

We are Drowning in a Sea of Least Publishable Units (LPUs)

David J. DeWitt
Jim Gray Systems Lab
Microsoft Corporation
Madison, WI
dewitt@microsoft.com

Ihab F. Ilyas
QCRI
Qatar
ikaldas@qf.org.qa

Jeffrey Naughton
Computer Sciences
Department
University of Wisconsin
Madison, WI
naughton@cs.wisc.edu

Michael Stonebraker
CSAIL
MIT
Cambridge, MA
stonebraker@csail.mit.edu

ABSTRACT

Our field is drowning in a sea of conference submissions. We assert that the sheer number of papers has begun to seriously hurt the quality of the work that the field is doing and that the field is going to implode unless we take action to remedy the situation. In order to improve the quality of the papers being published we must reduce the number being submitted. This will require a change in the culture of our field where “more” is being equated to “better” by both hiring and promotion committees. In this panel we will explore some ideas for correcting the situation.

Categories and Subject Descriptors

K.7.4.m [The Computing Profession Miscellaneous]: codes of good practice

General Terms

Human Factors

Keywords

Conference submissions, reviewing practices, tenure expectations,

1. INTRODUCTION

For at least a decade, the major DBMS conferences have been overwhelmed by paper submissions. This has led to:

* **Very large program committees** (order 200 people) to spread the reviewing load around.

* **Spotty reviewing.** After all with a sea of papers, it is hard to assign them carefully to area experts

* **Rejection of speculative papers.** Program committees have dealt with the sea of submissions by rejecting papers that are poorly written, have speculative conclusions or are otherwise lacking in rigor

Our community has reacted to this state of affairs by:

* **Submitting only “safe” papers**, i.e. ones that are technically deep in a very narrow vertical domain with lots of mathematics and theorems.

As a result, our conferences are full of boring, narrow papers which have little or no long term impact. Moreover, researchers with speculative papers will go elsewhere to see their work published. This will clearly lead to a “death spiral”.

Some of us have looked at the resumes of many fresh PhDs recently. It is clear that the competitive ones must have at least a half dozen publications. In contrast, a quarter of a century ago, competitive PhDs had one or two. Moreover, tenure cases are made based on 20 or more publications. A quarter of a century ago it was 6 or 7. Hence, publishing requirements have escalated by half an order of magnitude. We are hard pressed to believe that graduate students and assistant professors are 5 times more productive than 25 years ago. Clearly research and academic administrators think quantity not impact. The obvious conclusion is that this escalation of publishing requirements encourages our field to divide their work into smaller LPUs.

In addition, our field is at least three times larger than 25 years ago. These two effects drive up the quantity of papers by a factor of 20 or so and lead to the issues mentioned above. We are led to ask “What should we do?” The rest of this short paper and panel discussion summary lays out a collection of ideas. We group them into ones that are safe and incremental, and ones that are more revolutionary.

2. INCREMENTAL IDEAS

Attend any research session at a major DBMS conference and you will largely note poor attendance and boring, narrow, deep presentations consuming 30 minutes. An obvious answer is to move to “lightning talks” of say 10 – 15 minutes. After all, the purpose of the talk is to present a summary of the paper and encourage people to read the paper. This can easily be accomplished in 10 minutes. Moving in this direction would allow conferences to accept 2-3X more papers. This would increase the supply of “slots” dramatically.

Another idea is to drop research presentations completely in favor of research posters. That way the 10 people who are interested in narrow topics can talk with the author about the work. Again, this would allow the number of slots to be dramatically increased, without sacrificing interactions.

A third idea is to introduce another class of papers. After all, we now have industrial papers, demo papers, and research papers. Why not have a 4th class? The idea would be that 30 minute speaking slots would only go to papers of broad interest, and the remainder would use a lightning format or a poster format. Why stop there? A 5th class of papers could be for controversial or speculative ideas.

Of course, moving in this direction will force administrators to work harder. They can no longer just count publications; instead they will have to figure out the actual impact of any given job applicant. In our opinion, this is a good thing to do.

Another idea is to require everybody who wishes to attend a major conference to write a position paper or abstract. This has been used effectively by HPTS, for example. Based on the submitted materials, the program committee would select an interesting collection for presentation. In effect, this would mimic the style of HPTS.

Yet another idea is to construct a specific forum for assistant professors to impress the “gray beards”. After all, this is one of the purposes of the major conferences. One idea would be to have ten or more “gray beards” sit in a meeting room for (say) 90 minutes and listen to 18 five minute pitches. That way young researchers can influence a large collection of decision makers. A similar forum could be constructed for advanced graduate students. An extension of this idea is to have specific sessions for recruiting, whereby people looking for a job get (say) 5 minutes to say “look at me”.

Finally, “birds of a feather” sessions have completely disappeared in our conferences. These could easily be reinstated as discussion sessions on a particular topic, organized and led by an interested researcher.

All of these ideas would have the impact of dramatically increasing the publication slots at major conferences. It would also dramatically lower the number of papers that are “recycled” to the next conference. Since many papers are submitted 3 or 4 times, this would lower the “churn”. All of these ideas do not “upset the applecart”, in that none of these require dramatic changes in the major conferences. It is unclear, however, if these incremental changes will boost the quality or the impact of our papers. We now turn to a more revolutionary idea.

3. A MORE RADICAL IDEA

Our current major conferences are zoo-like (many hundreds of people), at which it is difficult to find people you want to interact with. They are wide-scoped, which means most papers are of interest to only a few people. Hence, most conferences have several parallel tracks. Our feeling is that this “mega conference” format should be discarded, and that it should be modified as follows.

First, conferences should limit attendance to (say) 200 people. An obvious way to do this is to require everybody who wishes to attend to write a one-page idea. A program committee could decide attendance based on these submissions.

Second, conferences should be single-track and probably way more narrowly focused. One can, of course, co-locate multiple of these conferences at the same time. It is silly to present a paper to 15 people (3% of the attendees) while the other 97% are elsewhere.

In essence, let’s have a dozen (or so) more narrowly focused conferences, instead of three “zoos”.

This would free conferences for their requirement to “keep score” for administrators, and they could focus on more interesting aspects around getting a bunch of similar researchers together.

4. CULTURE IMPACT

The solutions discussed by this panel or by its audience may cause or may require culture changes, which can be classified as follows:

Community Culture: in the way graduate students are trained in terms of deliverable, and expectations. It is hard to steer a big ship and the transition has to be well managed for any of the solutions to be sustainable.

Administrators Culture: in how to educate and to deploy these changes with hiring and promotion committees. We worked hard for years so the names of these venues are known and are currently used as an “easy” measure for impact.

5. ACKNOWLEDGMENTS

Our thanks to the SIGMOD 2013 Organizing Committee for accepting this panel proposal.