How Not to Write a Good Systems Paper

Roy Levin and David D. Redell

詹剑锋 讲解

为什么第一节课讲？

- 系统是一门综合性的研究方向
  - 脏活/累活(如何出彩?)
- 可以指导系统方向的读研生涯

论文分类

- 描述一个实际系统.
  - a global survey of an entire system.
  - a selective examination of specific themes embodied in the system.
- 描述一个未实现的系统.
  - utilizes ideas or techniques that you feel the technical community should know.
- 理论研究领域的某个主题.
  - performance modelling or security verification.

论文评估准则

- Original Ideas (新概念系统)
- Reality
- Lessons
- Choices
- Context
- Focus
- Presentation
- Writing Style

Original idea

Idea足够新吗？

- 至少一个新idea!

- There is no point in submitting a paper to a conference or journal concerned with original work unless the paper contains at least one new idea.

如何确认你的工作是否有创新?

- 你必须熟悉 **state-of-the-art和state-of-the-practice.**
- Perhaps the most common failing among the submissions in the first category (real systems) was **an absence of new ideas**;
- the systems described were frequently isomorphic to one of **a small number of pioneering systems well-documented in the literature**.

能不能简明地说清楚你的ideas?

- If your paper is to advance the state of knowledge, your reader must be able to find the new ideas and understand them.
- Try writing each idea down in a paragraph that someone generally versed in the relevant area can understand.
  - **If you can't, consider the possibility that you don't really understand the idea yourself.**
- When you have the paragraphs, use them in the abstract for the paper.

确切地描述你要解决的问题

- Your reader cannot be expected to guess the problem you faced given only a description of the solution.
- **Be specific.**

- Be sure to explain why your problem couldn't be solved just as well by previously published techniques.

值得写长篇大论吗?

- Frequently, papers describing real systems contain one or two small enhancements of established techniques.
  - The new idea(s) can be described in a few paragraphs; a twenty-page paper is unnecessary and often obscures the actual innovation.
- Since construction of a real system is a lot of work, the author of the paper sometimes **unconsciously confuses the total effort with the work that is actually new**.
  - ("My team worked on this system for two years and we're finally done. Let's tell the world how wonderful it is.")
- If the innovation is small, a small paper or technical note in a suitable journal is more appropriate than an SOSP submission.

是否和相关工作有显著差异？

- An obvious extension to a previously published algorithm, technique, or system, does not generally warrant publication.
- You must show that your work represents **a significant departure from the state of the art**.
- If you can't, you should ask yourself why you are writing the paper
  - Why anyone except your mother should want to read it.

论文引用的重要性

- You will have difficulty convincing the *skeptical* reader of the originality of your efforts unless you specifically distinguish it from previously published work.
  - This requires citation.
- Furthermore, you will find it harder to convince your reader of the superiority of your approach if he has read the cited works and you haven't.

论文引用需要注意的问题

- The answers to these questions help alert you to blind spots in your knowledge or understanding.
  - What is the oldest paper you referenced? 1960年代?
  - The newest?
  - Have you referenced similar work at another institution?
  - Have you referenced technical reports, unpublished memoranda, personal communications?

- Remember that citations not only acknowledge a debt to others, but also serve as an abbreviation mechanism to spare your reader a complete development from first principles.

与已有工作比较是否clear and explicit?

- You cannot simply say: "Our approach differs somewhat from that adopted in the BagOfBits system [3]."
- **Be specific**:
  - "Our virtual memory management approach uses magnetic media rather than punched paper tape as in the BagOfBits system [3], with the expected improvements in transfer rate and janitorial costs."

实现的价值

- Implementation experiences supporting or contradicting **a previously published paper design** are extremely valuable and worthy candidates for publication.

- Designs are cheap, but implementations are expensive.

中科院计算所

Reality

论文描述的系统是否已经实现？

- Your reader has a right to know at the outset whether the system under discussion is real or not.

实现的价值

- 如果系统已经实现,如何使用该系统?
  - What has this usage shown about the practical importance of the ideas?
- A multiple man-year implementation effort does not of itself justify publication of a paper.
- If the implemented system contains new ideas, it is important to explain how they worked out in practice.

如何写未实现的设计论文

- If the system hasn't been implemented, do the ideas justify publication now?
  - This can be a difficult question for an author to answer dispassionately, yet any reviewer of the paper will make this judgment.
- It is always tempting to write a design paper describing a new system, then follow it up in a year or two with an "experience" paper.

如何写设计论文？

- The successful papers of this genre nearly always include initial experience in the closing sections of the design paper.
- The subsequent experience paper then deals with the lessons learned from longer-term use of the system, frequently in unanticipated ways.
- Reviewers are very skeptical of design-only papers unless there are new ideas of obviously high quality.

Lessons

从该工作中获得什么收益？

- If you didn't learn anything, it is a reasonable bet that your readers won't either
- You've simply wasted their time and a few trees by publishing your paper.

What should the reader learn from the paper?

- Spell out the lessons clearly. Many people repeat the mistakes of history because they didn't understand the history book.

How generally applicable are these lessons?

- Be sure to state clearly the assumptions on which your conclusions rest.
- Be careful of generalizations based on lack of knowledge or experience.
- A particularly common problem in "real system" papers is generalization from a single example.
- **When stating your conclusions, it helps to state the assumptions again.** The reader may not have seen them for fifteen pages and may have forgotten them. You may have also.

中科院计算所

Choices

解释每一个设计与实现选择

- Why were the choices made the way they were?
- 好的论文不仅描述,而且解释.
  - save future researchers from following the same blind alleys.
- You also want to record potentially interesting side-streets you didn't happen to explore.
  - Make sure to state clearly which is which.

选择的确切动机

- Did the choices turn out to be right, and, if so, was it for the reasons that motivated them in the first place?
- If not, what lessons have you learned from the experience?
  - How often have you found yourself saying "this works, but for the wrong reason"?
- Many papers present a rational argument from initial assumptions all the way to the finished result when, in fact, the result was obtained by an entirely different path.
  - This kind of "revisionist history" borders on dishonesty and prevents your readers from understanding how research really works.

中科院计算所

Contexts

What are the assumptions on which the work is based?

- The skeptical reader is unlikely to accept your arguments unless their premises are stated.

- Make sure you get them all; it's easy to overlook implicit assumptions.

前提条件的合理性

- For "unimplemented systems" papers, this amounts to asking whether the assumptions of the design can hope to support a successful implementation.
- Many paper designs are naive about the real characteristics of components they treat abstractly,
- For theoretical studies, it must be clear how the assumptions reflect reality,
  - e.g., failure modes in reliability modelling, classes of security threats in security verification, arrival distributions in queuing systems.

Formal model

- If a formal model is presented, does it give new information and insights?
- Simply defining a model for its own sake is not very useful.

Focus

介绍性材料是否多余?

- "Real system" papers are particularly guilty of irrelevant description.
  - For a distributed file systems, the physical characteristics of the connection between computer and communication network are probably not germane.
- Avoid the temptation to describe all major characteristics of your system at the same level of depth.
  - Concentrate instead on the novel or unusual ones that (presumably) will be the focus of the original technical content of the paper.

是否包含充足的材料支持读者

- Do not assume that the reader has read every referenced paper within the last week and has them at his fingertips for instant reference.
  - "We adopt the definition of transactions from Brown [4], layering it onto files as described by Green [7, 18], with the notions of record and database introduced by Black [10] and White [12] and later modified by Gray [6]".
  -
- On the other hand, don't burden your reader unnecessarily with lengthy extracts or paraphrases from cited works.

Presentation

表达需要注意的问题

- Are the ideas organized and presented in a clear and logical way?
- Are terms defined before they are used?

Are forward references kept to a minimum?

- Readers get annoyed when they repeatedly encounter statements like "Each file consists of a sequence of items, which will be described in detail in a later section".
  - It's all right to ask him to do this once or twice, but only when absolutely necessary.

- Even if you can't afford the digression to explain "item" at this point, give the reader enough information to attach some meaning to the term:
  - "Each file consists of a sequence of items, variable-sized, self-identifying bit sequences whose detailed interpretation will be discussed below under 'Multi-media Files'."
  - Your reader may not yet understand your concept of files completely, but at least he has some glimpse of the direction in which you are leading him.

Have alternate organizations been considered?

- Theoretical papers, particularly of a mathematical character, are generally easier to organize than papers describing systems.
- The expected sequence of definition, lemma, theorem, example, corollary works well for deductive argument, but poorly for description.
- In "real system" papers, much depends on the intent: global survey or selective treatment.
  - Frequently, difficulties in organization result from the author's unwillingness to commit to either approach.
  - Decide whether you are surveying your system or focusing on a specific aspect and structure the paper accordingly.

Was an abstract written first?

- Does it communicate the important ideas of the paper?
- Avoid the passive voice and include a simple statement of assumptions and results.
  - "We designed and implemented a user interface following the ideas of Keysworth and discovered that converting the space bar to a toe pedal increases typing speed by 15%. However, accuracy decreased dramatically when we piped rock music instead of Muzak (tm) into the office."
- Leave discussion and argument for the paper. It helps to write the abstract before the, since it focuses your attention on the main ideas you wants to convey.

Is the paper finished?

- Reviewers can often help you to improve your paper, but they can't write it for you.
  - Moreover, they can't be expected to interpolate in sections marked "to be included in the final draft".
- Similarly, in a paper describing a system, a reviewer cannot tolerate the omission of important explanation or justification.
- Omitting sections with a promise to fill them in later is generally unacceptable.

Writing Style

写作要点

- Is the writing clear and concise?
- Are words spelled and used correctly?
- Are the sentences complete and grammatically correct?
- Are ambiguity, 俚语, and cuteness avoided?

Presentation的重要性

- If you don't have sufficient concern for your material to correct errors in grammar, spelling, and usage before submitting it for publication, why should you expect a reviewer to read the paper carefully?
- Some reviewers feel that this kind of carelessness is unlikely to be confined to the presentation, and will reject the paper at the first inkling of technical incoherence.
- "Please let me convince you that I have done interesting, publishable work."

Summary

小结

- These questions can help you write a better technical paper.
- Consult them often as you organize your presentation, write your first draft, and refine your manuscript into its final form.
- Some of these questions address specific problems in "systems" papers; others apply to technical papers in general.
- Writing a good paper is hard work, but you will be rewarded by a broader distribution and greater understanding of your ideas within the community of journal and proceedings readers.

Reading lists

- Armando's Paper Writing and Presentations Page
- http://www.eecs.berkeley.edu/~fox/paper_writing.html

中科院计算所

谢谢