

The Conference Reviewing Crisis and a Proposed Solution

H. V. Jagadish
University of Michigan

jag@umich.edu

ABSTRACT

In Computer Science, we have developed a vibrant conference culture, which has served us well thus far. However, with the growth of our field, the number of submissions to many conferences has sky-rocketed, leading to a downward spiral in reviewing quality and author satisfaction. This article proposes to break this downward spiral for the database community through JDMR, a journal for short “conference style” papers with rapid turn-around. An initial step toward this vision has been taken by VLDB.

1. THE CURRENT SITUATION

Our community has established very highly regarded conferences such as SIGMOD and VLDB. However, as our community has grown, these conferences are struggling to scale up. The number of submissions to these conferences has been on a steady increase over the years, more than doubling in the past decade. The program committee size has also grown in proportion.

The enormous size of our program committees leads to huge variances in reviewing. An individual PC member sees only a very small piece of the set of submissions. Their role becomes essentially that of a reviewer. Originally, the notion of a conference program committee used to be that the PC was actively involved in the selection of the entire program at the conference, and was at least aware of every part of the program. This was how SIGMOD and VLDB used to be in “the old days”. This continues to be the case in many other prestigious conferences today, such as SOSP and SIGCOMM. With our large PCs we have lost the normalization across accept decisions that a PC-based decision allows. Conference organizers are aware of this problem, and feel constrained to raise the PC size further.

As submissions increase in number, with pressures not to increase PC size, we have a tremendous number of reviews required of each PC member within a short period. Twenty reviews within two months is not an unusual load. While most PC members do the best they can, there is definitely a fatigue factor that limits the care with which reviewing is done. This leads to a variance in reviews, in addition to the variance in decision-making described above.

From an author’s perspective, this variance in review and in decision-making lead to unreasonable and disheartening rejections (and some unwarranted acceptances too, but authors aren’t the ones to complain about these). A rational thing to do under these circumstances is to resubmit, to the next conference, taking another try at spinning the roulette. This resubmission exacerbates the difficulties caused by too many submissions described above. In fact, the high resubmission rate may be a leading cause of the high submission rate, which in turn leads to the reviewing limitations described above, leading in turn to even more resubmission, and setting up a vicious cycle. See, for example, [2] and other papers at the Workshop on Organizing Workshops, Conferences and Symposia (WOWCS).

2. PRELIMINARY STEPS

Having recognized the problems above, there are two significant steps that our community has taken towards addressing some aspects of the above problems. One is roll-over between conferences, and the second is author feedback. Both of these are generally viewed as being useful, but also as being “painful” in that they require considerable additional effort on the part of conference organizers and program committees. There is no consensus in the community today regarding this cost-benefit trade-off. There are those who feel the benefits are worth the cost, and others who feel that the complications caused are just not worth the small benefit provided.

Roll-over between conferences (currently SIGMOD and VLDB) permits authors to submit papers rejected from one conference “with memory” to the other. The second conference assigns one fresh reviewer, in addition to the reviewers from the original conference. The re-

submission includes a list of changes the authors have made in response to the reviews. This mechanism has been designed to address papers that could have been accepted if only some specific issue had been addressed better. Since roll-over is managed as an exception to the normal conference review process, it comes with a high cost to conference organizers and program committee chairs.

An author feedback phase, comprising a mini-round of review/rebuttal, is increasingly being used by our conferences today. There usually are major constraints on what the authors are permitted to say, and very tight time limits on account of trying to make room for this within the reviewing cycle. Authors are not permitted to update their submissions during feedback. Unfortunately, this places authors in a position where they have to defend everything they did, even against legitimate criticism in light of which they would have modified their paper if they could. Most authors attempt to rebut every negative reviewer comment, whether right or wrong. Reviewers, in consequence, pay less attention to the feedback than one may expect. Anecdotally, it appears that reviews are changed only occasionally in response to author feedback.

In short, these two innovations are small “band-aids” that take small steps towards addressing the crisis. The JDMR proposal below is meant to provide all of the benefits of these two schemes, and more, while incurring none of the cost, since the mechanisms will be built in to the standard review process rather than being spliced in after the fact.

3. JOURNAL OF DATA MANAGEMENT RESEARCH

My proposal is to establish a Journal of Data Management Research (JDMR). JDMR negotiates with existing conferences to “participate” in JDMR. For participating conferences, JDMR manages the reviewing burden. Currently, VLDB is the only participating conference.

Authors will submit to JDMR and not directly to the participating conferences (such as VLDB). JDMR will review and accept papers for publication in JDMR, and presentation at the appropriate participating conference.

3.1 JDMR Reviewing

The fundamental goal is to provide journal-style reviewing with conference-like turnaround. Having served as an Associate Editor for many journals, I can tell you that finding the right reviewers is hard work. First, one has to decide who would make the best reviewers for a paper. Then one has to contact these reviewers and get them to agree to review. Often people are busy and decline, requiring others to be contacted. Others may be slow to respond, leaving the editor in limbo for a while. This whole process is time-consuming, adding

as much as a month to the review cycle time. Furthermore, the typical reviewer is not expecting the review request, and often is willing to do the review only if given sufficient time: requests for quick turnaround are often refused on account of too many commitments in the immediate future. Conferences avoid these problems by having a program committee comprising people who have committed to do work for the PC in advance. Review assignments are made on a best effort basis within the pool of reviewers available on the PC. What this means is that some papers may not get the most knowledgeable expert reviewers. This is not just a weakness for conferences but also a strength – since non-experts (in the narrow area of a paper) may serve as reviewers, good conference papers must be written in a manner that makes the key contributions accessible to any one with a general knowledge of data management. Furthermore, even in the journal review system, the quality of the review depends heavily on who the associate editor is able to recruit for the review.

Keeping these strengths and weaknesses in mind, JDMR will have a standing Review Board similar to the program committee of a major conference. I expect that service on the JDMR Review Board will get the same level of recognition as service on a conference program committee. The expectation for each review board member will be set at a maximum of 15 papers per year, again comparable to conference PC load. However this load will be spread out, and I am hoping we can limit to no more than 3 new papers in any month.

There will be a standing Editorial Board for JDMR, comprising several Associate Editors, with staggered two year appointments. Each Associate Editor will be responsible for approximately 45 submissions per year, a role roughly corresponding to that of an Area Vice Chair at ICDE. The responsible Associate Editor selects the reviewers for each paper, including reviewers outside the review board as needed on occasion. The goal will be to have the entire process of one round of review, including assignment, review, discussion, decision, and notification, completed within a period of one month.

Acceptance decisions will be made independently for each submission, based on an overall JDMR acceptance standard, which will be comparable to recent VLDB conferences. There will be no quotas, and no comparison between concurrently submitted papers with independent authorship.

3.2 Crisis Revisited

I claimed that JDMR provides all the benefits of current initiatives to address the conference reviewing crisis. JDMR provides full journal-style reviewing, with multiple rounds. So the benefits of roll-over are already built in to the base process. The paper is available for presentation at the next participating conference.

With the JDMR journal review process, an actual revision can be sought, obviating the need for a separate feedback process. Authors receive first round reviews with suggestions for change as well as questions from the reviewers. Authors may prepare as substantial a revision as they wish and submit for a second round of review. Any review rebuttals from the authors will get significant consideration, since these will typically be limited to cases of real disagreement (or reviewer misunderstanding).

Turning now to the major concern, with sequential re-submission, the expectation is that authors are more likely to be satisfied with the decision made on their paper, and therefore less likely to resubmit. I am not aware of any one who has compiled statistics regarding resubmissions. (Besides confidentiality issues across conferences, we have the further complication that some papers are improved, even substantially, between submissions, so we would have to be careful in specifying exactly what we count). Nonetheless, my estimate, based on anecdotal evidence, is that more than half the submissions to any conference are revised versions of submissions rejected from other conferences. Let us say that each paper is submitted 3 times before being accepted or abandoned. We will say, equivalently, that the resubmission rate is 3. Suppose that the new reviewing mechanism drops the resubmission rate to 1.5. (The goal, obviously, is to get to a resubmission rate of 1. But that limit is unachievable). Let us look at the consequences. Today, of 600 submissions to a typical leading conference, about 200 are original and 400 are resubmissions. With a resubmission rate of 1.5, we will have only 100 resubmissions for a total of 300 submissions to the conference. If 90 papers are accepted, this gives an acceptance rate of 15% today, and an acceptance rate of 30% in the new system. This “magical” increase in acceptance rates is possible without accepting more papers, and in fact by accepting pretty much exactly what is accepted today. In other words, without sacrificing quality.

3.3 Policies And Implementation Details

We will need to establish many policies for the journal. A few salient issues are listed below.

Extremely prestigious journals, such as *Science* and *Nature*, are able to maintain very short turn around times. They also have papers that are even shorter than our conference papers. Both shorter papers and shorter review times can be achieved without a loss in quality. (In fact, they may even be positively correlated). It is a question of setting the right expectations and transforming the culture of our community. This issue will be addressed aggressively.

JDMR will have twelve monthly deadlines each year. Papers will be distributed for reviewing in a “mini-batch” once a month. Decisions will be made, also in mini-batch mode, once a month.

A systemic issue for our conferences is the repeated submission of substantially the same work. With the quick turn around JDMR strives for, this may become even more of a problem. For these reasons, JDMR will have a strict policy of no resubmission of rejected papers to JDMR for a period of one year from the date of submission. A paper is considered a resubmission if the majority of the material in it was included in a previous submission, whether accepted or rejected. (In other words, there will be no possibility of resubmitting a rejected paper as “new” as encouraged by some journals when a round or two of revision is insufficient to bring a paper to acceptance). Authors are free to submit manuscripts rejected from JDMR, with or without improvement, to other venues that they deem suitable.

For most papers, two rounds of review should be plenty. Since JDMR will not have a “revise and resubmit as new” option, and since we hope to have quick review turnaround, we will be open to additional rounds of reviewing as needed.

Often, there is discussion of conference-quality versus journal-quality work. I personally believe the difference is qualitative, but not necessarily in quality. JDMR is meant to present precisely the kind of work that is currently published at top database conferences such as VLDB, with the same tradeoffs between freshness of idea and completeness of exposition. However JDMR is a journal in that it has a multi-round review process with year-round submissions. The expectation is that the acceptance threshold for JDMR will be similar to that for the VLDB conference today, modulo the possibility of having one round (or more) of revision. This round of revision may be a “major revision,” applicable in the case of papers selected for roll-over or for acceptance with shepherding at present, or “minor revision,” applicable for most papers accepted to the conference today, with small improvements suggested by the referees.

The expectation with JDMR is to manage submission numbers typical of conferences (several hundred each year) and to remain as selective as the prestigious conferences currently are. The review process for JDMR must reflect these realities. A major strength of the conference review process is the consensus-building among reviewers through a discussion phase. In contrast, associate editors typically make executive decisions, taking the reviews into account, in the case of traditional journals. JDMR will use a reviewer discussion phase to permit the editor to make a more informed decision, with reviewer consensus where possible.

Often, the most useful description of a paper is not in the abstract, and not appropriate for authors to say for themselves. For example, it is usually important to place a piece of work in the context of other work in the area – an expert may be able to do this in a few sentences, which would likely be in a tone that is not

quite right for a serious paper abstract. Because of this, JDMR will designate a paper “champion” reviewer to write a paragraph that can be posted as a public review when the article is published in JDMR. This paragraph can usually be molded from the summaries of the reviews, and hence should not take a great deal of time to write. The champion reviewer will be encouraged to sign the public review, and reveal identity. However, this will not be required: if the champion strongly desires anonymity, the public review can be published as a statement from an anonymous member of the JDMR review board.

4. CURRENT STATUS

What has been described thus far is one person’s vision of an ideal future. This vision has been under public discussion for almost a year now, and some version of the above has been available on the web for any one to see. Many worthwhile comments have been made, and have been incorporated. A first step toward realizing this vision has been taken recently, by VLDB, and is described in this section.

The VLDB Endowment has created a new online journal, *PVLDB (Proceedings of the VLDB Endowment)* to include all material currently published as the Proceedings of the VLDB Conference. All papers presented at the VLDB 2008 Conference are included in Vol 1 of this new journal. Volume 2 will correspond with the 2009 Conference.

For 2009 and 2010, there will be a “journal track” for paper submission and reviewing in parallel with the traditional (Core DB and IIS) tracks. Authors may choose to submit papers to the track of their choice. Irrespective of the track chosen, all accepted papers will be treated equally, both in terms of conference presentation and in terms of publication in the next issue of *PVLDB*.

The plan is to have approximately one issue of *PVLDB* published every quarter, with papers that have been accepted in that quarter. The September issue will be thicker in that it will include papers accepted through both the journal track and the traditional conference PC. All papers accepted to *PVLDB* in a year will get presentation slots at the annual VLDB conference. For 2009, the acceptance cut-off date is May 29. For 2010, the date is not yet set, but is expected to be around May again.

The steering committee (appointed by the VLDB Endowment) will be responsible for determining the policies of *PVLDB*, negotiating with all affected parties, and establishing the mechanisms for review. The current membership of the steering committee is: Serge Abiteboul (VLDB PC Chair 09), Peter Apers, Phil Bernstein (EIC, *VLDB Journal*), Elisa Bertino (VLDB PC Chair 10), Peter Buneman (VLDB PC Chair 08), H. V. Jagadish (EIC, *PVLDB*), Martin Kersten, and Meral

Ozsoyuglu (EIC, *TODS*).

The VLDB Endowment has ultimate responsibility for *PVLDB*, including the appointment of steering committee and Editor-in-Chief.

An initial review board with almost 100 accomplished database researchers has been appointed [6], and the VLDB “journal track” is now open for submissions [7]. As a prospective author, you should choose to submit to this track for all the flexibility it offers, in return for accepting some uncertainty with respect to the new and as yet untried process.

5. SOME POSSIBLE WORRIES

I have discussed the JDMR vision intensively, with many people, and have received much valuable feedback that has strengthened the proposal. In addition to a great deal of enthusiasm from many, I have also heard some concerns. I list below the major possible worries, and my take on them.

JDMR is not really a journal: To the extent that JDMR is closely associated with conference presentation, and is interested in “conference-style” papers, it is a journal-conference hybrid. I personally believe it is more a journal than a conference proceedings, on account of year-round submissions and multi-round reviews. But there are those in the community who believe strongly that such a hybrid should not be called a journal. This is a dialog in progress.

Conference proceedings should not just be renamed a journal: The Computer Science community has many first rate conferences, and considers publication in them to be extremely prestigious. This is not the case in most other disciplines. In the US, there is an impactful report from CRA[1] on this question, and promotion committees at most Universities, at least in CS, recognize the importance of conference publications. Many, who have fought hard for this recognition, are wary of schemes that may dilute their arguments regarding the first-class nature of our conferences. On the other hand, conference publications are still not given their full due in many universities outside the US, and even in the US in departments other than CS, such as IS departments. Also, Thomson’s ISI index[5], used widely for citation and impact analysis, does not include our most prestigious conferences. Finally, it is not clear what the impact of conference versus journal publication is when inter-disciplinary evaluations are made, such as for awards. So there is still much ground to be won.

Whatever be one’s position on the difficult (and potentially contentious) issues regarding how best to get conference publications due recognition, there should be little disagreement regarding JDMR. It is not merely a conference proceedings masquerading as a journal: it has legitimate claim to being called a journal on account of year-round submissions and multi-round reviews.

Existing highly prestigious journals may suffer: JDMR is addressing conference-style papers, so to the first order it should not compete with existing journals. However, the number of papers (both reviewed and accepted) in top conferences is several times greater than in top journals. This makes JDMR a big player, so even partial competition may be impactful on current journals. Certainly the Editors of current leading journals are worried. I think this is appropriate – they have to watch out for unexpected impacts and make sure to retain the prestige and importance of their journals. All I can say is that JDMR is not consciously trying to compete with them. The EDitors in Chief of two leading journals are on the PVLDB steering committee to help minimize any possible competition and to maximize synergy.

We may lose the quick (within 3 month) decision that conferences give us today: In fact, for the majority of submissions, JDMR will provide quicker (1 month) feedback. Revisions will not be sought as a matter of course – we will strive to make clean decisions to reject or accept (subject to minor revisions) in the first round to the extent possible. Where revisions are sought, the authors will be given a clear road map of what they need to do to get the paper ready for acceptance. Authors should know that they have a high probability of having their paper accepted if they make all the revisions suggested.

The review period allowed is too short to permit a thorough review: The time available per review is substantially greater than that for typical conference program committees. Also, no one I know actually takes weeks actively reading and reviewing a paper. Rather, most of the review time is spent with the paper “on stack” waiting until the reviewer is able to get to it. As such, there should be no negative impact on quality of review at all. There is a change of expectation with regard to reviewer scheduling – my hope is that enough reviewers will agree to do this because this is how they would like to be treated as authors. To the extent we are successful in changing community expectations, we all win.

Submissions may not arrive staggered round the year: Since it is hard to predict how much revision work may be required before acceptance, authors really cannot plan on a specific deadline. This in itself should cause some spreading out of submission dates.

My goal is to have multiple conferences participate in JDMR in the future. Acceptance for publication in JDMR would then be delinked from presentation at a conference. Once JDMR is recognized as a prestigious publication venue in itself, this separation of presentation from publication will not cause problems.

At the time their paper is accepted to JDMR, authors will be asked which participating conferences they would like to have their paper considered for. Each participat-

ing conference will make its choice on its schedule. Authors will be obligated to present their paper at any conference that accepts their paper for presentation, from amongst those in which they have expressed an interest.

How conferences choose papers for presentation will be up to the individual conferences. I envision a small conference program committee (say 20 people) will look at the pool of candidate papers (all accepted already to JDMR) along with the reviews and discussions surrounding the acceptance of these papers and any public commentary on these papers since their publication in JDMR. Additional reviewing by conferences is not expected. A small program committee thus becomes practical, and these individuals are likely to feel a much greater commitment to the success of the conference program than members of a PC with more than a 100 members. This will probably result in a more interesting conference program.

Conferences will be free to establish their own rules regarding prior presentation (can an article accepted to JDMR be invited for presentation at both SIGMOD and VLDB), regarding age limits (can an article more than a year old be invited for presentation, for example because the previous year’s conference committee did not realize the impact of the work in the article), and so on. Conferences may also choose to have multiple tiers of presentation, such as short and regular.

If 2 or 3 conferences participate in JDMR, then we will get sufficient smoothing of review load throughout the year that this concern goes away. However, for the present, VLDB is the only participating conference.

If authors miss being accepted in time for one year’s VLDB, they have to wait a whole year: This is indeed a limitation of the current situation with only the VLDB conference participating. This difficulty will be ameliorated if other conferences join JDMR in the future. Note though that review periods are tightly controlled and guaranteed, so the typical reason authors may fail to be accepted in time is because they were unable to complete a satisfactory revision in a timely way. Based on first round reviews, authors will often be able to determine how much time a satisfactory revision will take, and decide upon an appropriate course of action accordingly. In the worst case, authors may choose to withdraw and submit elsewhere – all they would have lost is about one month in time to obtain three high quality reviews.

JDMR may create a barrier to publication at multiple venues: Today, authors have multiple shots at publishing the same material, in different conferences, with quite likely different reviewers in each. They are thus able to deal with the randomness in the review and acceptance process by “purchasing multiple tickets to the lottery” as it were. If multiple conferences participate

in JDMR, then authors are deprived of some chances at the lottery. However, to the extent that the number of acceptances to JDMR is the sum of acceptance of these conferences, the odds of winning the lottery are now higher for the author, and improvements in the reviewing process should decrease the randomness. Furthermore, there will always remain other publication venues not participating in JDMR. This is not even an issue at present, since the VLDB Conference is the only participating venue.

Having a smaller PC choose papers for conference presentation raises the possibility of cliquish behavior, or at least the appearance thereof: First, this is not an issue with PVLDB, since all papers accepted for publication in PVLDB are also accepted for presentation at a VLDB conference, and vice versa. Second, considering the JDMR vision, this statement assumes that the prestigious event is conference presentation, with JDMR publication acting as a first stage filter. This cannot be correct, since there will be no official conference proceedings, and the only citable publication will be in JDMR. Finally, a smaller committee does not mean 3 people: it can be made as large as required to get broad representation. The only requirement is that it be small enough to work as a committee that will be able to participate in meaningful dialog and generate more interesting and better balanced conference programs.

The real problem with conferences is the poor quality of some reviews and JDMR does nothing to address this: It is true that JDMR does not directly address review quality, other than through permitting author rebuttal. However, spreading reviews out over time permits reviewers to review better, since they are not being hit with a large burst of reviews all at once.

A deeper issue with review quality has to do with reviewer evaluation and reuse. Conferences do not usually have a formal system of reviewer assessment in place leading to program committee selection based primarily on academic reputation (and other issues such as “balance”). In contrast, many journals do attempt to keep track of reviewer performance (both timeliness and review quality). JDMR will be able to track reviewer performance like journals, and feed this into the choice of review board membership in the long run.

6. CONCLUSION

This article developed a proposal that

1. Improves the conference reviewing process by:
 - (a) Providing a more “journal-like” experience to authors, with an ability to rebut reviewers and make improvements to the submission, thereby providing a lower perceived “random factor”.

- (b) Reducing the total number of papers in the system by making it both harder and less desirable to resubmit.
- (c) Improving refereeing quality by spreading out referee requests rather than bundling them together as happens with conferences today.
- (d) Simplifying the system by eliminating the need for complicated protocols in place today for roll-over, for author feedback, etc.

2. Does all of the above while maintaining, and even enhancing, the prestige and quality of our very well-regarded major conferences.

3. As a corollary of the above, provides a prestigious venue for publication of “conference-style” papers that can legitimately be called a new breed of journal.

An initial step toward this JDMR vision is being implemented as *PVLDB*. See [4] for details.

Acknowledgments

This effort has been many years in the brewing. Peter Apers led this effort at the VLDB Endowment before me. Martin Kersten wrote a thought provoking piece on a “Paper Pond.” The PVLDB Steering Committee, of course, has greatly improved upon an original vision and developed an implementation plan. I am also grateful to the VLDB Endowment Board for its unwavering support.

7. REFERENCES

- [1] CRA Board of Directors. Best Practices Memo: Evaluating Computer Scientists and Engineers For Promotion and Tenure. *Computing Research News*, Sep. 1999.
- [2] J. Crowcroft, S. Keshav, and N. McKeown. Scaling internet research publication processes to internet scale. In *Proceedings of the USENIX Conference on Organizing Workshops, Conferences, and Symposia For Computer Systems* San Francisco, California, April 15, 2008.
- [3] Konstantina Papagiannaki. Author feedback experiment at PAM 2007. *ACM SIGCOMM Computer Communication Review*, v.37 n.3, July 2007.
- [4] *Proceedings of VLDB Endowment*. <http://www.vldb.org/PVLDB>
- [5] Thomson Reuters. Science Citation Index. http://www.thomsonreuters.com/products_services/scientific/Science_Citation_Index
- [6] VLDB (Journal Track) Review Board. <http://www.eecs.umich.edu/db/pvldb/reviewboard.txt>
- [7] VLDB (Journal Track) Submission Site. <http://cmt.research.microsoft.com/PVLDB>