

ICSE Program Co-Chairs Guiding Document

Version 1.0

July 3, 2017

Authors

Gerardo Canfora
Marsha Chechik
André van der Hoek
Martin Robillard
Willem Visser
Jon Whittle

The document consists of three parts. Part I provides basic context. Part II stipulates the process that the ICSE Steering Committee has adopted for making changes to the review process for the research papers track. Part III contains additional (and non-mandated) wisdom amassed from previous PC Co-Chairs regarding various aspects of the review process.

Appendix A serves as a “living document” capturing the latest agreed-upon review process.

Part I: Basic Context

Motivation

This document is meant to guide PC Co-Chairs of future ICSEs, providing them the basic framework of best practices of how to serve as a PC Co-Chair, how to enact change in the review process, and how to contribute to making the lives of subsequent PC Co-Chairs easier. The document grew out of discussions of the ICSE Technical Review Subcommittee of the ICSE Steering Committee, and has been discussed and approved by the Steering Committee at large.

Priorities

The ICSE review process is a difficult ‘beast’ to manage, beholden to a variety of stakeholders. PC Co-Chairs have to keep in mind the community of authors, the reviewers, the conference itself, and even the broader community beyond ICSE when it comes to the conference’s reputation. Given this, it is useful to consider the following priorities:

1. *Quality first.* The overarching purpose of the review process is to develop a quality program for ICSE. This quality must be carefully protected, as it is the key factor determining the reputation of ICSE, both within our community and outside of it. First and foremost, thus, the review process is concerned with selecting high-quality papers that will form the core of the ICSE conference program.
2. *Authors second.* ICSE is critically dependent on the willingness of authors to submit their papers to the conference. It has a heritage of being broadly inclusive, both in terms of evolving topics of interest and in terms of nurturing the community. It must not be forgotten, then, that the ICSE review process serves its constituency and should do what it can to promote an inclusive atmosphere.
3. *Reviewers third.* ICSE is also critically dependent on the willingness of reviewers to serve the conference. Without qualified reviewers who invest their time to review papers, there would be no conference. While serving as part of the review process is an honor, the review process thus should also take into account what it asks of reviewers.

It is impossible to achieve a perfect review process that optimizes all of these priorities; tradeoffs exist with nearly every decision that is made. Limiting the number of submissions for sure benefits reviewers, might increase quality, but certainly disadvantages some authors. Providing three reviews to every paper increases the satisfaction of authors with the review process, but at a cost of increasing the review load. Many other examples exist.

Part II: ICSE SC Policy for Changing the Review Process

Process

Recognizing the difficulties associated with choosing a satisfactory review process, and reconciling the fact that PC Co-Chairs need the ability to tailor the process while the community has a need for longer-term stability, ICSE adopts the following process:

- *Approving changes.* Any desired changes to the ICSE review process that are visible from an author perspective (e.g., double blind or not, number of reviews per paper, rebuttals, review ratings) or represent a significant departure from previous practice (e.g., program board model versus program committee model, online PB or PC meeting), should be first presented to the ICSE Technical Review subcommittee, with a detailed description of the problem being addressed, an analysis of pros and cons, a careful articulation of the potential impact of the proposed changes (especially in light of the above priorities, both positive and negative impact), and a plan for measuring this impact. Changes will be vetted by the technical review subcommittee (which by definition includes one of the PC Co-Chairs who may be proposing the changes). Changes that the ICSE Technical Review subcommittee considers relatively innocuous can be approved by the ICSE Technical Review subcommittee together with the Steering Committee Chair without involving the entire Steering Committee. Changes that are more controversial, or have significant impact, will need full approval of the Steering Committee.
- *Promoting openness.* Particularly controversial or different choices of review process shall be subject to a phase of community input. This phase will take place after consideration of the ICSE Technical Review subcommittee, but before consideration by the ICSE Steering Committee at large so to maximize input to the discussions of the Steering Committee. The chair of the ICSE Technical Review subcommittee shall be responsible for soliciting input from the broader community through a timely, public, and online solicitation. This could take the form of an online survey, an e-mail request for comment, or presentation at a townhall meeting (use of Facebook or other social media is problematic and discouraged, since it excludes a portion of our community). Of note that is that it should be made clear to the community that the ICSE Steering Committee is seeking its input and suggestions, but that the community is not being asked to vote on any particular proposal – that remains the task of the ICSE Steering Committee. Managing expectations, then, is key to any round of community input. It is strongly suggested that the ICSE Steering Committee also closes the loop with respect to the community: decisions that touch authors, and especially the decisions for which community input was actually sought, should be communicated back to authors in a timely and ideally open manner (for instance, through a townhall or e-mail to the ICSE mailing list).
- *Maintaining stability.* To provide the community with stability, and to be able to properly assess the impact of a change over time, changes to the review process will remain in effect for a minimum of three years, with two exceptions:
 - A change will be retracted if it is clear that it had significant, unanticipated negative consequences.
 - A change can be amended with a next suggested change if the next change is determined by the ICSE Technical Review Committee to be a friendly amendment, in line with the intent of the original change and attempting to further improve its intended effectiveness.
- *Conducting assessment.* Any change shall be accompanied with a plan for assessing its impact, both positively so and negatively so. This plan shall be put into action by the PC Co-Chairs (in consultation with the ICSE Technical Review subcommittee, if so desired), with the results presented and reflected upon in the final report from the PC Co-Chairs and discussed, as needed, during a subsequent ICSE Steering Committee meeting.

While this process constrains future PC Co-Chairs by virtue of decisions made in previous years, this process balances the needs of the various stakeholders, helps future PC Co-Chairs chairs in providing a blueprint of the review process upon which they can build, and encourages careful exploration of alternative review models when so needed.

Current Review Process

For the process described above to be effective, the current review process should be carefully documented, as well as updated each year that involves changes to the process.

Appendix A is meant to serve as the ‘living statement of the ICSE review process as decided upon by the ICSE Steering Committee (by delegation to the ICSE Technical Review subcommittee)’; it thus should be updated accordingly when changes are approved. The chair of the ICSE Technical Review Committee is responsible for doing this, in cooperation with the PC Co-Chairs who suggested changes that were approved.

Part III: Additional Amassed Wisdom

Timeline

The above process requires a longer lead time than usual for PC Co-Chairs, because changes to the review process might need to be reviewed and discussed by the ICSE Steering Committee as a whole (with particularly significant changes having been discussed twice in the past by the Steering Committee).

The timeline below builds in time for one round of discussion by the Steering Committee as a whole (should that be needed); if particularly challenging changes are being proposed, a longer lead time than the timeline below might be needed, in which case it is suggested that the changes are resubmitted to the Technical Review subcommittee in time for discussion by the ICSE Steering Committee in May of year $n-2$.

Year	Month	Action
$n-2$	Jul	Submit possible changes to Technical Review Subcommittee
$n-2$	Jul/Aug/Sep	Invite review team (option 1; see common challenge #1)
$n-2$	Sep or Nov	SC meeting to discuss possible changes (if deemed necessary by Technical Review Subcommittee)
$n-1$	Jan	Invite review team (option 2; see common challenge #1)
$n-1$	Feb	Publish call for papers
$n-1$	May	Meet with review team in person
$n-1$	Aug/Sep	Paper submission deadline
$n-1$	Dec	Paper accept/reject decisions
n	Jan	Review team and author survey

<i>n</i>	Apr	PC Co-Chairs review report drafted, so it can be shared at ICSE SC meeting
<i>n</i>	May	ICSE!
<i>n</i>	May	SC meeting to discuss review report (as needed)
<i>n</i>	Aug or Oct	PC Co-Chairs review report finalized and officially published on the ICSE SC web site
<i>n</i>	Aug or Oct	Upload all used materials (standard e-mails, slides, review process) to online archive for future PC Co-Chairs

Common Challenge #1: Selecting the Review Team

Irrespective of the review process being followed, selecting the review team that actually conducts the reviews of the papers and assists the PC Co-Chairs in deciding whether to accept or reject a paper is a non-trivial undertaking. Per the Memorandum of Understanding and per accepted good practice, it is important to strive for a balance in composition of the review team, topics covered, expertise, university and industry affiliation, geographic area, seniority, and gender. In the past, each set of PC Co-Chairs has generally developed its own approach to performing this selection process. Yet, some tips from previous PC Co-Chairs include:

- It is typically useful to develop a matrix of topics versus the planned review team, together with anticipated number of paper submissions in each topic (based on previous years, adjusted per the PC Co-Chairs expectations). This matrix can help in many ways, from finding out which has over-capacity in some topic, to holes in coverage, to – perhaps most importantly – identifying candidate reviewers during the paper assignment process. Examples of this type of matrix from previous years are also valuable to future PC Co-Chairs.
- Nearly every year, some topic suddenly is much more popular than in previous years, or some new topic emerges that just was not on the scene before. It is difficult to anticipate these kinds of swings and trends, though to at least think about ‘what might be coming’ tends to help in building some overcapacity in the review team.
- Some reviewers are highly specialist, others are generalists. It is often useful to balance the presence of these kinds of reviewers in the review team.
- The more people one can select one trusts, the better off one is. This, however, should be carefully balanced with the need to not show ‘favoritism’ or ‘clubism’. That is, for the health of the conference and the perception of being broad and inclusive, it is important to select a broad review team that does not appear to ‘just’ be friends of the PC Co-Chairs.
- It is important to, every year, bring first-time members (especially junior faculty and post docs) onto the review team to help spread knowledge of how the ICSE review process is conducted and build community. At the same time, experience has shown that including too many new members on the review team in any given year can seriously derail the process.
- Experience is important, especially at the level of the Program Board. New PC members, thus, should have reviewed for other conferences and/or workshops before, and new Program Board members should have served at least once on the ICSE Program Committee within the past five years.

Two questions that inevitably come up in the selection of the review team include how many times a single person can serve on the team (in a row / in a given set of years) and whether there is a 'blacklist' of reviewers from past years who might have underperformed. To the first point, in order to avoid someone becoming 'too much of a regular', ICSE adopts the official SIGSOFT policy (<http://www.sigsoft.org/policies/pgmcommittee.html>):

For recurring events, at least one-third of the program committee membership should change. Most program committee members should not serve for more than three consecutive terms and in no instance should an individual serve more than four consecutive terms.

ICSE does not have a blacklist of reviewers, on purpose: should such a list ever become public, or should it become known that a certain person is on it, it could be disastrous and lead to major personal fallout. In addition, a person might have had an 'off year' and be a stellar reviewer again in future years. ICSE PC Co-Chairs, therefore, are strongly encouraged to converse with the PC Co-Chairs of previous years to identify review team members who underperformed. Yet, they should use their own judgment to actually decide whether or not to invite anyone, including people so identified.

Finally, some PC co-chairs like having their review team in place early (option 1 in the above table), whereas others prefer to await the performance of potential review team members in the review process of the ICSE immediately prior (option 2). Either way can, and has been shown, to work. Regardless of approach, however, PC Co-Chairs should be aware that it sometimes happens that a review team member has to renege on their commitment, often because of personal circumstances.

Common Challenge #2: Assigning Papers to the Review Team

One of the most time intensive and critical activities for PC Co-Chairs is to assign papers to members of the review team. Getting the right reviewers for each paper helps to create a much better experience for all involved. Clearly, a perfect assignment is impossible, and the experience of every pair of PC Co-Chairs has been that automatic assignments do not work. As a result, each team has spent somewhere between three and five full days to create the review assignment. A whole variety of methods have been used, from purely manual to taking an initial assignment from CyberChair or EasyChair and then looking at each paper and improving the set of reviewers manually. Regardless, PC Co-Chairs should budget significant time for this step of the process and have a plan in place before starting.

Several other tips from previous PC Co-Chairs include:

- It often helps to have assistance during the paper assignment process, for instance from a grad student who only checks for conflicts on proposed (changes in) assignments or from a staff person who records all of the decisions being made. It saves time: even a few seconds per paper, with 500 or more papers, creates significant savings.
- Past PC Co-Chairs have experienced a range of different behaviors when it comes to members of the review team:
 - Some reviewers bid high on exactly the number of papers they anticipate reviewing, hoping to get this precise assignment.
 - Some reviewers block inordinate amounts of papers, including many that easily fall within their expertise.

- Some consider their expertise to be ‘X’ only when a paper exactly matches their own, very specific research agenda, and bid accordingly.

As a result, most PC Co-Chairs discount somewhat the bids and expressed expertise and other preferences of reviewers. In the interest of making a fair, balanced assignment across reviewers and in the interest of giving each paper the best possible set of reviewers, it is common to break blocks, consider a reviewer an expert even if they do not so say, and other such factors.

- Desk rejection is an acceptable practice. ICSE receives its share of submissions that are clearly out of scope, incomplete, or of inferior quality. While the overhead on PC Co-Chairs is not insignificant in identifying these papers among the many that are submitted, it is usually fairly easy to spot such papers and to desk reject them with a simple explanation of why. Note that the length of a paper is not always an indication: quality results do not always need the full page limit to be described.

Common Challenge #3: Ensuring Review Quality

The quality of reviews produced by the review team is of utmost importance. Reviews and final decisions on acceptance/rejection should be fair to the submitted paper and its contribution, balanced in considering both strengths and weaknesses, and, perhaps above all, constructive. Long term, ICSE will benefit from nurturing a strong community of researchers and practitioners who choose to participate by submitting and attending. This requires careful messaging in all aspects of the review process. An important goal for ICSE is to help people improve in their work, research skills, and contributions. Providing authors with insightful, constructive feedback is essential to tending to a rising level of research overall.

Some reviewers see themselves as ‘gatekeepers of the exclusive quality of ICSE’. While this has an element of truth, given that ICSE seeks to produce high-quality proceedings each year, PC Co-Chairs should be mindful that the level of gatekeeping varies per PC member (some are overly zealous, others far too lenient). The review process exists precisely for this reason: PC Co-Chairs use it to establish a uniform quality level across which to judge all submissions.

Past PC Co-Chairs have voiced the following suggestions:

- Set review expectations clearly with the review team. These expectations should be communicated repeatedly, starting with the initial invitation to the review team and frequently thereafter (especially at the start of each new review phase). A particularly important aspect is to stress positive reviews and emphasize that the review team should be looking for reasons to accept papers, rather than to reject.
- Review each and every one of the submitted reviews personally, and flag any problematic reviews with the responsible Program Board and/or Program Committee member. Almost every PC Co-Chair, if not every PC Co-Chair, has kept an eye on reviews as they come in. Scanning the reviews this way enables flagging reviews that are too short, perhaps contain offensive language, do not follow review guidelines, and so on. Finding such reviews early, and communicating with the relevant reviewers, helps avoid further issues later on.
- Set a deadline to receive some portion of reviews halfway through the review period. Some reviewers ‘love’ holding on to all of their reviews until the very end. For some reviewers, this should not be seen as a problem, as they balance their reviews across one another. For others,

however, it may mean a full set of inferior reviews right at the deadline. To avoid the latter and be able to proactively manage review quality, some PC Co-Chairs have set a policy to receive half of the reviews halfway through the review period.

- Make sure that discussion summaries focus on the key reasons why a paper was accepted or rejected. Rather than summarizing individual reviews, or highlighting parts of those reviews, the discussion summary should reflect the consensus reasons that explain to the reviewer the decision by the review team.

Common Challenge #4: Managing the Review Process

With many individuals involved in the review process, management of the overall review process, and especially progress toward timely completion of the review process, is crucial. Suggestions from past PC Co-Chairs in this regard include:

- Build in slack in the review timeline (but do not necessarily let the review team know up front). Most PC Co-Chairs include up to week of slack in the due dates for reviews, discussions, and/or decisions. It is inevitable for a variety of circumstances that a few folks will need this extra time, and having it eases the stress on all involved.
- Communication with the review team is essential to having a smooth experience. This communication both involves setting of expectations and deadlines (usually communicated in the invitation to join the review team, at an in-person lunch-time meeting at ICSE n-1 where the entire review team is asked to join, and in subsequent e-mails), but also proactive management of possibly emerging issues. A friendly check in on how someone is doing if their reviews seem lagging, for instance, can help alleviate concerns, or clarify that the concerns are warranted and need to be addressed.
- Setting the example helps build a strong review culture. PC Co-Chairs in the past have, as one example, read the early set of returned reviews in detail and used excerpts from these reviews to communicate with (part of) the review team to illustrate good and/or bad review practices that one should watch for; this, then, led to the review team itself improving its practices and making suggestions to one another.
- Help individuals make difficult decisions. It is not uncommon for a review team member to become ill at a crucial time or otherwise face personal circumstances that may make it difficult to fulfill their reviewing duties. It is also not uncommon for the review team member to feel bad in such circumstances and wish to fulfill their obligations nonetheless, even though their ability to do so may be tenuous. In such cases, communication is key, but sometimes, gently letting a review team member 'off the hook' early on when one becomes aware of the circumstances might be preferred over waiting until later when the member cannot perform.
- It is inevitable that author/reviewer conflicts arise at a later time than the bidding and assignment process (even with, and perhaps especially so, with double blind reviewing). It is ok to re-assign a paper to another reviewer as a result. In general, past PC Co-Chairs have realized that a perfect load distribution among all reviewers is impossible, and that +1 or -1 reviews as compared to the average review load is an acceptable range and accommodates reassignments.

Common Challenge #5: Double Blind Reviewing

- *[The next Technical Review Committee Chair and the PC Co-Chairs who are introducing double blind should capture their insights here.]*

Appendix A. Latest ICSE Review Process

The latest ICSE review process is governed by the following primary decisions, as initially established for ICSE 2018 and agreed upon by the PC Co-Chairs for both ICSE 2019 and 2020.

- No limit exists in the number of papers an individual author may submit.
- The review process starts with the PC Co-Chairs considering any papers for desk rejection.
- All papers that are not desk-rejected receive equal treatment (to be fair to the authors):
 - Each paper receives at least three reviews.
 - Rebuttals are included and feed into the final stages of discussion of the Program Committee before the Program Board meeting (to enable authors to answer problematic reviews and/or clarify questions that have arisen).
 - Authors receive decision summaries that summarize the key factors leading to the decision of accept or reject.
- The review process is conducted by a Program Board with Program Committee (to deal with scale):
 - Program committee members are the primary reviewers of the papers; program board members read the papers but do not serve as reviewers themselves (to avoid significant 'sway' in argument, since Program Board members are present in the in-person meeting but Program Committee members are not).
 - The program board shall meet in person.
- The review process is double blind.
- The review process concludes with the PC Co-Chairs producing a review report that provides a summary and assessment of the entire effort. This report typically rests on personal thoughts and observations of the PC Co-Chairs, but also questionnaires of the Program Board, Program Committee, and authors.