

“Identify the Champion”

An Organisational Pattern Language for Programme Committees¹

Oscar Nierstrasz

Software Composition Group, University of Berne²

Abstract. The peer review process for technical contributions to conferences in computing sciences is very thorough, and can be as stringent as the review process for journal publications in other domains. The programme committee for such a conference will typically convene at a meeting, where submitted papers are discussed, and accepted or rejected for presentation at the conference. Experience shows that discussions are more focussed, and the entire process runs more smoothly if most of the time is devoted to those papers that are actually “championed” by some committee member. In order to make this work effectively, however, the notion of “championing” must be introduced early in the review process. This paper presents a set of process patterns that help to achieve this goal.

Introduction

Programme committees (PCs) for many computing science conferences adhere to standards that are more common for journal publications in other domains. A great deal of time and effort is invested in reviewing submitted papers. In most cases, the PC meets physically for one or two days, thus entailing considerable expense. Occasionally “virtual meetings” are held by telephone or e-mail, but the dynamics of the meetings are essentially the same: each submission must be judged on its own merits, and accepted only if there is general agreement that the paper meets the standards of the conference.

Clearly, in order for a paper to be accepted, it helps if there is some PC Member who “champions” it at the meeting. Although the idea of championing seems to be central, it is seldom formally incorporated into the review process. When it is, experience shows that it has a dramatic effect in focussing attention on the key issue — i.e., accepting the best papers — and making meetings much more effective. In order for this to work well, however, it is important to instil the idea of “championing” throughout the review process. We present here a small pattern language³ that captures successful practice in several such conference review processes.

1. Copyright 1999 Oscar Nierstrasz. All rights reserved. “Identify the Champion,” in *Pattern Languages of Program Design 4*, N. Harrison, B. Foote, H. Rohnert (Ed.), vol. 4, Addison Wesley, 2000, pp. 539-556. Permission is granted to copy for non-commercial purposes. Available as HTML or PDF from: <http://www.iam.unibe.ch/~oscar/Champion/>

2. *Author’s address:* Institut für Informatik (IAM), Universität Bern, Neubrückstrasse 10, CH-3012 Berne, Switzerland. *E-mail:* oscar@iam.unibe.ch. *WWW:* <http://www.iam.unibe.ch/~oscar/>

3. It is a “language” (as opposed to a “catalogue”) in the sense that one should combine the individual patterns to form a complete solution to the problem addressed.

Target audience

This pattern language is targeted first of all to PC Chairs who must define the review process and chair the meeting, and second of all to PC Members, who need to understand the key problems and issues in the process. Finally, the language may be of interest to authors who are submitting papers to such conferences.

Overview

Pattern	Intent
<i>Identify the Champion</i>	Make the paper review and selection process for a scientific conference more efficient by focusing programme committee members' attention on whether or not they will "champion" a submitted paper during the programme committee meeting. (This pattern is a prerequisite for the others.)
<i>Experts Review Papers</i>	Distribute papers to PC members who are likely to be competent champions. (These are two competing patterns.)
<i>Champions Review Papers</i>	
<i>Make Champions Explicit</i>	Make PC Members aware of championing before they start reviewing papers.
<i>Identify the Conflicts</i>	Analyse the review results and group papers into interesting categories.
<i>Identify Missing Champions</i>	Detect and avoid archetypical problems that arise during PC meetings.
<i>Champions Speak First</i>	Maintain order during the PC meeting.
<i>Consensus on PC Papers</i>	Handle the special case of papers submitted by PC Members.

Pattern: *Identify the Champion*

Context

You are the PC Chair for a computing science conference with high scientific standards. You should define a review and selection process for selecting the best papers from those submitted.

Most review processes are superficially the same, but the details can be quite different from conference to conference.

The PC Chair collects submissions and distributes them to individual PC Members for review. PC Members may return papers to the PC Chair in case of conflict of interest, or lack of expertise. Each submission is typically reviewed by 3 or 4 PC Members. Review forms are collected and sorted, and submissions are ranked prior to a PC meeting where the papers are discussed and either accepted or rejected for presentation at the conference.

A PC meeting typically lasts 1-2 days, during which, for example, 20-30 papers are selected for presentation at a conference from a much larger number of submissions, for example, 50-200. For each paper, the PC Chair invites the PC to present arguments in favour of and against acceptance. In principle all papers are discussed, but often the papers with the lowest scores are summarily rejected, and are only discussed if someone explicitly requests it. Anonymous extracts of the review forms are typically returned to all authors to help them improve their papers, regardless of whether they were accepted or not.

Problem

How should you design the review process so that the PC meeting will succeed in selecting the best papers?

Forces

There is a large number of forces at play. Here we list some of the most important ones. Some forces that are specific to only certain phases of the review process are listed under the other patterns that follow.

- Each PC Member should have roughly the same number of papers to review, and papers should be reviewed by domain experts.
- It is hard to assign an objective score to a paper, since numerical ratings, or qualifications like “strong accept” or “good paper” will be interpreted differently by each reviewer.
- PC Members have limited time to devote to evaluating 15-20 papers.
- PC meetings are expensive to organize.
- It is hard to know in advance which papers will generate the most discussion.
- Without strong guidelines, a PC meeting can quickly degenerate into chaos.

Solution

Make the paper review and selection process for a scientific conference more efficient by focusing programme committee members’ attention on whether or not they will “champion” a submitted paper during the programme committee meeting.

Be sure to distribute the papers to PC Members who are likely to champion them. Organize the review forms, the ranking and sorting of reviews, detection of conflicts, and the review meeting itself around the identification of champions. Use rating schemes with explicit operational meaning, such as “I will champion this paper”, rather than implicit, subjective meaning, such as “strong accept”, or “better than average”, or “5”. Group papers around presence or absence of champions and detractors rather than ranking them by weighted scores. Drive all discussions and decisions by identifying the champion.

Rationale

In practice, only those papers that are successfully championed by some PC Member present at the meeting will be accepted. Since this happens anyway, the entire review and selection process will be much more efficient and effective if championing is made explicit in the process.

Identifying the champion forces PC Members to focus on their behaviour during the PC meeting rather than on their subjective impressions while reading the paper.

Examples

The notion of championing is commonly applied during PC meetings to keep discussions focussed, but it is rarely made explicit in the rest of the process. Some recent PCs that have made championing explicit include ICSE 98,¹ ECOOP 98,² OOPSLA 98³, FSE 6⁴ and PLoP98.⁵

Related Patterns

Jim Coplien [3] describes several related patterns drawn from various OOPSLA PCs, including *AssigningPapersToReviewers* [5], *SortedPaperList* [2] and *PaperChampion* [4]. The earlier

1. International Conference on Software Engineering.
2. European Conference on Object-Oriented Programming.
3. Object-Oriented Programming Systems, Languages and Applications.
4. Foundations of Software Engineering.
5. Pattern Languages of Programs.

patterns, however, do not explain how the notion of championing can really drive the whole review process.

Dick Kemmerer, the PC Chair of ICSE 98 [personal email communication] points out a variant of *Identify the Champion* that has some distinct drawbacks:

Several weeks before the PC meeting I told the PC that we would be following a procedure close to what you described in the paper, and gave them all a pointer to the paper, asking them to read it before the meeting. I immediately got an objection from one pc member who said he did not want to use the "champion" approach. ... I mention this because there are folks out there that have a preconceived notion of the "champion" approach. At these PC meetings a paper would not be accepted *unless someone was willing to put their name on it* (i.e., their name appeared on the paper as "recommended by"). ... The problem seems to be that people are less likely to want their name to appear in print as being the endorser than to be a champion as per your approach.

Shepherding is a pattern of guiding a paper (or rather, its authors) through rough terrain so that it can reach ground where it can be truly championed without reservations. Some conferences (like PLoP) make heavy use of shepherding, whereas others avoid it.

A variant is *Conditional Acceptance* in which the PC Chair exercises discretion over publication of the final version of the paper. Yet another variant is *Mentoring*, in which authors who seek advice in preparing a paper for submission may be assigned a "mentor," who is usually a PC Member.

Write to the Program Committee, [1] is a pattern that authors can apply to increase the chances that their papers will be accepted. The idea is to write in such a way as to win over a potential champion by catching his or her attention and providing good ammunition to argue for acceptance during the PC meeting. A good understanding of the dynamics of PC meetings helps in applying this pattern. A related tactic is to try and identify specific PC Members who are likely to review the paper, and write in such a way as to win them over as champions. (It almost goes without saying that this tactic can easily backfire!)

Pattern: Experts Review Papers

Context

You are a PC Chair, and have decided to apply *Identify the Champion*. You are expecting 100-200 submissions, each of which should be evaluated by 3-4 PC Members.

Problem

How do you distribute papers to the PC?

Forces

- Papers should be reviewed by someone competent to evaluate their contribution.
- It can be hard to guess who is the best person to review a paper.
- The most convincing champion for a paper is a domain expert.

Solution

Try to match papers to domain experts in the PC.

One easy way to get a rough match is to first get all PC Members to identify the domains and keywords that correspond to their main areas of expertise and their secondary areas of interest. Then, in the Call For Contributions (CFC), ask Authors to explicitly state which of these apply to their submissions. When submissions arrive, use this information to match papers to PC Members.

It helps to scan each paper to get a feeling who should review it. If related work by some PC Member is explicitly mentioned in the paper or the references, then that may be a good person to review it *unless there is a conflict of interest* (i.e., the PC Member is a co-author or colleague!).

Rationale

An expert who likes a paper is typically more willing to champion it than a non-expert.

Resulting Context

If an expert is less than enthusiastic about a paper, that is likely to kill its chances for acceptance, especially if the positive reviews are from PC Members who are competent in the domain, but not experts.

Pattern: *Champions Review Papers*

Context

You are a PC Chair, and have decided to apply *Identify the Champion*. You need to distribute submitted papers to PC Members who are competent to evaluate them.

Problem

How do you distribute papers in such a way as to maximize each paper's chance that it will find its champion?

Forces

- Papers should be reviewed by someone competent to evaluate their contribution.
- It can be hard to guess who is the best person to review a paper.
- The most convincing champion for a paper is a domain expert.

Solution

Let PC Members pick the papers they want to review.

One (not very good) way to do this is to distribute a list of titles and authors to all PC Members and have them select which papers (i) they would like to review (i.e., they think they might champion) (ii) they feel competent to review, and (iii) they do *not* want to review (either because they do not feel competent, or because they have a conflict of interest). There are two problems with this approach: (i) PC Members will have too little information to go on, and (ii) it will cost at least a week to get the responses back (if you're lucky).

A better way is to request Authors in the CFC to *pre-register* their intent to submit a paper, at least a week in advance of the paper deadline, by sending the PC Chair an e-mail containing the title, authors, contact author's coordinates, an abstract, and keywords (see *Experts Review Papers*). After the pre-registration deadline has passed, send the list of abstracts electronically to the entire PC, and ask them to categorize papers as above based on this information. When the paper submission deadline passes and all papers are in, papers should already be assigned to PC Members, and can immediately go out.

Rationale

Asking PC Members to "bid" for papers to review reinforces championing.

Examples

Many conferences, ICSE in particular, traditionally use a bidding procedure.

Related Patterns

Assigning Papers To Reviewers [5] discusses some of the main ideas behind both *Experts Review Papers* and *Champions Review Papers*. The key difference is that the patterns presented here focus on giving papers the best chance of being championed.

Pattern: Make Champions Explicit

Context

You are a PC Chair who plans to apply *Identify the Champion*. You must design review forms that will help you prepare a successful meeting.

Problem

How can you make it easy to tell which papers will be championed in advance of the meeting?

Forces

- Reviews are necessarily subjective, since numerical ratings, or qualifications like "strong accept" or "good paper" will be interpreted differently by each reviewer.
- A PC Member who gives a paper a high score is not always willing to champion it.
- Just as surprisingly, PC Members who give mediocre scores to paper sometimes turn out to be strong champions.
- Review forms with lots of numerical grades (i.e., for originality, soundness, readability, etc.) are hard to convert into a single, meaningful score.
- Review forms that are too complex annoy reviewers, and much of the information they record is typically ignored during the PC meeting.
- PC Members often delegate papers to "subreviewers," who may not be fully aware of the details of the review and selection process.

Solution

Ask PC Members explicitly on the review form whether they intend to champion the paper.

It is very tempting to introduce very fine-grained scales of appreciation on review forms, such as scales from 1-10 for various criteria, including originality, soundness, presentation, etc. These kinds of rating typically have the opposite of the desired effect, namely they waste the reviewers' time and they make it more difficult to tell who is willing to champion a paper.

In practice the only ratings that are really critical for the PC meeting are (i) a score indicating whether the paper should be accepted, and (ii) a score indicating the reviewer's expertise. The other issues, though important, normally appear in the written commentary (if they are relevant) and are not essential to running the PC meeting.

The most important thing about the paper's score is to make the *operational* semantics of the score clear. It frequently happens that a PC Member gives a paper a strong accept "because it was the best of the papers I had to review," but not because it was particularly good. The essential semantic categories are:

- A: Good paper. I will champion it at the PC meeting.
 B: OK paper, but I will not champion it.

C: Weak paper, though I will not fight strongly against it.

D: Serious problems. I will argue to reject this paper.

These four positions cover the interesting ones taken by PC Members during discussion. Finer gradations of appreciation are typically uninteresting. Note that it is not important how the scores are labelled — for example, they may still be numeric (i.e., from 1 to 5 or from 0 to 10), or textual (i.e., strong/weak accept/reject), but their meaning must be clear.

The most important positions are A and D, as these are, respectively, the champions and detractors. B and C are fence-sitters, but will also supply arguments for or against. The difference between a B and a C is that B is basically in favour of a paper, but is not willing to champion it, whereas C is not impressed by a paper, but could be convinced if someone else champions it.

Separate ratings of high or low confidence are not especially useful, since low confidence tends to show up anyway as a B or C position.

A separate rating for the reviewer's expertise, on the other hand, is essential to detect the *inexpert champion* situation. The following ratings are used *only* for conflict detection and *not* to rank papers:

X: I am an expert in the subject area of this paper.

Y: I am knowledgeable in the area, though not an expert.

Z: I am not an expert. My evaluation is that of an informed outsider.

Note that it is not necessarily the intention that all reviewers be experts. In conferences with broad scope, it can be useful to have some non-expert reviews to evaluate a paper's accessibility to a general audience. Only in rare situations, however, should a non-expert consider championing a paper.

The scores and expertise ratings would normally *not* be revealed to authors, as this information is purely procedural, and has no function after the selection process is done.

The remaining parts of the review form will vary, but typically include: (i) a *Summary* in which the reviewer briefly summarizes the paper and its main contributions, (ii) *Points in favour/against acceptance* which help focus the discussion at the PC meeting, (iii) *Additional comments for the authors*, and (iv) *Additional comments for the PC*. There may also be parts useful for running the meeting or the conference itself: Paper number, Authors, Title, Referee. Should the author be invited to present a demo? Is the paper suitable for receiving an award? etc.

Rationale

The most important thing for a reviewer to decide is whether he or she thinks that the paper is worth defending at the PC meeting, *not* whether it is a great paper or not. *Make Champions Explicit* helps put reviewers in the right frame of mind.

PC Members who delegate papers to subreviewers are aware that they must be prepared to play the role of champion/detractor on the basis of the review. They are warned in advance if they are the only champion for a paper.

Examples

ICSE 98, ECOOP 98, OOPSLA 98, FSE 6 and others.

Pattern: Identify the Conflicts

Context

You are a PC Chair who has applied *Make Champions Explicit* in designing your review forms. It is now a week before the PC meeting and all (or most) of the reviews are in.

Problem

How should you order or group the papers in preparation for the meeting?

Forces

- It is hard to convert the results of several reviews into a single meaningful score.
- When papers are ranked numerically, there will often be high-ranking papers that are rejected, and low-ranking papers that end up being accepted.
- The PC meeting runs more smoothly if "similar" papers are grouped together, i.e., all those that can be quickly accepted, those that will surely be rejected, those that will generate debate.
- It is important to identify controversial papers as early as possible.
- It is hard to guess which papers will be controversial.

Solution

Group papers according to their *highest* and *lowest* scores. Take care to identify papers with both extreme high and low scores. Do not attempt to rank papers numerically.

The purpose of ranking and classifying papers is to give some structure to the PC meeting by grouping together papers that are likely to require the same kind of debate. Whereas numerical rankings typically fail to achieve this, grouping by presence or absence of champions quickly gets to the point.

A particularly simple and effective way to group papers is to assign a two-letter code to each paper, consisting of the highest and the lowest scores, and to sort the papers by this code. If we are using scores A-D, as described in *Make Champions Explicit*, this yields 10 groups of papers, of which 7 are interesting:

AA, AB: All reviews are positive, and there is at least one champion. These papers will almost certainly be accepted.

AC: This means that all reviews are A, B or C. This is a likely accept, since there is at least one champion, and no strong detractor. The only question is whether the reservations of the C review are serious or not.

AD: This is a serious conflict, and will certainly lead to debate. Note that this does not distinguish between cases where, for example, we have three As and one D, or one A, one B, one C and one D. In practice, the important positions are the extremes.

BB: All reviewers are fence-sitters. Everyone likes the paper, but no one is willing to be a champion. The discussion should determine whether the B's are really A's or not.

BC: These tend to be borderline papers, since no one is willing to be either a strong advocate nor a detractor. Such papers are often put on a "slush pile" and resurrected or discarded after the rest of the programme has been defined.

BD: These papers are likely to be rejected. There is no strong champion, but there is a strong position against acceptance. Such a paper might still be accepted if the B decides after all to champion it.

CC, CD, DD: These papers are almost certain rejects. Papers may be resurrected from this group only under exceptional circumstances, for instance, if it turns out that none of the reviewers were experts, but another PC Member who is an expert in the domain reads the paper during the meeting and decides to champion it.

Note that this classification scheme works independently of the number of reviews each paper receives. What is significant are the high and low scores.

Rationale

Identifying the extremes highlights potential controversy. Where there is no controversy, the PC can typically come to a quick decision.

Examples

Dick Kemmerer reports:

As far as using the approach at ICSE98 goes, ... I had six categories instead of four. ... The categories were: ... Will argue for acceptance (A); Inclined to accept (B); Not opposed to acceptance (C); Not opposed to rejection (D); Inclined to reject (E); Will argue for rejection (F)

Having the two middle categories caused some problems, and I would use only four if I were to do this again.

Before the meeting I separated the papers into three groups:

Group 1 (Likely Accepts): AA, AB, AC

Group 2 (Mixed): AD, AE, AF, BB, BC, BD, BE, BF, CC, CD, CE, CF

Group 3 (Likely Rejects): DD, DE, DF, EE, EF, FF

In addition, all PC Member papers were postponed until all other papers were decided.

We first discussed the Group 1 papers and they were all accepted with little time devoted to them. Next we rejected all of the Group 3 papers without discussing them, although I stated that any PC Member that wanted to bring one of them up was welcome to do so. This did not happen. The middle group as you predicted took the most time.

I also distributed reviews for papers with conflicting reviews to the reviewers beforehand for online discussion. When reviews were changed I entered the updated review before the PC meeting. I also informed each PC Member what papers they would champion before the PC meeting, so that they were prepared.

At the OOPSLA 98 PC meeting, a similar scheme was used (ratings from A-F), with very similar results.

Related Patterns

SortedPaperList [2] has a similar intent, but considers instead the total number of accepts or rejects that a paper has received.

Pattern: Identify Missing Champions

Context

You are a PC Chair who has applied *Make Champions Explicit* in designing your review forms. It is now a week before the PC meeting and all (or most) of the reviews are in.

Problem

What problems should you detect in advance of the meeting?

Forces

- PC Members have limited time to devote to evaluating 15-20 papers.
- PC meetings are expensive to organize.
- Not all reviewers will be experts.
- Not everybody will make it to the meeting.
- Some reviews will be late, or missing.

Solution

Identify which papers are likely to be championed by whom, and make sure that champions are prepared for the meeting. If a potential champion is not an expert, or for some reason cannot attend the meeting, take some compensating action (like soliciting an extra review).

Since PC meetings are expensive to organize (think not only of the travel and hotel costs, but of the salaries paid for those attending!) and cannot be repeated, and the selection process depends so heavily on the identification of champions, it is important to detect potential problems *before* the meeting takes place. This means that reviews should be returned to the PC Chair well in advance — typically at least a week before the PC meeting. To reduce delays, to facilitate analysis of the results, and to permit automatic preparation of review packages, it is a good idea to distribute review forms electronically.

The following situations are variants of “missing champions” that should be detected early to help PC Members better prepare for the meeting:

Late or missing reviews: This is the most basic problem to check for. PC meetings often start with only one or two reviews received for some of the papers. In such cases it is often necessary to get someone to review papers “on the fly” at the meeting, which is clearly a half-measure. Each paper should receive a minimum of three reviews for a fair review process. Check which papers are missing reviews, and pressure the PC members to deliver them. If you doubt that the review will arrive on time, try to solicit an extra review.

Missing champion (or detractor): The PC Chair should check whether anyone who cannot be present at the meeting happens to be the only champion or detractor for a paper. An email exchange in advance between the reviewers may help them come to a consensus in advance of the meeting, or at least to clarify the source of disagreement.

Absent reviewers: This is a variation of the above, in which *none* of the reviewers are present to present their views of the paper. In a large conference with 200 submitted papers and 20 or 30 PC Members, it is almost inevitable that a couple of papers will fall into this category. These papers should be discussed by email. If necessary, another review should be solicited by an attending PC Member.

Unprepared champions: Very often a PC Member is surprised to discover at the PC meeting that he or she is the only champion for a submission. An unprepared PC Member may buckle under negative criticism of the paper and withdraw support. PC Members often pass on papers for evaluation to “subreviewers.” This can be an efficient way to review large numbers of papers, as long as the PC Member carefully checks the papers and reviews before the meeting. If a PC Member ends up being a champion for a subreviewed paper, it is essential that he or she be warned in advance in order to validate or overturn the review.

Inexpert champions: Sometimes during the PC meeting it turns out that the only champions for a paper have low confidence because they are not experts in the problem domain, while the experts either were not assigned the paper, or are only lukewarm about acceptance. Typically non-experts will back down from an “accept” position if there is dissent from an expert. In such cases it can be useful to solicit an extra review from an expert in advance of the meeting.

Low overall expertise: If none of the reviewers is an expert, then the selection process can break down regardless of the scores given by the reviewers. In such cases the PC Chair should solicit an expert review in advance of the meeting.

Finally, papers submitted by PC Members, or for which PC Members have a conflict of interest (i.e., papers submitted by close colleagues) require special treatment during the meeting, but do not constitute problems as such.

Rationale

Identifying champions in advance helps everyone be better prepared for the PC meeting.

Champions who are unable to attend the meeting have a better chance to influence the meeting if they are identified explicitly well in advance.

Pattern: Champions Speak First

Context

You are a PC Chair who has applied *Make Champions Explicit* in designing your review forms. You have also used *Identify the Conflicts* and *Identify Missing Champions* to group the papers and identify potential problems in advance. You are now ready to run the PC meeting.

Problem

How do you focus attention at the meeting?

Forces

- Each PC Member present has only reviewed a fraction of the papers.
- Meeting participants feel obliged to talk about all the papers they reviewed.
- Meandering discussion wastes time.

Solution

Discuss the papers in groups, according to *Identify the Conflicts*. For each paper, first invite a champion to introduce the paper, and then to present reasons why it should be accepted. Then invite any detractors to explain why they think it should not be accepted. Finally open the general discussion, and try to reach a consensus. For any paper, if there is no champion, *the paper should not be discussed*.

If no conscious effort is made to quickly identify champions, much time can be wasted discussing papers that have no chance of being accepted. Very often, when a paper comes up for discussion, a PC Member will start by saying, “Well, I didn’t like this paper because ...” This is not very useful, first of all because it does not tell the rest of the PC what the paper is about. Second, it does not lead to effective decision making, since the purpose of the meeting is to *accept* papers, not to reject them (i.e., it is more productive to concentrate on discussing papers that have a chance of being accepted than those that don’t). Long unfocused discussions with delayed decisions may exhaust all reviewers. In the end, the decision taken may depend on who has the most stamina.

It is good to set some ground rules to keep discussions focused. For each paper, the champion, if one exists, or the closest there is to a champion, should introduce the paper by briefly summarizing it and presenting the points in its favour. Then, the detractor (or whoever has the strongest negative points) should speak next. Finally the remaining reviewers can back up these arguments, or fill in missing points. If there is a detractor, then the champions and detractors typically play the roles of defence and prosecution in a trial, and the rest of those present play the role of the jury. Frequently either the champions or the detractors become convinced by the arguments of the other, and a consensus is quickly reached.

If no consensus is possible, it may be necessary to ask the PC to vote. In this case *all* PC Members present who have participated in the discussion should vote (since they act as a jury).

PC Members should also be reminded of the criteria for acceptance. These may be more stringent, or more lax, depending on the nature of the conference, or may be quite specialized. Typically, an accepted paper should have a clear, original contribution, and fulfil the usual criteria of readability, completeness, etc. Originality is a strong criterion, and many papers fail to be accepted if they do not clearly demonstrate new results. A champion/detractor should address these specific criteria.

If there is no clear champion for a paper, the discussion should focus on checking why no one wants to champion it (i.e., to try to smoke out a reluctant champion). If no champion can be identified, the paper can be quickly rejected.

Delaying a decision on a paper is almost always a bad idea. A decision should only be delayed if something will happen in between that may change the outcome, i.e., if an expert will check the paper for originality. Borderline papers (BC grouping) may be delayed until the other papers have been considered.

On the other hand, if the champion for a paper is not present, then the discussion and decision *must* be delayed until that person either arrives, or can be consulted by e.g., telephone or e-mail.

It is highly recommended to supply each PC Member attending the meeting with copies of all the reviews for which they do not have conflicts. This makes it easier for everyone present at the meeting to actively participate in the decision-making process, even for papers they have not personally reviewed.

The actual order in which papers are handled does not seem to matter much in practice, as long as they are handled in relatively coherent groups. It is common to start by accepting as many of the "easy" papers as possible, and by rejecting all of the unchampioned papers before starting in on the controversial papers.

Rationale

Discussions tend to be shorter and more focused if they can only take place when a champion is identified. Delays only take place if there is a chance that a new champion can be identified.

Explicitly encouraging PC Members to champion papers provides opportunities to draw reluctant champions out in the discussion. (Each reviewer can be explicitly asked, "Are you willing to champion this paper?")

Resulting Context

"Champions Speak First" can stifle debate if applied too rigorously. One must be careful not to discourage reluctant champions.

Examples

This pattern is so common that there are few PCs that do not apply it in some form. The key difference is whether the PC meeting has been prepared in advance by applying the related patterns mentioned in the Context.

Related Patterns

PaperChampion [4] presents a similar idea, though without advance preparation by champions.

Pattern: Consensus on PC Papers

Context

You are a PC Chair running a PC meeting according to *Champions Speak First*.

Problem

How should PC papers be handled?

Forces

- Sometimes weak papers are accepted to conferences where one of the Authors is a PC Member.
- A conference will not be taken seriously if it appears that PC Members can get their papers more easily accepted than other Authors.
- You may want to apply more stringent requirements to PC papers than to other submissions.

Solution

PC papers should be accepted only if there is at least one champion and there are no serious (expert) detractors.

Whenever a PC paper is discussed (and, typically, whenever a PC Member has a conflict of interest with a paper being discussed), the PC Member concerned leaves the meeting, and is only called back when a decision has been taken.

Rationale

By making sure there are no detractors, PC papers will be accepted only if there is a consensus. This generally ensures that such papers are “at least as good as” the best papers accepted to conference.

Acknowledgements

Many thanks to Serge Demeyer, Rachid Guerraoui, Dick Kemmerer, Tony Simons and Jim Coplien who suggested numerous improvements to the presentation and the contents. Thanks as well to the PCs and PC Chairs of ECOOP, ESEC, ICSE and OOPSLA who have helped make this pattern explicit.

I am also indebted to Mike Beedle, the PLoP 98 shepherd for this paper, and to my PLoP 98 workshop group, who suggested numerous simplifications and encouraged me to decompose *Identify the Champion* into a pattern language. I especially thank Bob Hanmer & David Cymbala who helped me factor out the individual patterns.

Many thanks as well to Ian Chai and colleagues at UIUC who workshopped the paper and sent additional comments. Finally, I am grateful to the PLOPD4 reviewers who suggestion several simplifications and cuts.

Bio

Oscar Nierstrasz is Professor of Computer Science at the Institute of Computer Science and Applied Mathematics of the University of Berne where he leads the Software Composition Group. He is interested in all aspects of component-oriented software technology. He has served on the programme committees of ECOOP, OOPSLA, ESEC and many other conferences, and as the Programme Chair of ECOOP '93 and ESEC '99.

References

- [1] Kent Beck, OOPSLA 93 Panel, <http://www.acm.org/sigplan/oopsla/oopsla96/how93.html>
- [2] Kent Beck, “SortedPaperList,” <http://c2.com/cgi/wiki?SortedPaperList>, in [3]
- [3] Jim Coplien, “OopsIaProgramChairPatterns,” <http://c2.com/cgi/wiki?OopsIaProgramChairPatterns>
- [4] Jim Coplien, “PaperChampion,” <http://c2.com/cgi/wiki?PaperChampion>, in [3]
- [5] Ralph Johnson, “AssigningPapersToReviewers,” <http://c2.com/cgi/wiki?AssigningPapersToReviewers>, in [3].