

# What does an Associate Editor actually do?

Graham Cormode

G.Cormode@Warwick.ac.uk

## ABSTRACT

What does a Associate Editor (AE) of a journal actually do? The answer may be far from obvious. This article describes the steps that one AE follows in handling a submission. The aim is to shed light on the process, for the benefit of authors, reviewers, and other AEs.

## 1. INTRODUCTION

Journal publications are an important part of the propagation of results and ideas in computer science. Papers in prestigious journals reflect well on their authors, and serve to provide a full, detailed and peer-reviewed description of their research. Yet, the process from submission to decision is opaque. A researcher typically submits their paper to a journal and then waits months (sometimes many months) before receiving a set of reviews and a decision on whether the journal will pursue publication of the submission. It is far from obvious to the researcher exactly what is going on during this time.

The purpose of this article is to shed more light on this process, by describing the typical sequence of events from the perspective of the associate editor. The hope is that this serves multiple purposes:

- To help authors understand the process, and allow them to make their submissions with this knowledge.
- To help journal reviewers understand their role in the process, and how they can be most effective in helping to determine the right outcome for a submission.
- To help me (and, by extension, other associate editors) think of the process more clearly, and optimize our role within it.

The editorial structure of a journal varies between titles, but in general there is an editorial board which consists of an Editor-in-Chief (EiC) and multiple Associate Editors (AE). The role of this board is to determine which papers to accept for publication in the journal.

In general, the EiC receives new submissions and allocates these to AEs for handling through the review and decision process. The complete range of tasks performed by the EiC is not necessarily known to the AE: there are many “behind-the-scenes” tasks performed that they do not get to see<sup>1</sup>.

This article focuses on the role of the AE in the editorial process, in order to answer the question “What does an Associate Editor actually do?”. The answer is far from obvious: for example, one thing the AE does *not* typically do is “edit” papers in the popular sense of the word<sup>2</sup>. Rather, the AE’s main task is to make editorial recommendations to the EiC about what decision should be made on submitted papers.

To accomplish this, the AE has a seemingly simple set of responsibilities: to obtain referee reports for each paper they are assigned, and use these to make their recommendation for the paper, in a timely fashion. The execution of these tasks however requires quite a substantial amount of effort; moreover, this effort is concentrated in areas that might not be initially obvious. To explain this, I will describe the detailed sequence of steps that I follow between receiving a new assignment and providing my recommendation. A standard caveat applies: this description reflects my perspective and processes, informed by input from others (for example, [5]). Different AEs will no doubt have different approaches to the job. The author takes no responsibility for any loss, damage, or injury that may result from following any advice in this article.

**Outline.** In Sections 2 and 3, I outline the two main components of the AE’s job: initial handling and selection of reviewers for a paper (Section 2), and obtaining a decision for a paper (Section 3). In Section 4, I offer some suggestions for reviewers, authors, and associate editors in turn.

<sup>1</sup>In more blunt terms, I don’t fully know what the EiC does.

<sup>2</sup>The person who does make edits to accepted papers is the sub-editor, although in my experience this primarily involves the insertion or removal of commas.

## 2. SECURING REVIEWERS

**Step 0: pre-processing.** When a paper is submitted to a journal, it receives some attention before being assigned to an AE for handling. The EiC, and possibly an editorial assistant, will look over the paper. The general goal of this step is to check that the paper is suitable for further processing: Does it meet the formatting requirements? Is it generally on-topic for the journal? does it have a clear, novel technical contribution? Is it possible to open the files? Is it written in the language used by the journal? If the paper passes these checks, then the EiC will identify an AE to handle the paper, and assign it to them. The choice of which AE will handle the paper may depend on many factors: whether it falls within the AE's area of expertise, the relative workload of the AEs, avoiding potential conflicts of interest between the authors and the AE, and so on<sup>3</sup>.

In most journals, the paper is handled via a web-based manuscript system (with a generic sounding name like *ScholarCentral* or *ManuscriptOne*), which tends to enforce a particular workflow. The web-based manuscript system (WBMS) will generate email alerts to each participant when they have a task to perform. So when a paper is assigned to me, the WBMS will generate an email message telling me that I have work to do.

**You've got email.** My process on receiving a new paper to handle is as follows: I first sigh<sup>4</sup>, realizing that this means more work to do. Then I am overcome with excitement about the prospect of guiding a fresh paper through the journal submission process.

I next take a print out of the main paper and any cover letter. As soon as possible, I run a hot bath, and immerse myself in the water and in the paper<sup>5</sup>. I then read the paper to get an idea of what it is about, roughly what techniques it is using, and what papers are most relevant to the work in hand.

My objective in this phase of the process is to identify a set of researchers to contact and ask them to provide a review of the submission. As such, my approach is quite different to when I am reviewing a paper myself. As an AE, I do not find it necessary to comprehend every last detail of the paper, or even to grasp all of the ideas presented. Rather, my goal is to find experts who can understand the paper in detail, and provide commentary on its significance and novelty. Consequently, I try to avoid forming a strong opinion about whether the submission should be accepted: the bulk of that work will

<sup>3</sup>I suspect that a whole new article could be written about the job of the EiC, and I would encourage someone to do so.

<sup>4</sup>Or, according to taste, shriek, cry out, rend my clothing, or ask "Why me?"

<sup>5</sup>People often ask me why I read papers in the bath. I patiently explain that it would be hopeless to try to do this in the shower.

be on the reviewers. However, based on my initial reading of the paper, I will have a sense of the general level of the paper.

Sometimes it is clear that the paper does not meet the standards of the journal. In such cases, an AE may provide an "administrative reject" decision (also known as a "desk reject"). I do this when I am certain that the paper stands almost no chance of eventually being accepted. In particular, I want to be able to provide the authors with a supportable reason for the reject decision and feedback that they can make use of. Reasons I consider suitable to motivate an administrative reject include if the submission is presented so badly it is impossible to understand any of what is being said; if the results very clearly duplicate prior work; if the topic of the paper seems very much out of scope for the journal; or if the submission includes text that appear in other previously published papers and thus violates the journal's plagiarism policy. In my experience, submissions meeting any of these criteria are not common, perhaps because the EiC catches them before they are assigned to an AE.

There are still some papers which I believe are borderline for the journal, but which do not match any of the above conditions. In these cases, I can invite reviewers to review the paper, even though I think its prospects are poor. It is better to allow a seemingly poor paper a fair chance with expert reviewers, than for an AE who is not an expert in its area to deny it any chance. This gives the authors of the paper a fuller set of reviews, which is hopefully of use to them. The tradeoff is that I am asking reviewers to give their time to review what may be a poor paper. My rationale is that reviewing is part of the service we owe the community in return for submitting our own papers, and we cannot always expect high-quality papers to read. Moreover, it should be a relatively quick task for an expert if the submission is indeed of low quality to make an assessment and to prepare a short review highlighting the deficiencies. I can invite fewer reviewers (say, two), if I think that there is a good chance that they will both provide negative reviews.

As a third option, I sometimes desk reject based on a fixable issue, such as problems with figures or formatting. In the feedback to authors, I let them know that it is permissible to resubmit a corrected version of the paper. I also indicate that I believe that such a revision is unlikely to meet the high standards of the journal. This leaves the door open for the authors to resubmit, while indicating heavily that they would do well to reconsider their choice of venue.

**Picking Reviewers.** After getting a sense of the paper, my next step is to identify a set of potential reviewers to invite. I think about the paper as I understand it, and

which researchers are active in that area or related areas. I cast my mind over papers I have read, presentations I have seen, and conversations I have had to identify who is suitably expert on the topic. There doesn't have to be an exact match – perhaps the application is unusual, but a reviewer has used similar techniques.

I also draw ideas for reviewers from the paper. Does the paper make extensive reference to some prior work? Does it compare to a method described in a previous paper? Then there is a good chance that I will invite the authors of these papers (assuming that they do not overlap with the authors of the current submission) to perform the review. I may do some speculative searching – are there keywords or problem descriptions from the paper that I can find other papers about online? In particular, can I find papers on similar topics published in the same journal – since I feel the authors of those works owe a review back to the journal.

After brainstorming for a while, I usually have a list of half a dozen potential reviewers. I do some additional research on them to ensure that they are well-placed to help. Before inviting each reviewer, I check their homepage and their entry on DBLP. I look at the titles and venues of their papers, and years in which they have been active in this area, and also descriptions of their current role and activities.

*Other commitments.* I tend to avoid asking people who indicate that they are the head of a large research group, chair of their department and active in running a start-up at the same time. Such people tend to be too busy to perform reviewing tasks<sup>6</sup>. Advanced graduate students can be a good fit because they know their focus area very well, and have very few other pressing demands on their time<sup>7</sup>; however, it is sometimes hard to tell which students are mature enough in their area without a personal recommendation. So the bulk of reviewing falls upon faculty and researchers who don't appear too busy, or don't yet realize how busy they are.

I avoid asking EiCs and AEs of any journal to perform a review: they are usually far too occupied with the submissions for their own journal. In particular, I avoid asking an AE from the same journal to assist<sup>8</sup>.

*Still Active?* The editor's curse is to find someone who has worked on some highly related topics, only to discover that their last publication was in 1999. Usually this means that they have left research for another career,

---

<sup>6</sup>They often appear to be too busy even to respond to review requests.

<sup>7</sup>Graduation can wait.

<sup>8</sup>I hope they realize that this is why I turn down their corresponding review requests. Ideally, the EiC would always assign the paper to the most expert AE on that topic. However, I have gradually come to realize that EiCs are less omniscient than one might at first imagine them to be.

retired, or abandoned this area of study<sup>9</sup>. In some cases, I identify a reviewer who would be perfect to help with a paper, only to discover that they are no longer alive, which I find most inconsiderate.

Following this analysis of reviewers, I pick a shortlist of 3 or 4, and start to send out invitations. The WBMSS typically has a default invitation template describing the expectations. I personalize this invitation, to give some indication of why I have invited the reviewer: for example, because I think the submission relates to their expertise on a topic, or because it compares to their system, for example. My hope is that this personal touch will make them more likely to accept the invitation. The invitation can also indicate if the paper is a resubmission, an invited submission or an extended version of a conference paper.

I might include the submitted manuscript with the invitation. When I am invited to review, I often find it helpful to quickly scan the submission, to determine how relevant it is and much effort it will be. When suitable, I like to give other reviewers this opportunity. However, I must admit, when a paper seems particularly long and technically dense, I may avoid sending it, for fear of scaring off the potential reviewer.

**Dealing with rejection.** Inevitably, some invitations to review will be met with rejection. Indeed, in my experience about half of responses are negative. This can be for many reasons, of varying validity: the invitee is too busy, does not consider themselves an expert on the subject matter, does not find the paper interesting, or just doesn't feel like it on the day. A negative response does not annoy me (unless I feel that the paper really was spot-on for the reviewer). What does irk me are two things:

*Tardiness* – it should not take a long time to respond to a review request. If people are actively at work, I would hope to hear a reply within a couple of days; if traveling or otherwise tied-up, I would still hope to hear within a week or so<sup>10</sup>. It pains me when an invitee sits on a request for weeks, and then declines (possibly only after a reminder). Even when the invitation is accepted after a long pause, this can be a troubling sign, as it indicates that the review itself may be similarly delayed.

*Lack of alternative suggestions* – my favourite type of response is actually a very fast negative response that comes with a list of suggested alternate reviewers. This means that the invitee has thought about the invitation, understands that they are unable to commit to it, but has

---

<sup>9</sup>One does not like to name names, but on multiple occasions I have had papers which refer heavily to the work of S. Brin and L. Page. However, these two stopped publishing in the 1990's, and have not responded to any of my requests for reviewing. I can only assume that these promising researchers have given up on academia, and followed a less rewarding career in industry.

<sup>10</sup>Everyone checks their email while on vacation, right?

considered it enough to come up with a list of others who may be able to help. This is particularly valuable when the area of the submission is less familiar to me. As a reviewer, I suggest alternates when I am unable to assist – unless I really don't know the topic. As a result, when an subject matter expert declines, I often follow up with an email pushing for some suggested alternate reviewers. I encourage people to feel obliged to provide alternates when declining an invitation.

**Adding more reviewers.** When reviewers decline an invitation, I need to find more reviewers to invite. Sometimes I have some back-ups already picked, or can take advantage of suggestions from those who have declined. I avoid having more than four “active” invitations at one time, in case all reviewers accept: it is redundant to have a large number for one paper. Often though, I need to find some new candidates. This is perhaps the toughest part of the job, as it means further head scratching to come up with good candidates. It is quite dispiriting when a large number of reviewers have declined to review a paper. The worst case is when the paper is quite specialized, and all the natural candidates have been tapped. It is particularly galling when, after prompting for other reviewers, the suggestions consist of candidates who have already declined. At this stage, the AE can feel that the task of finding enough suitable experts to evaluate a paper may be impossible. However, with persistence, enough reviewers will eventually agree.

**Reaching Acceptance.** When sufficient reviewers have agreed to review a paper (usually three or four), and dates for the review have been agreed, the initial phase of the process is complete. I can sit back, relax, and wait for the reviews to arrive.

### 3. GETTING TO A DECISION

**The whooshing sound they make as they fly by.** When I first started working as an AE, I imagined that the bulk of the effort was in weighing up the reviews for a paper, and synthesizing these to come up with a careful, considered decision and rationale for it. This a much less significant part of my work than I had thought. Indeed, it seems that much of the effort of the AE is in reminding, cajoling and threatening reviewers who have agreed to provide a review, but who fail to fulfil their promises.

In the ideal situation, reviewers will perform their task within the allotted time (typically, six weeks to a few months), and deliver a carefully thought-out, clearly expressed review. Indeed, most reviewers do an excellent job in this regard, and I am truly grateful to them. However, there are many cases where things do not follow this outline, and more active involvement is required.

The WBMSS usually includes a “due date” for each review (which can be set by the AE), and may automatically remind the reviewer as the deadline approaches and is passed. In addition, around the time of the deadline, I send a personalized reminder, as this is harder to ignore than an automated message. I do not keep detailed statistics, but while many reviews are received on time, it is a sad fact that a large fraction are late. A little tardiness is forgivable, but after more than a week, it starts to become a problem. Many journals strive to have a rapid turnaround time for submissions, and delayed reviews are the biggest obstacle to achieving this goal [4, 5].

Checking this requires more of my attention. I have to keep an eye on which reviews are late, and send reminders to reviewers, requesting that they make good on their promise, and deliver their review. The pressures that I can bring to bear are limited: I can send increasingly plaintive requests, or express my displeasure or anguish at the continued delay; I can try to provoke guilt or regret in the reviewer; but there are few direct actions I can take against the tardy reviewer. Persistence is my only weapon. In a few cases I have given up on receiving a review when the other reviews received were sufficient to reach a decision.

**The reviews are in.** When I do receive a review, I read it carefully, and check that there are not any obvious problems with it. Problems in reviews are rare, but occasionally it may be clear that the reviewer's standards are not calibrated for the venue (too harsh, or too lenient); or that the recommendation does not align with the content of review (e.g. many major flaws highlighted, but an “accept” recommendation). Reviews can sometimes be improved by clarifying what is expected from a revision, and ensuring that the discussion is as objective as possible. The AE can ask a reviewer to revise or elaborate their review. Very rarely, there may be inconsistencies across reviews that are resolved by an (email) discussion with the AE in the middle.

**The Big Decision.** When there are sufficient reviews for a paper, I can make a decision. The typical number is three, but more or fewer is possible. I am happy to recommend rejection for a paper on the basis of two reviews which agree on this outcome, or even one in extreme cases. For a positive recommendation, I prefer to have received three reviews, even if they are not unanimous. Collecting four reviews is reasonable (and acts as insurance against one reviewer going awol); more than four is unusual except for very selective journals.

I usually find it fairly swift to make a decision: reviews often agree on the general level of quality and interest in a submission. Some normalization is needed based on the standards of the journal, but in general it is quick to weigh the comments and scores of the re-

viewers, and reach a consensus. The process is guided almost exclusively by the reviews—my opinions of the paper carry almost no weight at this point<sup>11</sup>. The first decision is a binary one: Is there any prospect of publishing this paper in the journal? Does it show enough potential and interest? If not, then the recommendation is to “reject” the paper. This recommendation is accompanied by a justification, summarizing the reasons for rejection: I identify the main reasons from reviews that led to the decision. It may include more or less encouragement to submit to another venue, especially if the submission was ultimately judged out of scope or below threshold for my journal. The authors may appeal a reject decision, either to the AE or the EiC, but without evidence of serious unfairness this is unlikely to alter the outcome. A rejected paper is sometimes resubmitted to the same journal, after some revisions. Most journals will try to catch this, and either reject automatically, or assign it to the same AE to handle.

If the paper is not rejected, there are three possible recommendations: “accept (as is)”, “minor revision”, and “major revision”. The exact semantics of this vary depending on the journal, but as a rough guide, a major revision will be returned to the same reviewers to get their opinion on the new version; a minor revision will be scrutinized by the AE; and an accept will move straight into the publication queue. However, the AE has a lot of leeway: a minor revision may be sent out to reviewers; and a major revision may be sent only to a subset of reviewers, or new reviewers may be added. I won’t spell out all the situations that can arise, but the underlying issue is the same: before giving an “accept (as is)” decision, I want to be certain that the paper represents a sufficient contribution for publication in the journal. When the reviews indicate some notable questions or concerns, I want to be assured that these are suitably addressed before recommending the paper for publication. Sometimes I can do this myself (based on the revised submission, and any cover letter or list of revisions, and comparing these to the original reviews); or I may seek the opinion of the original reviewers on such questions.

**Recommendations and Decisions.** You may notice that an AE makes a “recommendation”, not a “decision”. This is deliberate terminology: it is the EiC who makes the decision, not the AE, who merely recommends an outcome. I will let you into a secret: I have not encountered cases where the EiC’s decision did not follow the recommendation of the AE, although this does happen. I find that this is a useful way of thinking about

<sup>11</sup>Occasionally, an AE may enter their own review for a paper they are handling on a topic are familiar with, especially if the invited reviewers have not done a timely job. Then this review is weighed up with the others.

the process. It reminds me that I have to justify my recommendation both to the authors and to the EiC; I am not making decisions at my whim. Once I submit my recommendation on a paper to the EiC, I can again sit back: my work – for now – is done.

**Revisions.** For revisions, the process starts over again – selecting reviewers, obtaining reviews, and making a recommendation. Typically, one invites the same set of reviewers, although there is the option to add new reviewers (if additional input is needed), or drop some (for example, if they were entirely satisfied with the previous version). There can be multiple rounds of revision, but if major issues remain after a first revision, it is common to move towards a reject. Once a reject or accept is reached, the AE’s involvement with the paper is concluded.

## 4. RECOMMENDATIONS

Based on this description of the process, I have a number of recommendations and requests for those involved in the journal review process:

### 4.1 Recommendations to authors.

It is easy to imagine that a journal will immediately recognize the novelty and importance of a submitted paper, and that the editors will quickly identify experts who can judge the merits of the submission. However, the reality is perhaps less ideal: there is no guarantee that the EiC will be able to match the paper to the best AE for the paper, or that the assigned AE will be able to identify and secure the most expert reviewers. Authors can help this process along:

**Suggest suitable Associate Editors.** It is often appropriate to suggest an AE to handle the paper. Take a look at the editorial board, and see which AEs have familiarity with the area. The suggestion usually can be communicated to the EiC as part of the cover letter, or within the WBMS.

**Suggest suitable Reviewers.** Before my experience as an AE, I did not think it was necessary to suggest reviewers: the journal staff should easily be able to identify an expert set of reviewers. Proffering suggestions seemed to imply that the nominees were my cronies. Now I realize that it is very valuable to suggest reviewers: there is no guarantee that the AE will be a leading expert in the domain of the paper, and I find that reviewer suggestions are useful input to me as an AE. I carefully evaluate suggested reviewers, and only follow up if it is clear that they are suited for the paper, and do not have conflicts of interest with the authors<sup>12</sup>. I tend to invite only one or two suggested reviewers, and fill out

<sup>12</sup>In particular, it is important to avoid inviting the authors to review their own paper, which is not unprecedented

the rest of the panel with “independent” reviewers, to avoid any issue of bias. Authors should realize that their suggestions may not accept the invitation, and there is little value in suggesting a “big-name” researcher who is too busy. Lastly, some journals also allow authors to indicate “non-preferred” reviewers. I can think of few situations where this is of use to authors, and it seems that there should be some clearly articulated explanation.

**Think about your citations.** Think carefully about which works you cite, and whether there are any important references missing. An AE will often look to the bibliography for potential reviewers to invite. So authors should realize that their bibliography is another list of “suggested reviewers”. They should also reflect on how fairly they describe and compare to prior work, since the authors of those works may be called upon to judge the submission.

**Optimize your revisions.** As noted above, the revision will be handled by the same AE as the original submission, and will typically be read by the same reviewers. It is therefore sensible to optimize the revision accordingly. Make a cover document containing each review, and indicate how you respond to each point: what changes were made, and where. It is OK to disagree with a reviewer comment, so long as you explain why. It is also helpful to indicate which sections have changed in the paper, via highlighting<sup>13</sup>. This takes extra work, but this type of effort can make the review process go much more smoothly, and hence speed the paper to publication.

## 4.2 Recommendations to Reviewers.

These are perhaps less recommendations than pleas:

**Respond swiftly and decisively to requests.** As an AE, my goal is to provide well-informed decisions to authors in a timely fashion. This starts with responding to the initial review request. Please don’t sit on a review request for weeks: it is usually only the work of a few moments to determine one’s current level of commitments, and availability to accept a new task. As noted above, a swift response is often appreciated, even if it is negative. Please also provide alternate reviewer suggestions as a matter of course. Often, I receive a request and I think “Why are they asking me? Why don’t they ask X?”. The reason may be that the AE does not know that X is the expert on this topic – so please inform them of this! You can also use declining a review request as an opportunity

<http://barcorefblog.blogspot.com/2012/10/fake-peer-reviews.html>

<sup>13</sup>This has the advantage that it will focus the attention of the reviewers on just those parts of the paper; otherwise, they may re-read the whole paper, and come up with additional comments and things to change.

to advance the career of a more junior member of your community, by suggesting someone less well-known.

**Honour your commitments.** When you accept to perform a review, you are making a commitment to deliver the review by the date agreed. This commitment should be taken seriously. It is easy to devalue the importance of review work – after all, it is “voluntary” work. However, I view reviewing as an obligation: when we submit papers, we expect them to receive appropriate and timely reviews, and so we should perform reviews similarly. It is tempting to think of reviews as less important than the many other demands on our time, (our own research, teaching, and funding deadlines) and allow the review to get progressively later and later. But this is quite unprofessional. It delays the process for authors, who need to get timely decisions in order to publish their work and progress their careers.

It goes without saying that you should do a good, careful job in reviewing the paper. For guidance on this, there are several good articles on the topic [3, 6, 1]<sup>14</sup>.

You should *always* accept a request to review a revision of a paper. The work involved should be much less than to perform an initial review (especially if the authors have suitably optimized their revision). If you asked for changes, then you should at least look at the response.

**Accept a reasonable number of requests.** It is hard to load-balance incoming review requests: sometimes, many arrive in close proximity. However, as indicated above, it is important to be an active participant in the review process, and do your fair share. One heuristic is to perform 3 – 4 reviews for each submission you make (assuming that each paper does have multiple authors), but more senior people may need to do more.

Be aware that a journal review brings different expectations to a conference review. A journal review is expected to be in greater depth, and to more carefully scrutinize the whole paper. Consequently, the review should attempt to evaluate the paper in full, or be explicit about which sections could not be verified. Journal papers may also be (much) longer than a typical conference submission, so one to several months is allotted to perform the review – do not interpret this as permission to leave the review to the last minute.

## 4.3 Recommendations to Associate Editors.

The above discussion has outlined the workflow I tend to follow in handling a paper. Implicit in this are several recommendations and considerations:

**Be considerate of authors.** Your goal as an AE is to oversee a fair and timely handling of submissions to your journal. So try to ensure that each submission has

<sup>14</sup>As well as some that are laughable, e.g. [2].

a fair chance, by identifying and inviting suitable reviewers, and using these to make good decisions on papers. In some cases, the most considerate thing to do is to swiftly reject a paper, rather than enter it into a lengthy review process, taking up reviewers' effort, and ultimately reaching the same outcome.

**Be considerate of reviewers.** Try to identify reviewers who are suited to the paper, and try to avoid asking the same reviewers to help with a lot of papers. Be understanding when reviewers need more time to review a paper, while firmly reminding them of their obligation. Remember that reviewing is a mark of service to the community, and an indication of the esteem with which the opinion of the reviewer is held, so be sure to allow junior researchers the opportunity to participate in the review process. This can also be a learning opportunity for them to see firsthand how peer review works in practice, and to calibrate their opinions against the reviews of others.

**Be considerate of yourself.** When I started as an AE, I had high aspirations: I would read each paper in detail, and provide my own review and comments in addition to those of the invited reviewers. This lasted for exactly one paper. For journals with high throughput, you may handle 20-30 papers per year, on a wide variety of topics, and it simply is not practical, nor a good use of your time, to try to do too much. Stick to the core tasks, and you will be doing the community a service.

By way of guidance, here are my estimated times for handling a submission. Of course, these can vary: an obviously unsuitable paper may be faster to handle.

Read and think about paper: 1-2 hours
Search for and invite initial reviewers: 1 hour
Handle review responses, and find replacement reviewers: 1-2 hours.
Receive and process reviews: 0.5 hours total
Chasing reviewers to deliver their reviews: 1 hour
Re-visit paper, and formulate recommendation: 1 hour

## 5. CONCLUDING REMARKS

This is the end of what I have to say.

**Acknowledgments.** I thank Jian Pei for many helpful comments and suggestions.

## 6. REFERENCES

- [1] Mark Allman. Thoughts on reviewing. *ACM SIGCOMM Computer Communication Review (CCR)*, 38(2), April 2008.
- [2] Graham Cormode. How not to review a paper: The tools and techniques of the adversarial reviewer. *SIGMOD Record*, 37(4):100–104, December 2008.
- [3] Ian Parberry. A guide for new referees in theoretical computer science. *Information and Computation*, 112(1):96–116, 1994.
- [4] Richard Snodgrass. CMM and TODS. *SIGMOD Record (ACM Special Interest Group on Management of Data)*, 34(3):114–117, September 2005.
- [5] Richard T. Snodgrass. ACM TODS associate editor manual. <http://tods.acm.org/editors/manualFeb2007.pdf>, January 2007.
- [6] Toby Walsh. How to write a review. <http://www.cse.unsw.edu.au/~tw/review.ppt>, 2001.