

QUO VADISC? OUTCOMES OF THE PODC/DISC CONFERENCE MODEL SURVEY

Stefan Schmid

Abstract

At the ACM PODC 2020 business meeting and on Zulip, a task force was formed to propose options for changing the PODC/DISC deadline schedule and/or moving to a publication model based on journal-style reviewing. This article summarizes the models identified by the task force and their underlying rationales, and reports on the results of a recent survey among community members, conducted by the task force.

1 Introduction

The task force commissioned to explore possible alternative PODC/DISC conference and publication models consists of seven active members of the community:

1. Christian Cachin
2. Faith Ellen
3. Yuval Emek
4. Rachid Guerraoui
5. Christoph Lenzen (chair)
6. David Peleg
7. Jukka Suomela

In September 2020, the task force conducted an extensive survey in the community. The survey was announced on the PODC mailing list. Despite the relatively short time frame and the extensive survey (reading and responding to the specific questions took around one hour), 69 anonymous community members participated in the survey.

In this article, we first present and discuss the models the task force identified, and then report on the survey results. The article may serve as a basis for the upcoming community meeting at DISC (Oct 12 - 16).

2 The Models Suggested by the Task Force

This section reviews the current situation and the possible models, as identified by the task force.

2.1 The Current Situation

The task force survey focuses on the two main annual conferences:

1. PODC (under ACM), around July-August
2. DISC (under EATCS), around October

Publication practices of the community are currently centered around conference abstracts rather than journal papers. The task force identified the following potential shortcomings of the current situation:

- The spacing between PODC and DISC deadlines is imbalanced.
- The full versions of many papers never get published (and often are not written). The full versions of papers on arXiv are not peer-reviewed, which can make it problematic to build on top.
- In many schools, promotion and tenure committees treat conference papers as less important than journal papers. A lack of journal papers can negatively affect the career development of young researchers in the community.
- Acceptance decisions are not only a question of merit, but may also depend on the availability of slots.
- Randomness in the decision process can cause unpredictable publication delays for solid, but not outstanding papers. It may also result in authors being less willing to improve their rejected papers before resubmitting them.

2.2 Possible Alternative Models

The task force proposed three alternative models. The three suggested models all share the following properties:

- Unifying PODC and DISC into two annual conferences of equal standing, called The Summer Symposium on Principles of Distributed Computing (S-PODC) and The Winter Symposium on Principles of Distributed Computing (W-PODC). While the task force considers the specific names S-PODC and W-PODC as placeholders, the idea is that they should capture the focus on fundamental aspects of distributed computing systems. A new Society for Principles of Distributed Computing could be established to oversee the organization of the two conferences. In its most minimal form, the new society simply assumes the roles of the steering committees of PODC and DISC.
- Having 2 or 4 equi-spaced deadlines.

2.2.1 Model 1: Conference Reviewing

Summary: The first suggested model requires the smallest change from the current state of affairs. A comparative discussion about the relative advantages of having 2 or 4 deadlines will appear at the end of this section. The description below is for 4 deadlines; the 2-deadline variant is self-explanatory.

Details:

- Each conference is associated with two review cycles so that in total, we will have four equally spaced deadlines per year.
- Each conference is associated with one CFP that announces its two review cycles.
- Papers accepted in a review cycle associated with conference X are presented during conference X. (This presentation will be relatively short, 5-10 minutes per paper.) They are also included in the proceedings of conference X, which means that the proceedings include papers from two different review cycles.
- The authors of each accepted paper upload a manuscript and a pre-recorded video presentation of their work to the conference website shortly (a few weeks) after the notification.
 - Regular papers: manuscript = full version; video length 25 minutes
 - Brief announcements: the manuscript contains (at least) the submitted content; video length 10 minutes
- The proceedings of conference X are produced after the second review cycle associated with conference X (around the actual dates of the conference).

- The appointment of a PC chair for conference *X* means that he/she is appointed for the two review cycles associated with conference *X*.

To be determined:

- Regarding the appointment of (non-chair) PC members, the task force sees three alternatives:
 1. Each PC member serves on one review cycle; the entire PC (possibly excluding the chair) is replaced from one review cycle to the next. (This is relevant mainly for the 2-deadline model; with 4 deadlines, it might not work well.)
 2. Each PC member serves on the two review cycles that are associated with the same conference; the entire PC is replaced from one conference to the next.
 3. Each PC member serves on two review cycles; half of the PC is replaced from one review cycle to the next (a “rolling PC”). It is recommended that the appointment of PC members that serve under two PC chairs will be handled by both of them.

Optionally, each submission may be associated with some PC members responsible for it, requiring PC members to remain active on all the submissions under their responsibility until the fate of these submissions is decided. This implies that submissions arriving on deadline *X* should preferably be assigned to new PC members who joined the PC on round *X*, to minimize the “overflow” of PC duty.

- Relations with ACM and EATCS: It needs to be determined if the community wants to stick with the current umbrella organizations. If so, one could associate the summer conference with ACM and the winter conference with EATCS (or vice versa). This means that the production of the proceedings of conference *X* will be handled by ACM/EATCS, depending on whether *X* is the summer or winter conference.
 - The task force suggests to have a unified format for the proceedings of the two conferences, but this may be difficult given that ACM insists on their own format. An effort needs to be made to ensure that the space bounds are more or less equivalent.
- Enforcing limitations on re-submissions: If the number of deadlines per year are doubled, one should consider enforcing certain limitations on re-submitting rejected papers. There are two alternatives (which could also be combined):

1. Hard limitation: e.g., a paper that has been rejected twice cannot be submitted again (to either of the two conferences).
2. Soft limitation: a submission of a previously rejected manuscript must be accompanied by a cover letter that includes the previous reviews and explains how they have been addressed — why is the currently submitted manuscript better than the previously rejected one(s)?
 - This rule could be applied to any submission that includes material from a previously rejected manuscript. This could be applied to resubmissions rejected from other venues too.

2.2.2 Model 2: Hybrid Conference+Journal Reviewing

The goal of Models 2 and 3 is to implement a transition of the community into a publication model based on journals rather than conference proceedings.

Summary: Model 1 augmented with the requirement that the full versions of all accepted conference papers have to be submitted for publication in the journal associated with the conferences soon after the respective conference.

Details:

- The conference CFP will state that the authors of accepted papers are also expected to submit their work to the Distributed Computing journal, max. X months after the conference presentation. The authors can also choose to publish in a different journal, as long as they do it in a timely manner.
- When the PC members accept an invitation to join the PC, they also agree to referee or find appropriate referees for the journal submissions for the papers they already handled in the conference. The PC members are committed to handle the first round of reviewing in max. X months and revision rounds in max. X months.
- The Distributed Computing journal follows its usual protocols and its own standards, and it can also freely use reviewers outside the PC. If the paper is not strong enough or good enough, it can get rejected.

To be determined:

- How to ensure that, for a short contribution, exactly the same text can serve as the “conference submission” and as the “journal submission”. A possible solution would be to negotiate with the publishers to avoid copyright issues. The conference version could be a short abstract with a link to the full version.

- Do we expect that external reviewers of the conference papers will also help with journal reviews? Or do we expect that the PC members who used external reviewers for the conference papers will find new external reviewers for the journal submissions if needed?

2.2.3 Model 3: Journal Reviewing

Summary: This model is based on establishing a journal for the Society for Distributed Computing, called Transactions on DC (TDC), or some other suitable name, and requiring authors to submit a full paper, jointly to the conference and the journal. The journal will serve as the society's main archival medium, effectively replacing the proceedings in their current form (although we may still have thinner proceeding volumes, based on 1-2 page abstracts). Once a submission is accepted, it gets to be presented in the conference (specifically, the nearest one) as well as published in its full version in TDC.

Details:

- Authors are required to submit a full paper, which is evaluated thoroughly as in journals.
- To reduce the reviewing load, the reviews shall consist of two or three stages:
 - Stage 1 (quick evaluation): The PC (possibly using sub-reviewers) evaluates significance and interest (following PODC's current standards). This stage should take 3-4 weeks, and its outcome is either "reject" or "proceed to a full review". Typically, submissions that pass on to the next stage are likely to be accepted (except in rare situations such as an error or another serious problem).
 - Stage 2 (full review): The PC assigns referees who will go over the entire paper and verify it, including correctness. If the reviewers recommend acceptance (possibly subject to minor revisions), then the paper will enter the "accepted papers list" of the current PC. But the reviewers may recommend a major revision, in which case an additional stage is needed, and the paper will not be included in the current list of accepted papers.
 - Stage 3 (revision): If the reviewers recommend a major revision, the authors will be asked to revise it. The revised submission will be sent to the reviewers again, who will recommend either acceptance or rejection. When the paper is accepted, it will enter the "accepted papers list" of the current cycle.

- Each step (quick evaluation / full review / revision) will be given a strict time bound.
- Rejections, as well as “smooth” acceptances, are made within the time frame of a single review cycle. In contrast, in case of a revision, the entire process might take longer than 3 months, and thus span more than one review cycle.
- A submission made on some deadline X may end up in either of the two conferences, depending on how long it took to review / revise it until it is accepted. In other words, when authors make a submission on deadline X, they do not know which conference they will end up in. This has several implications. In particular:
 - The call for papers for some deadline X should be issued by the DC society jointly for both conferences.
 - Consecutive PC’s affect each other. In particular, each PC passes on a “commitment” to the next PC, in the form of a collection of submissions that passed Stage 1+2 and are now under revision. In other words, each PC starts with two piles of submissions: “old” submissions, currently under revision, and “fresh” submissions.

To be determined:

- Should TDC be a full-fledged journal, and accept also other types of submissions, e.g., of full papers based on extended abstracts that have appeared in other theory conferences such as FOCS / STOC?
- What happens to Distributed Computing (DC), the current main journal of our community? Do we want DC to serve the role of TDC? Currently DC has limited volume, and it is perhaps slightly different in spirit and goals from the envisioned TDC.
- Who will run TDC? In the current model, our main journal (DC) is run by Springer-Verlag, and our proceedings are produced for us by ACM (for PODC) and Springer-Verlag resp. more recently LIPIcs (for DISC).

2.2.4 Discussion: 2 vs 4 Deadlines

All three models can work with 2 or 4 equi-spaced deadlines a year. The task force identified the following advantages of 2 deadlines:

- Simpler PC organization (separate PC for each conference, no need for staggering, fewer PC meetings per year).

- Might fit better with Model 3, where longer time could be needed to complete the evaluation of some submissions.

An advantage of 4 deadlines is:

- Shorter waiting from the time a paper is ready for submission until it is accepted.

2.3 Independent Methods

The task force considers the following changes useful regardless of whether the current system is continued or one of Models 1 to 3 is adopted.

- No excuses regarding full versions: they are being submitted as well (under all models), and will be archived in a public repository in case of acceptance. Updates are fine, but the original version is guaranteed to be available.
- Upload a recorded talk within a few weeks after acceptance, which is archived and published.
- Ask authors who submit papers to pledge to be available as subreviewers.
 - Suggestion: the submission process includes a mandatory “topic preferences” form (that each co-author must fill out).
 - May help with reliable and fast recruitment of reviewers, so the task force considers this particularly important when choosing Model 3 with 4 deadlines.

2.4 Relation with Publishers

The basic assumption is that the community should publish with gold open access. This means that there must be funds to pay for the production costs. Currently, with ACM, individuals who want open access for their paper pay for it directly; this is expensive (more than 1000 USD). LIPIcs currently publishes DISC proceedings with gold open access and the costs are much smaller than with ACM (less than 100 USD). However, LIPIcs is only an option for conference proceedings, not for a journal.

If there is an independent society, then cost may also be absorbed by that society and its members. For example, volunteers could run an instance of OJS (Open Journal Systems) and the society pays for minimal residual cost (i.e., a web server and backups). The IACR (<https://www.iacr.org>) currently

uses this model for ToSC (<https://tocs.iacr.org>) and TCHES (<https://tches.iacr.org>).

An issue with the current situation of scientific publishing with gold open access is that it may introduce both labor cost (for work done by volunteers) and monetary cost (for work that cannot be done oneself). Such costs can be charged to authors and attendees (per paper or per conference), as currently done by DISC with LIPIcs. On the other hand, it could be charged to readers, with a paywall around the library, as currently done by Springer and by ACM (only those who pay for the online library can view). No matter which option is taken, it is usually in the researchers' interest to ensure that ownership of publications remains with them and that copyright is not given away to a commercial publisher.

2.5 Comparison Between the Models

The task force summarizes the relative pros and cons of the models from their perspective as follows. While some points were controversial, the following gives the overall impression. The term “Model 0” is used to refer to the current state of affairs.

2.5.1 Model 0 (status quo)

The main advantage the task force sees in favor of the current system is simply that there is no need for change. The perceived disadvantages identified by the panel are listed under “main shortcomings” at the beginning of the document.

2.5.2 Model 1 (conference reviewing) vs. Models 2 & 3 (journal reviewing)

Model 1 makes minimal changes, with the goal of recalibrating PODC and DISC deadlines to be equally spaced. As we currently perform conference reviews within 3 months, the 4 deadline version would spread out the reviewing more equally throughout the year, while reducing the (average) time between a paper being ready for submission and its publication. However, members of the program committee have work assigned to them for a longer period of time. Models 2 & 3, on the other hand, in addition seek to change the current approach to reviewing to be, in essence, that of a journal. The hope would be to have higher quality reviewing with all the associated improvements over the current system (see main shortcomings), at the expense of a higher review workload. Opinions regarding how much additional workload would actually be incurred varied, as there is also the hope of reducing the number of times papers are evaluated from scratch and/or resubmitted without substantial changes.

2.5.3 Model 2 vs. Model 3 (conference reviews first or immediate journal reviewing)

Model 2 proposes to have the journal submission follow the current conference reviewing and publication process. The idea is that it guarantees that a decision on (conference!) acceptance or rejection is taken within a fixed time bound. Accordingly, this model is of highest interest with 4 deadlines. In addition, there is an opportunity to improve on the original submission based on reflection and the feedback received at the conference. Model 2 appears to be more compatible with maintaining ties to ACM than Model 3. However, implementation issues relating to short papers whose journal versions are essentially identical to their conference versions might arise. There were concerns that Model 2 would cause an even higher reviewing load on the community, because accepted conference papers would receive journal reviews for modified versions after a time gap of several months.

The main idea of Model 3 is that we would not only have a thorough reviewing process, but acceptance decisions would be based on a full evaluation from the start. In contrast to Model 2, no separate, additional “conference style” reviews would be performed. However, an initial screening process is foreseen to avoid fully reviewing submissions that would not pass the bar assuming correctness. The possibility of not receiving a final decision within a single review cycle in case of a major revision is a disadvantage. However, it is an advantage that a paper requiring major revisions is not accepted prematurely.

With 4 deadlines, there is no more time than today for a conference review, so care has to be taken to not overload reviewers in Model 3. While 2 deadlines (rather than 4) could be considered to mitigate this issue, this would further add to the delay between papers being ready for submission and receiving a first acceptance notification. However, the task force agreed that Model 3 with 4 deadlines is feasible if (i) the community fully supports it and (ii) authors submitting papers are required to pledge to serve as subreviewers.

3 Results of the Survey

In September 2020, the task force shared a document with the contents above on Google Drive, and distributed a survey form via the PODC mailing list, asking for inputs from the members of the community. Although the survey was open only for a relatively short time period (September 15-24), to have the results ready for discussion at the DISC conference, and despite the non-trivial time investment required (according to the email by the task force, 60-90 minutes), a significant amount of input was received.




Figure 1: Left: Which geographic region do you feel most closely associated with? Right: How long have you been active in the field?

Let us start with some statistics. The survey was completed by 69 community members:

- 10 female, 50 male, 9 no answer or other
- 40 from Europe/Middle East, 21 from North America, 3 from East Asia, 2 from India, 1 from South America (see Figure 1 (left))
- 39 faculty members with tenure, 11 faculty members prior to tenure, 12 postdocs, 4 industrial researchers, 1 graduate student, and 2 others
- 11 participants had more than 20 PODC/DISC papers, 11 between 11 and 20, 29 between 4 and 10, and 16 between 1 and 3
- 61% of the participants have more than 10 years experience, 35% have 4-10 years experience, and only 4% have a shorter experience (see Figure 1 (right))

Regarding the popularity of the models, Figure 2 shows the feedback received on the question:

- For each model and number of deadlines combination, do you think the combination is good, bad, or neutral?

Especially Model 1 with 4 deadlines seems to receive interest. Also Model 2 with 4 deadlines and Model 3 with 2 deadlines receive good feedback, however, there are more concerns about the detailed implementation, as we also see in the individual text answers that some participants provided (see below for more details).




Figure 2: Model and deadline acceptance overview

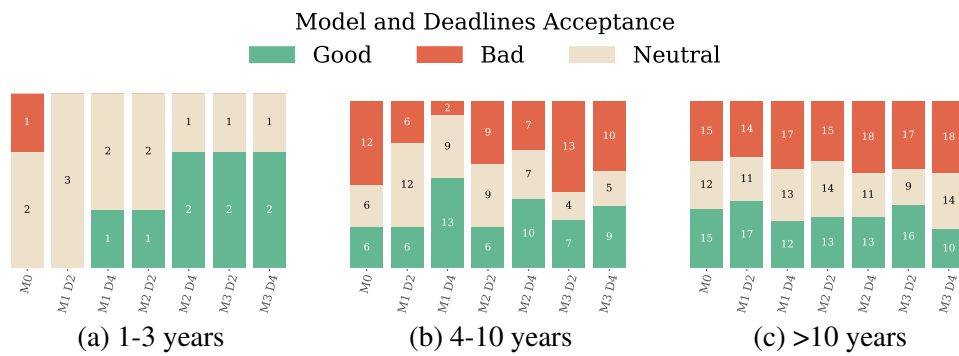


Figure 3: For each model - #deadlines combination, do you think the combination is good, bad, or neutral?

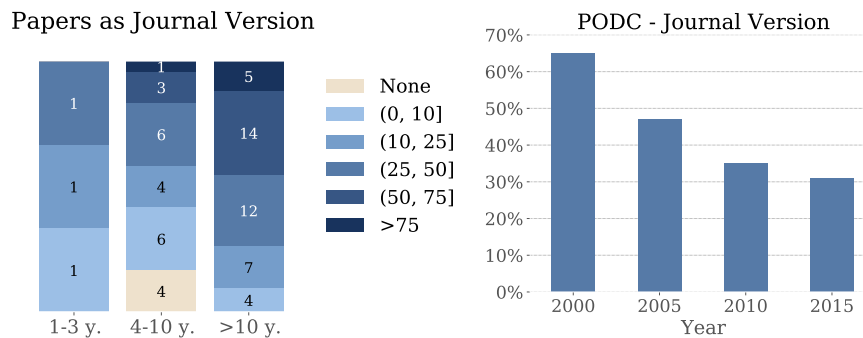


Figure 4: Left: Percentage of journal versions today. Right: PODC journal statistics.

Figure 3 shows the detailed evaluations of the models according to the seniority of the participant. Unfortunately, the number of junior participants is small, so conclusions are difficult, however, the preferences seem to differ across the age groups.

Regarding journal publications, the community was asked:

- Approximately what percentage of your conference papers have journal versions?

Figure 4 (left) shows the results, categorized again according to seniority. The figure confirms that only a small fraction of participants regularly publishes their contributions in journal versions. To complement this information, David Ilcinkas kindly shared with us his statistics on the proportion of PODC papers which were finally published in a journal: Figure 4 (right) shows the results for four specific years (2000, 2005, 2010, 2015), indicating a “dejournalization”.

Figure 5 shows the input received regarding the questions:

- *Full version*: Should we require the submission of a full version that will be publicly archived (unless it has been publicly archived already) upon acceptance? The full version would remain confidential if the submission is not accepted.
- *Video*: Should authors be required to upload a video presentation within a few weeks after acceptance, which will be archived and published?
- *Pledge to review*: Should authors of submitted papers be asked to pledge that they will serve as reviewers (up to X reviews per submitted paper)?
- *Additional mechanisms*: Should we seek to develop additional mechanisms for incentivizing high review quality and improving load distribution? The

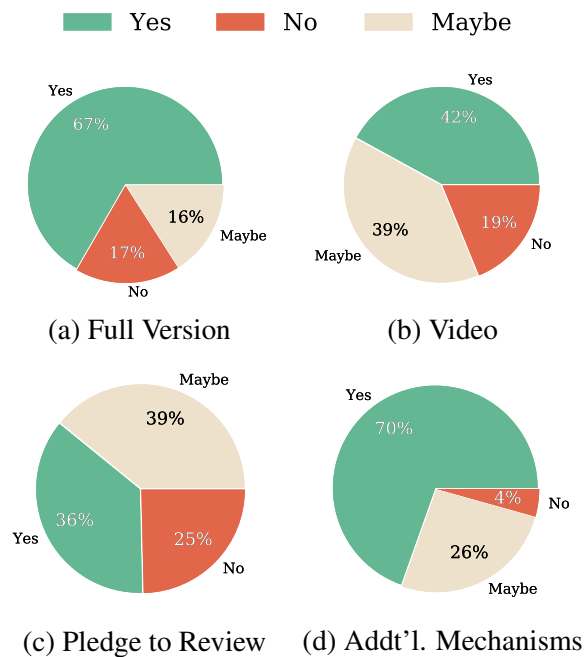


Figure 5: Model independent methods

assumption here is that proposals would be worked out and reviewed by the community (if they cannot be tried out with low effort and risk) first accepted.

The figure suggests that there is a demand in the community to implement additional mechanisms to improve the review process, and developing mechanisms accordingly could be a priority. Also the requirement to submit full versions seems to find significant support. There are more than twice as many proponents of the idea to require a video than there are opponents. Also the idea to pledge to review finds more proponents than opponents, but it is clear that further discussions are required on how to implement this change.

4 Discussion

The participants were also given the opportunity to comment on the pros and cons in written form. Overall, many participants seems to agree that the review quality should be improved, that journal versions (or at least full arXiv versions) are often desirable, and that adjusting deadline can make sense. However, there are also valid concerns that these models may result in more work for the community.

The fact that Model 1 is popular may also be related to the various open questions that Models 2 and 3 raise. Judging from the individual answers, there is more uncertainty about these models, which is not unexpected, given the larger change and more open questions. For example, it is important to understand, e.g., the implications of a name change (e.g., on the conference ranking), how to best organize the live gatherings at the conference (should it be more than a “Highlights on” event?), how to enforce the specific rules (e.g., journal submissions), etc. As one participant pointed out, Model 3 may have the advantage that there are already similar examples where conferences have their own proceedings such as VLDB, or how OOPSLA, ICFP and POPL are linked to the PACMPL journal. Someone also pointed out that other communities (e.g., EMSOFT) are already using Model 2. Some participants also propose intermediate models, e.g., a Model 1.5 where there are conference publications, but there are one or two rounds of rebuttals (orthogonal from the discussion of number of deadlines and their spacing).

It is clear that while there is a significant interest in the community in innovative models, many more discussions on the implementation details are needed, perhaps also considering to perform changes in multiple stages. The ongoing pandemic may naturally introduce additional models and I suggest to repeat a similar survey in the near future, hopefully with less time pressure and at a larger scale, also encouraging the participation of more junior researchers, which are under-represented in the evaluation above.

Acknowledgements

Kudos to David Ilcinkas for sharing with us his statistics on the PODC journal versions (Figure 4, right) and to Thomas Fenz for his help with the figures.