use the term “archival publication” instead of journal.