Conferences vs. Journals in CS, what to do?
Evolutionary ways forward and the ICLP/TPLP Model

Position paper for Dagstuhl meeting 12452:
Publication Culture in Computing Research

Manuel V. Hermenegildo
IMDEA Software Institute and T.U. Madrid

Draft of Nov. 7, 2012

Abstract

We computer scientists seem to do it differently to other sciences: we publish mostly in conferences and our conferences are of a different nature. Our journal papers are long and take a long time to review and publish whereas often their papers are short and published quickly. And all this interacts with the tendency to evaluate researchers or departments frequently and in a mechanical way (via paper numbers and citation counts) instead of infrequently and deeply (by actually reading papers) and the fact that the current way in which bibliometry is done makes our papers invisible to the world. This position paper offers my viewpoint on what the problems are and why they are important, and also elaborates on some realistic ways forward. In particular, regarding the issue of conferences vs. journals, it proposes the model adopted a few years back by the logic programming community for its main conference (ICLP) and journal (TPLP, Cambridge U. Press). This model is based on the assumption that CS journal papers can be of two types: rapid publication papers (similar to those of other sciences and also close to our conference papers) as well as the longer journal papers that are traditional in CS. Then, the concrete proposal is to, instead of publishing the traditional conference proceedings, have the papers submitted instead to a special issue of an (indexed) journal which is ready online in time for the conference. The traditional conference reviewing process is also improved (following journal standards for rapid publication papers and special issues) to include at least two full rounds of refereeing and a final copy editing step. I argue that this model offers an evolutionary path that solves a good number of the incompatibility problems with other sciences of the current CS models, without giving up the essence of CS conferences. For this reason I believe that this model (or one of the closely related models being proposed) would be a clear path forward, and relatively straightforward for the community to adopt widely.

1 Problems with the CS publishing culture

It is a well known fact that our publication culture is very different from that of other sciences: we publish mostly in conferences while journals are the norm almost everywhere else. While in other sciences conferences are generally just communication vehicles for presenting work which is typically submitted to journals, CS conferences publish proceedings of papers that are fully

refereed publications, and are often perceived to be of better quality and more prestigious than most journals. Indeed many top-level CS researchers seldom publish in journals. Furthermore, the papers published in CS journals are also generally very different from those of most other sciences: our journal papers are very long and take a very long time (often years) to review and publish, whereas in most other sciences papers are generally much shorter and are reviewed and published in a much shorter time. A reason why the publication delay of CS journals is tolerated to some extent may be because, after all, the results are really already published in conference proceedings and the role of the journal paper is to complete the work offering full proofs, comprehensive experimental results, etc., in some ways similarly to a monograph on the topic.

Is there a problem? It is tempting to see no problems in the singularity that the CS publication culture represents. It is true that we have great conferences that in many cases we consider better than any journal, and thus publishing in their proceedings should not be a problem, and that, at least at certain levels, the model has worked well for us so far. Furthermore, it is perhaps even attractive for an innovative science as ours to be different. However, the recent thread of viewpoints, opinion papers and workshops, on the CS publication culture shows that not all is well [4, 6], and it has been argued that this culture gives rise to a number of distortions (see, e.g., [5, 8]). Without entering the argument of whether the conference system works for us or not, I will go in a slightly different direction and argue that perhaps some of the most serious problems with the CS model in fact have to do with the way it makes CS be mis-perceived *outside* CS.

A first problem: the perceived value of conference of proceedings papers vs. journal papers beyond our discipline. While the value of conference papers is often recognized at, say, the CS department level (although there is a worrying trend of cases also where it is not, e.g., in many universities in Europe) problems often start at the levels just above. It is a fact that a top-level CS department can easily push upward a tenure case with a good number of papers in first-level conferences and few or perhaps even no journal papers. And the case will go through the upper levels with perhaps some explanation by the department head, even if the dean and other department heads will probably think that it is really up to the CS department to waste one of their positions on a researcher with such an obviously low number of "real" papers (or even with no "real" papers) published. However, when something is at stake that crosses disciplines, for example a university-level award or a distinguished professor position, the low number of journal papers of a CS candidate may easily become an issue: colleagues from other disciplines may be less understanding with the atypical CS culture at this point, and be ready to see it as an excuse to push a candidate that is obviously worse than theirs, which has over one hundred journal papers. This also helps perpetuate the misconception that we are not a "real" science among the other sciences (how can it be a real science if there are no "real" publications?).

A second (and arguably even more damaging problem): the widespread and to some extent inevitable use of raw bibliometric evaluation. One may fairly ask why we should even worry about bibliometric evaluation when we all know that it has many serious flaws [1, 2]. I certainly believe that it is imperative that we reverse the current growing tendency to evaluate researchers *too frequently and in a mechanical way* (via paper numbers

and citation counts) instead of *infrequently and deeply* (i.e., by the revolutionary concept of actually reading and understanding papers, and judging their significance and impact on others in technical terms).

However, the unavoidable fact is that bibliometric evaluation is used more and more everywhere. There are probably many reasons for this, including that the alternative (really looking at people's work) is of course more costly and, more importantly, probably becomes unfeasible as larger and larger sets of individuals are compared. Thus, whether we like it or not bibliometric comparisons are the norm when evaluating and ranking departments, schools, universities, research centers, countries, etc., etc. Their use in one way or another is widespread at all levels of science policy (i.e., in the decision-making processes of those that make strategic decisions about our organization, funding, etc.). Even at the level of individuals it is sometimes very hard to avoid using bibliometry, for example at the first levels of filtering when there is a large number of candidates applying for a small number of fellowships.

Because bibliometry is done using the standard journal databases (i.e., Thompson ISI/JCR) our papers (and thus, CS as a whole) are invisible. Unfortunately for CS, all bibliometric comparisons are invariably made using the de-facto standard, i.e., the Thompson ISI databases and tools, and *in practice these databases only index journals* (and not all). Some of Thompson's databases (and the competing Scopus) do list some conferences but there are many CS conferences missing [3], which makes them quite inadequate in practice for CS bibliometry. In any case bibliometric comparisons are typically made using instead the "standard" databases, i.e., the ones listing only journals, since that is what is appropriate for other sciences. Given our conference-centered publication culture the very unfortunate side effect is that the immense majority of our papers are *invisible* to the indexing and comparison tools used by everyone outside CS. It is well documented that ranking CS researchers using these databases and tools gives clearly invalid results [2]. Then, surely comparing CS researchers, institutions, etc. to those of other disciplines using these tools is equally invalid —but nevertheless *this invalid picture is what they see of us.*

All this puts CS in a very unfavorable situation. For example, since SCI publication data is a vital factor in university rankings, CS departments can be perceived as not contributing much to the overall position of the university for the simple reason that they only in bring a small number of indexed journal papers (and perhaps it is not worth investing in CS as much as in other disciplines). A simple, concrete example of this effect, just to make the point: Mattern [2] has documented how at ETH Zurich a bar chart listing publications is used to compare departments. Of course this chart is generated using ISI tools, thus only reflecting SCI journal publications. The CS department inevitably displays a tiny-looking bar which does not reflect in any way its publication performance. A similar chart is published in the monthly newsletter at our university, and circulated to all faculty and to external stakeholders, which leaves the CS department in a similarly unfair position with respect to other departments. The fact is, CS publications, and thus CS as a whole, are mostly invisible to bibliometry.

We have won battles, but they are recurring and uphill, and we probably cannot win this war. It may be tempting to think that this is all just a local problem (e.g., European, or of lesser departments, or whatever) and that in one's department/university/funding agency/country the conferences vs. journals battle has been won. But I believe that thinking

in this way is dangerously complacent. First, I have seen many cases where after apparently winning the argument and establishing special criteria for CS, some time later the person holding the dean / university president / funding agency director / policy maker position changes and the first thing the new holder does when reviewing a policy that contains a strange special case which he/she does not understand (and actually gives the impression that it just a weak department or area somehow trying to cheat) is to promptly reset the situation back to "normal:" hey, privileges should not be tolerated! I have fought myself a number of such battles at different levels and actually won (temporarily, I am sure) many, but in the process I have also gradually come to the conclusion that, while we can win a battle here and there, we simply cannot win this war. The bottom line is that we are defending a position that is not shared by the vast majority of scientists across all disciplines and this will always be at the very least an uphill battle and one which, I believe, we cannot win in a stable, long-term manner.

We cannot always be there and in any case we do not (cannot?) have a good bibliometric alternative to offer. A vital factor in coming to the conclusion that we are fighting a losing war is that in many of the situations in which CS will fare badly due to the use of journal-based bibliometric data it is simply not possible for one of us to be there present to explain that we are special, that we publish in conferences instead of journals, that our journal papers are much longer and thorough, that they stay longer in the pipeline, etc., etc. Also, it takes time and energy to explain things at such meetings and to write white papers about the CS special case –it has surely taken a lot of my time, over the years, including writing this article– that we can put to better use. If you think this is not the case because you have never had to do it, think again: there is surely someone above you (at your university, at the national CS department's meeting, funding agency, scientific society, etc.) that has put and continues putting the required time and energy for you in order to keep pushing the stone back up the slope over and over again.

And, even worse, even in the unrealistic case that we could be present every time to explain our case and we had enough people willing to do it, it is still not clear that one can offer an alternative *bibliometric* way of making comparisons across disciplines that treats both CS and other disciplines fairly, and, more importantly, that every user of bibliometric tools at all levels will accept and switch to. It is simply not straightforward to use CiteSeer, CORE, DBLP, Microsoft Academic Search, Google Scholar, etc. data alongside the SCI numbers to provide equivalent data for CS and other sciences when elaborating rankings, charts, etc. Conversely, it is in general not clear that data from these more CS-friendly tools can be used consistently for all the other disciplines. Finally, even if these tools were to evolve and eventually cover all disciplines, at least in the medium term it appears highly unlikely that other disciplines will accept to be evaluated and compared using anything other that the most established data and tools (currently SCI). So, while we should certainly support the very worthwhile efforts towards getting the newer indexing mechanisms to be widely accepted, it does not seem wise to rely solely on the unlikely universal adoption of these new tools. Furthermore, the issue of conference vs. journal papers would still remain.

2 CS should switch as soon as possible to publishing in journals, but in combination with conferences

For all the previous reasons, while I have been myself for many years and in many instances a stark defender of the CS conference culture and the value of our conference papers, for quite a while now I have been convinced that we in CS need to switch as soon as possible to models that are more journal based. However, by this I do not mean to simply drop our conference papers and publish all work through *traditional, CS-style* journal papers. This is clearly not a solution, because these journal papers are quite different and serve different purposes than the conference papers: traditional CS journal papers are much more complete and much longer (50+ pages sometimes) than a conference paper and they take a very long time to review and publish (often in the order of years). While publications of this type, complete with the associated reviewing process, are necessary and useful in many cases, they are at odds with rapid communication of scientific results, which is what the conference papers do.

What do other sciences (e.g., biology or physics) do with the kind of papers that are similar to our conference papers? They publish them in journals: in fact, their journal papers are very similar to our conference papers in both content and publication time. They tend to be relatively short, published relatively fast (e.g., around 2-3 months from submission), and with a reviewing process that is not more strict or thorough than that of our good conferences –except for being able to implement in the same amount of time more than one round of refereeing (to which I will return later). The very long journal paper seems to be a CS phenomenon (and to some extent also in mathematics), and is probably another artifact of the existence of conferences that are refereed and considered real publications. In CS, conference papers have taken the space occupied by the shorter, faster journal publications of other disciplines. And the CS community probably tolerates the long reviewing periods of many CS journals because the basic results are typically already published seriously in the proceedings of a good conference (and perhaps for the same reasons reviewers of CS journals perhaps also feel less pressured to deliver reviews in a short time).

Interestingly, the changes necessary to overcome these limitations seem to already be happening.

Rapid publication journal papers instead of conference proceedings. CS journals have been trying to speed up reviewing times, and many have created special tracks of *rapid publication papers* (also called "technical notes", etc.). Such papers are guaranteed to be reviewed and published in a short time and are also more limited in length (typically around 15 pages). The papers contained in the special issues of our journals are also often in this category. This kind of CS rapid publication papers are proper journal papers which can very clearly be considered at the same level as the journal papers of other sciences and are certainly indexed as such by the standard indexing tools. At the same time, they are in fact also quite similar to CS conference papers in length and reviewing/publication timing, and, thus, dissemination capacity.

This clearly points to a way forward: *traditional CS conference papers should be published as "rapid publication" papers in (indexed) CS journals instead of in conference proceedings.* This would ensure correct indexing by bibliometric tools and I believe is the only way in which CS will be able to compete on equal terms with other sciences in the many cases in which comparisons are made (which, as argued before, will be inevitably be based on such

tools). This is independent of the preservation of the longer kind of papers, which would still remain, also published in journals and filling their quite different, CS-specific role.

Keeping our conferences. But, how can we convince ourselves as a community to switch and publish all papers in journals, even if the journal is as fast as (or faster than) a conference and the papers are of the same type as before? The change would mean turning our many high-quality conferences with a long tradition into something completely different, which is unlikely to be accepted by the corresponding communities. It is clear that it would be very hard to change our love for the look and feel of our conferences overnight. And, it is also not crystal clear that it is a good idea, since, as we discussed before, the CS model does work for us. At the very least conferences must remain to fulfill their important meeting place and information exchange roles. Thus, I believe that *any solution that is likely to be widely adopted needs to be evolutionary and somehow reconcile our traditional conferences in one way or another with journal publication.* I.e., some kind of merge of conferences with journals is needed.

Implementing such a merge is not immediately obvious since there are process differences that have to be bridged including the existence or not of a call for papers and program committee, single or multiple rounds of refereeing, etc. Fortunately, a few such proposals have already been implemented and are in operation in different areas of CS, with significant success so far. In the following, I will start by describing the solution used by the Logic Programming community (the ICLP/TPLP model), followed by some alternatives and comparisons.

3 The ICLP/TPLP model

During the time I was President of the Association for Logic Programming (ALP) I instigated the development of a model that would allow moving the publication of the papers that were traditionally published in the proceedings of the main conference in the area, the International Conference on Logic Programming (ICLP), to the main journal of the area, Theory and Practice of Logic Programming (TPLP). This journal is published by Cambridge University Press (CUP) and indexed by Thompson SCI and has several Open Access aspects, including that the papers are also published in arXiv and the ALP maintains pointers to the papers in its site.

In the design of this model (the "ICLP/TPLP model" from now on) we tried to maintain as much as possible of the traditional dynamism and flavor of the conference, while ensuring that the papers were published as bona-fide journal papers, guaranteeing the required level of quality and copy editing. The model was discussed and developed over a period of time in a close collaboration of the ALP steering committee with the Editor in Chief of TPLP, Ilkka Niemelä, and with CUP representative David Tranach, and after final approval by the community, it was then first implemented for ICLP 2010 by Torsten Schaub and myself as PC chairs. It has been in use since then and the upcoming 2013 edition will be its fourth edition.

The model is based on the tried and true notion of a *journal special issue*, in this case associated with a given year's conference. The notion of special issue allows having fixed times for refereeing rounds and keeping things within deadlines, in a framework that is well understood by editors, authors, and publishers. Furthermore, its is also compatible with the presentation of papers in a conference. In the following we describe schematically the workings of the model:

Submission

- There is a yearly *call for papers* which is issued well ahead of the conference (approximately with the same lead time as before the change to journal publication). This call is *simultaneously for submission of papers to a special issue of the journal (associated with that year's conference), as well as for presentations at the conference.* There is a submission deadline and dates for notification, final version, etc. The call also provides details regarding the (2 round) reviewing process and the fact that the papers will be presented at the conference, as explained below.

- There is an editor of the special issue that takes the role of PC chair. The traditional PC members are the special issue area editors.

Reviewing

- The special issue editor / PC chair assigns all papers submitted to the special issue (which must arrive by the submission deadline) to at least three special issue area editors / PC members. These in turn must find a reviewer for each of the papers assigned, or can review the papers themselves, and must in any case be able to discuss them.

- There is a **first round** of refereeing (involving for each paper at least three referees, with full reports) after which the PC discusses all papers. The outcome for each paper can be **reject**, **revise**, or **accept**.

- Papers for which a revision is recommended have to be resubmitted by a second deadline, together with an explanation of how the reviewer comments have been addressed. They then go through a **second round** (again, three referees, typically the same ones, and reports). The PC discusses all of these again. The outcome of this process is now **reject** or **accept**. A rejected paper can also be recommended for resubmission to the journal to undergo further refereeing, but now decoupled from the special issue (and, thus, from the conference).

- After that, all accepted papers go through a **final, copy-editing phase**, in collaboration with the publisher. Each paper is assigned to a single PC member during this phase. This last step can include additional 'shepherding' if required (done by the PC member assigned to the paper).

- The reviewing process including two full rounds and copy editing can be completed within the time frame of a traditional conference thanks to the fact that there is no printing time any more and the journal editor can produce an on-line version of the special issue in a very short time (see below).

Publication

- As mentioned before, there are no ICLP proceedings in the traditional sense any more. The full papers presented at ICLP are those accepted and published, as regular ("rapid publication") journal papers, in the TPLP special issue.[1]

[1]TPLP already had the two types of journal papers mentioned before: rapid publication papers, where the journal is committed to accepting or rejecting the paper in a short time (normally less than three months) and

- CUP creates the standard entry in the TPLP web site for this issue, complete with volume and issue numbers, table of contents, page numbers, and the papers themselves, and *makes it available just in time for the conference*. The actual physical issue may be printed and reach libraries some time after that.

- CUP can go from final version of the papers to this web page in a relatively short time (much less than the ˜2-3 months required by a standard publisher to produce and deliver printed proceedings). This allows extra time for implementing the two (or more) rounds of refereeing and the final editing step.

- All registered attendants at the conference get a password for on-line access to this web page during the conference and indefinitely from then on ("lifetime access"), which they can use at the conference and later to read papers on line, download them and read them in their laptops, or print them for personal use. The conference is also allowed to provide the papers in a USB stick for registered attendants (and has done so to date).

Presentations at the conference

- All papers accepted in the TPLP special issue are presented at the conference.

- The call for papers also includes a call for posters / extended abstracts, which, after a light review by the PC, get a smaller slot for presentation in the conference, but are of course not published in the journal special issue.

  During review process, papers that are rejected for the TPLP special issue but which are still considered to contain interesting material for discussion can be invited to submit instead an extended abstract to this poster / extended abstract track.

The posters / extended abstracts track, already present in the ICLP conference editions before the move to journal publication, allows the conference to reach further in its dissemination, early publication vehicle, information exchange, and community building roles, in a similar way to the conferences in other disciplines. This role is also played additionally by the many workshops that occur around the conference.

Long journal versions

Submission of an extended version of an already published rapid publication paper (including those ones published in one of the ICLP special issues) to TPLP as regular papers (or to another journal) is allowed. However, the new paper must contain a significant amount of new material (typically twice as much, justifying the difference in length). This is consistent for example with ACM's policy about Prior Publication,[2] which accommodates scenarios where the first publication of a paper is in a journal and the second publication of the longer and revised version of the paper is also in a journal. Note that this in no way allows getting "two papers for one:" essentially as much new material must be added as if writing a separate new rapid publication paper. The author does get a long paper and a short one (but this really means few if any any extra points in practice over publishing two of the shorter ones).

which are limited to 15 pages max, and regular papers where no such time guarantee is given and which are typically longer, often 40-50 pages.

[2] http://www.acm.org/publications/policies /sim_submissions

However, making it possible to also publish long versions allows and encourages people to really complete a topic (full results, full proofs, full experiments, ...).

4 Other models combining conferences and journals

A number of other models have been proposed and implemented that try to address the same concerns by combining CS journals and conferences in different ways.

Special issue after the conference, recommending papers to the journal

The traditional way of combining conferences and journals in CS has been to edit a journal special issue with selected papers from the conference. A variation is that the PC simply recommends a few authors to submit their paper to the journal. In this traditional model the journal publication then happens well after the conference has been held and thus a good time after the papers have been published in a formal proceedings for the conference.

It should be clear to all of us by now that this traditional model has not worked in practice, in the sense that only a very small subset of conference papers end up appearing in journals this way. Thus, this does not seem to be an effective way of solving the invisibility problem (otherwise, indexing of CS work would be working well and we would not be concerned about this issue). There are a number of reasons that prevent this approach from working in practice.

In the case of editing a special issue after the conference, if the papers in the journal are going to be essentially the same as those *already published in the conference proceedings*, then it is really double publication and it simply does not make sense. Clearly it is better in that case to eliminate the conference proceedings and *replace* them with a special issue, as proposed herein, or some other directly journal-based solution. If instead what the process expects is that all the selected 15-page conference papers will be turned into traditional CS 40-page journal papers in sync and in a short time, then it is simply not realistic. Many authors will not be able to do this in a preset amount of time, and others will not be able to do it at all. Even if some kind of intermediate solution is tolerated ("adding some material"), as is more often the case, authors still often do not have the time or even the desire (if they respect the conference more than the journal, as is often the case) to start adding to what they consider a pretty polished paper, published in a perfectly fine venue. Also, there will be many cases where the conference paper really included all results and is thus not really amenable to an extension.

As a result, the special issue may contain only a small subset of (not necessarily the best) papers of the conference, and all those that did not make it to the journal will remain invisible, not solving the problem. This solution also places a big burden on the PC chairs that have to take on the task of editing the special issue, at a time when they would rather forget all about this PC that sapped all their spare energy for the last few months. If instead of a special issue the connection between the conference and the journal is simply a recommendation by the PC chair to submit to a journal, the situation is not really much different.

These are in my opinion some of the reasons why, despite the availability of the traditional solutions of journal special issues or fast tracks after a conference, only a very small subset of good CS papers appear in journals.

A decoupled model: VLDB/PVLDB

The VLDB/PVLDB model is an approach, developed essentially concurrently with and independently of the ICLP/TPLP model, that clearly shares motivation and objectives. As a result, both models have arrived at similar conclusions and solutions in many areas, but at the same time the two models implement alternative solutions in some aspects, which are interesting to explore and discuss.

In the VLDB/PVLDB model a new journal (Proceedings of the VLDB Endowment – PVLDB) in order to implement the scheme. In PVLDB papers are limited in length and there is guaranteed quick turn-around for reviews. Submissions to PVLDB are accepted all year round. All submissions are reviewed by three members of the PVLDB editorial board (in single-blind mode) and the reviews are guaranteed to be provided within 2 months. When a submission needs a revision, the revised paper, which can and should be accompanied by detailed explanations of the authors, will be reviewed by the same reviewers that read the original submission. The final versions of accepted papers are checked by the responsible editors for journal-level quality assurance. Rejected papers are barred from being resubmitted to PVLDB for one year from the date of submission, not considering any invited revisions submitted. If accepted, the paper appears in PVLDB and the author is invited to present at the next VLDB (Very Large DataBases) conference. Papers submitted in February or March are guaranteed to be reviewed in time for the August VLDB conference (but may not have time for a round of revision and re-reviewing). The PVLDB experiment has now been in place since 2009, and appears to be going well.

One obvious difference between PVLDB/VLDB and ICLP/TPLP is the fact that PVLDB is a dedicated journal, publishing exclusively papers to be presented at VLDB, and created for the purpose of implementing the model, while TPLP is a preexisting journal, which publishes other (both rapid publication and long) papers in addition to the special issue whose papers get presented at the ICLP conference. The ICLP/TPLP solution was made possible by the fact that both ICLP and TPLP are respectively the official main conference and journal of the Association for Logic Programming and the collaboration of CUP. However, although convenient (because it saved the complication of creating and indexing another journal) and elegant (because it groups the two sets of papers in a single series) I do not view this coupling as a fundamental part of the ICLP/TPLP model, which could work also work well with a dedicated journal. I.e., the ICLP/TPLP model could be easily adopted by a different community even if an existing hosting journal is not available in the area, by simply creating a dedicated journal in the PVLDB style.

In my mind the most significant difference between the ICLP/TPLP and the VLDB/PVLDB model is the "decoupled" nature of the latter. In the VLDB/PVLDB model, instead of having a deadline as in traditional conferences or journal special issues, continuous submission is allowed and all the papers accepted in the (dedicated) PVLDB journal, up to a certain cutoff date, are then presented at the conference. This does have the advantage that authors can submit at any time and that reviewing time can be longer for some papers (depending of course on how close their submission of to the cut-off date). The counterpart is that this model loses more of the traditional traits and flavor of the CS conference process. In contrast, the ICLP/TPLP model preserves the call for papers and sets clear deadlines for the different steps of the reviewing process, in cycles that everyone understands and is used to respecting, without any need to find area editors that are assertive and disciplined enough to enforce strict deadlines asynchronously. It keeps the process fully predictable, while still improving

the conference process by implementing two or more rounds of refereeing. It also preserves the PC and the (possibly on-line) PC meetings, so that all papers are looked at and discussed at the same time by the same set of people, allowing direct comparisons, as is done now in our conferences.

Overall, the ICLP/TPLP model was designed to ease adoption by both the LP community and other communities and conferences, by preserving essentially intact the "look-and-feel" of CS conferences, from the point of view of both PC and authors, except for the improved, two-round reviewing process. My personal opinion is that in this sense the VLDB/PVLDB model could be perceived as too large a departure from the traditional conference model for it to become adopted by the majority of the top conferences in CS.

SIGGRAPH+ACM ToG

The SIGGRAPH+ACM ToG is also one of the early models implemented and has some interesting characteristics. Basically, papers accepted to ToG (ACM Transactions of Graphics) may also be presented at the SIGGRAPH or SIGGRAPH Asia conferences. Conversely, the proceedings of SIGGRAPH and SIGGRAPH Asia are printed as is as special issues of ToG. This is similar to ICLP/TPLP except that the traditional conference refereeing is untouched, i.e., there is one round of review, with three referees which have seven weeks to evaluate the paper, followed by a two-day rebuttal, and finally a PC meeting. A PC member is responsible for ensuring that the paper's authors take suggestions for improvement into account in the final version but there is no second round of refereeing as in the ICLP/TPLP model in which authors submit a revised version of the paper and explain how reviewers comments have been addressed, which get to be checked a second time by the reviewers.

5 Conclusions

I have offered my viewpoint on some of the problems present in the CS publication culture and why they are important, and have also elaborated on some realistic ways forward. In particular, on the issue of conferences vs. journals I have presented the model adopted a few years back by the logic programming community for its main conference and journal, the ICLP/TPLP model, and also compared it to some of the other interesting models presented to date. The ICLP/TPLP model is based on the assumption, supported already by many CS journals, that CS journal papers can take the form of rapid publication papers (similar to those of other sciences and also close to our conference papers) which can live alongside the longer journal papers that are traditional in CS. The concrete proposal is then to have papers submitted to a special issue of an (indexed) journal which is ready online in time for the conference, rather than publishing the traditional conference proceedings. It also proposes to improve the traditional conference reviewing process to follow journal standards for rapid publication papers and special issues by including at least two full rounds of refereeing and a final copy editing step.

I argue that this model offers an evolutionary path to solving a good number of the incompatibility problems with other sciences of the current CS conference publishing model, without giving up the essence of our conferences. Furthermore, I argue that its adoption would also be a net improvement within CS since nothing is lost by simply changing the publication medium and the move to two rounds of reviewing can only improve the quality of the published papers.

I believe that the model presented (or one of the closely related models being proposed) would be relatively straightforward for the community to adopt widely by conferences in CS and this adoption could finally make our papers come out of their conference-tainted reputation outside CS and their invisibility to bibliometric tools.

Acknowledgments

A large number of people participated in the design of the ICLP/TPLP model and even more in making it a reality. As mentioned in the text Ilkka Nielmela as TPLP editor in Chief and David Tranach on the side of Cambridge University Press were both highly involved in the design. Torsten Schaub had an important part as co-chair with me of the first (2010) edition of the new ICLP/TPLP publication model. We jointly worked out the many implementation details (including working out how to coax EasyChair into implementing multiple runds of reviewing) that inevitably needed to be solved. The help of Andrei Voronkov (EasyChair main author and guru) on these matters is also acknowledged. The PC members of the 2010 ICLP and the PC chairs and PC members of subsequent ICLP editions contributed greatly by taking on the extra burden of the rounds of refereeing and the proof reading of final copies. Manuel Carro as conference coordinator for the ALP has helped implement it and Gopal Gupta as current President of the ALP has continued to support the approach. Finally, the ALP steering committee members and the ALP community as a whole supported the move and contributed in a number of meetings over the years while the proposal was being discussed. While all these people contributed to the good parts of the design, any shortcomings, the opinions on this paper, and, very specially, any possible errors are my own responsibility.

This work has received funding from the Madrid Regional Government under CM project P2009/TIC/1465 (PROMETIDOS), and from the Spanish Ministry of Economy and Competitiveness under project TIN-2008-05624 *DOVES*.

References

[1] David Parnas. Stop the Numbers Game, CACM, Nov. 2007.

[2] Friedemann Mattern. European Computer Science Summit. 2008.

[3] Jeffrey C. Mogul, Tom Anderson. Before and After WOWCS: A literature survey, A list of papers we wish had been submitted. WOWCS'08 Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems.

[4] Vardi, M.Y. Conference vs. journals in computing research. Commun. ACM 52, 5 (May 2009).

[5] Fortnow, L. Time for computer science to grow up. Commun. ACM 52, 8 (Aug. 2009), 33–35.

[6] Vardi, M.Y. Revisiting the publication culture in computing research. Commun. ACM 53, 3 (Mar. 2010), 5.

[7] Denning, P.J., Keromytis, A.D. Letters responding to March Editor's Letter. CACM, 53, 5, 6 (May 2010).

[8] Halpern, J.Y., Parkes, D.C. Journals for certification, conferences for rapid dissemination. CACM, 54, 8, 36-38 (August 2011).