

On Submissions and Resubmissions

Olivier Danvy

Department of Computer Science, Aarhus University
danvy@cs.au.dk

Influential papers can be re-submissions. For example, John Hughes and Koen Classen were awarded the most influential paper of ICFP 2000 about QuickCheck. To conclude his reception talk at ICFP 2010, John quipped “Not bad, for a paper that was rejected at [name of conference omitted].” (He also added that the reviewers made valuable suggestions and encouraged them to revise and re-submit.) For a second example, Mads Tofte and Jean-Pierre Talpin were awarded the most influential paper of POPL 1994 about region inference. In the course of his reception talk at POPL 2004, Mads said that the initial version of their submission was not accepted at POPL 1993 because it did not handle recursion. They were told to handle recursion and come back next year, which they successfully did.

At ICFP 2011, and with the approval of the ICFP steering committee, I ran the following experiment about resubmissions:

Authors of re-submissions can attach a confidential *revision history*, i.e., an annotated copy of the reviews of the previous submission(s) specifying how each point in each review was addressed. If one of the reviewers of a previous submission is asked to review the re-submission (and assuming that he has kept a copy of his review even if he has discarded the corresponding submission), this re-reviewer sends his original review to the program chair, who sends back the annotated copy. This way, the re-reviewer directly sees how the previous review was addressed, and can start reviewing the re-submission on an informed footing. Nobody else knows that a submission is a re-submission or is brand new before completing his own review. No new reviewer of a re-submission sees the annotated reviews of the initial submission.

The declaration that [the submission was brand new, or was the revision of a previous submission] was optional. There were 12 such declared resubmissions. One submission presented itself as brand new, but was subsequently identified as a resubmission (which was not counted against it). So among 112 submissions, at least 13 were resubmissions, i.e., 11%.

Enclosing the revision history was also optional. Among the 12 declared resubmissions, 5 had annotated copies of previous reviews, and 7 did not:

- among the 7 that did not, at least 3 were reviewed by a reviewer of the previous submission;
all 7 were accepted;
- among the 5 that did,
 - 3 reviewers identified themselves as a previous reviewer, and asked me for a copy of the previous review;
 - 1 (sub-)reviewer suggested the resubmission be rejected on the ground that it still had problems, which I spotted while going through this review; I then asked the reviewer for a copy of his previous review, and I sent him the corresponding annotated reviews; this sub-reviewer then realized that his main issue had actually been addressed. This made him

change his mind and suggest the submission be accepted – which it was.

Among these 5, 4 were accepted.

At the PC meeting, the reviewers appeared to be neutral about whether a submission was brand new or was a resubmission. All re-reviewers, however, were interested about how their initial review had been addressed.

Conclusion: This experiment at ICFP 2011 suggests that attaching a revision history to a resubmission has a positive effect:

- It forces the authors to spell out how they have addressed (or not) the previous reviews.
- It makes it possible for the re-reviewer to read – or not – the annotated copy of the previous review, so that (s)he has the choice of starting either on a fresh footing, or on an informed footing. And if the reviewer chooses not to read the annotated copy of his/her previous review, it is with the knowledge that the points of this previous review *were* considered (though of course not necessarily addressed) by the authors.
- Also, the time invested in writing a review is better spent, particularly if there is a concrete prospect of seeing the points of the review being addressed in the future.

Spin off:

One week after the notification, authors of accepted papers could be offered the option of uploading an annotated copy of their reviews to the submission site. The reviewers then would have a chance to quickly check how their points were received and addressed, and the opportunity to get back one more time to the authors.

This option stands midway between the current situation, where authors are left on their own after their submission is accepted, and shepherding, where a PC member supervises the finalization of the submission. It is a lightweight version of going through the final version of accepted papers: the reviewers may not have the time for that, but they may have the time to talk one last time about their review.

An open issue: Should a reviewer be sent the complete revision history, or just a copy of his/her own previous review?

There are pros and cons for each alternative: the reviewer should not be influenced by other reviews a priori, i.e., before writing their own review; on the other hand, it is not infrequent that after posting their own review, reviewers have access to the other reviews, so why not see the whole revision history; on the gripping hand, reviewers and re-reviewers do not always agree with each other, and these disagreements should not penalize the authors.

Depending on the alternative, the authors would compose their revision history either to be read as a whole, or to be read separately by the re-reviewers.

Acknowledgments: Thanks to Phil Wadler for suggesting this note be written and to Andrzej Filinski and Julia Lawall for comments.