SOSP 2019 Diversity Workshop

How to read a (systems) paper

Rebecca Isaacs
Twitter

What we will cover

1. About papers

   Take-away: Knowledge about the publication process

2. Reading papers

   Take-away: A technique for efficient and effective paper-reading

3. Bonus material (if time): Discussing papers

   Take-away: What to do in a reading group

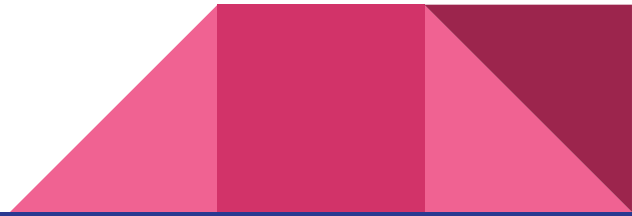

Qualifications and a disclaimer

Who am I to tell you this stuff?

- 13 years as a researcher at Microsoft Research
- Many papers written… some published, many rejected :-(
- Lots of experience on program committees
  - SOSP, OSDI, NSDI, EuroSys, ATC, ASPLOS, HotOS, HotNets, …

But, I only know about the systems community.

The paper mill

Why write and publish research papers?

- To disseminate research and make technical contributions
  - Have impact, influence the field
  - Give back
- A forcing function for high standards
  - Peer review ⇒ rigorous scrutiny of claims and assumptions
- For career advancement
  - Funding, tenure, evidence of intellectual achievement
- To be part of a community

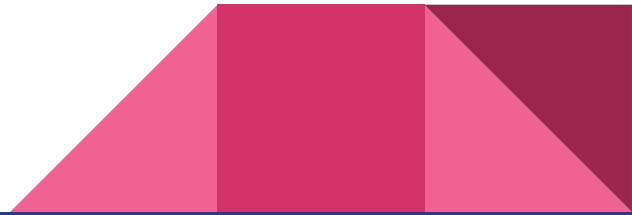

The peer review process

Question: why does a paper not look like a blog post?

Understanding how publication works helps us to understand why the papers are the way they are.

Peer review (1): The call

The conference* posts a CFP:

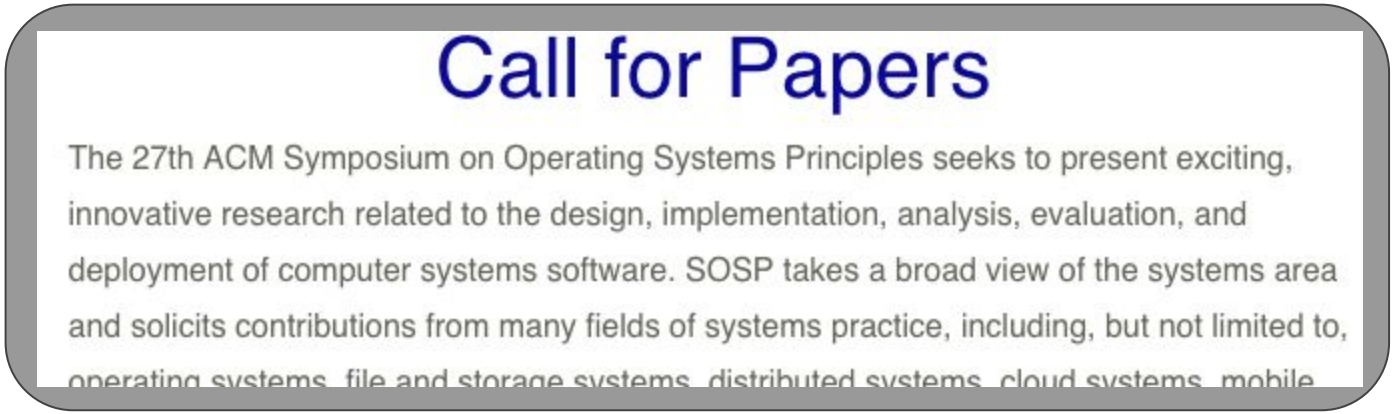

*In systems, conferences are more prestigious than journals**

**Some “best papers” are re-published in journals after appearing at the conference as a mark of honor

Peer review (2): Example timeline

The authors submit a paper by the deadline

- Typically (but not always) 12 pages and 99% complete
- Increasingly, reviewing is double blind: both authors and reviewers are anonymized.

The program committee reviews the submissions over ~3 months, then meets (often in person for 1-2 days) to choose the final program.

Accepted papers are then revised, sometimes extensively and usually with the help of a shepherd from the PC.

Peer review (3): The program committee (PC)

Selected (“invited to serve”) by the PC chair(s)

- the chairs are appointed by the Steering Committee
- all positions are unpaid

The chairs usually aim for a mix of expertise, experience, genders, races, geographical locations, academia and industry

- they rarely meet their goals entirely (people say no)

Peer review (4): Acceptance rates

For the best conferences, acceptance rates are 15-25%

- Highly competitive
- Who publishes?

From csrankings.org (systems and networking):

#	Institution	Count	Faculty
1	Massachusetts Institute of Technology	19.4	26
2	University of California - Berkeley	12.1	22
3	University of Washington	11.9	15
4	Stanford University	8.9	14
5	Carnegie Mellon University	8.7	21
6	University of Michigan	7.6	14
7	University of Wisconsin - Madison	7.5	9
8	Princeton University	7.3	8
9	Cornell University	5.4	11
10	Harvard University	4.5	5

Peer review (5): Reviewing

- Multiple reviewing rounds
  - 3 reviews per paper round 1, approx 50% dropped
  - 2-3 more reviews for remaining papers
  - all candidate papers discussed online before the meeting
  - 60-80 papers discussed in the PC meeting
  - single-track conferences accept around 40 papers
- The reviewer’s workload (“heavy”):
  - 25-30 papers over 3 months
  - typically 3 hours / paper

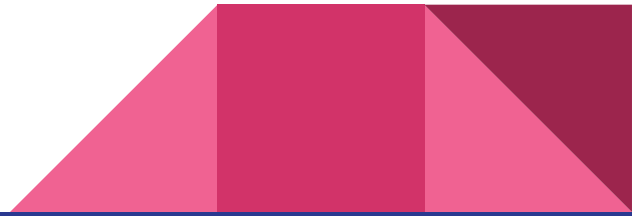

The peer review process (6): Reviews

- The good ones:
  - Constructive, insightful, detailed
- The bad ones:
  - Adversarial, incorrect, brief, vague

Let’s look at some examples!

Review example

From Frank McSherry's OSDI 2018 reviews

Strengths

The work adds indexes to timely differential dataflow systems, which are very interesting streaming computing engines. If indexes can improve performance and simplify the writing of timely dataflow programs, that's a great win.

Areas for improvement

Explain in plainer terms the specific problems of existing dataflow designs that this paper deals with. The current description of these problems is overreaching and confusing...

Comments for author

I was excited to read this paper, because... Unfortunately, this paper was very confusing to me. I couldn't grasp what new properties the indexes added, ...

Review example

From Dave Andersen's blog post, a super-bad review:

> Contribution: "The paper does not propose any new idea or methodology."
>
> Weaknesses: "The paper is very superficial."

Also see "How not to review a paper" by G Cormode

Implications of peer review

Many of the papers we read have been submitted, rejected and revised multiple times.

Look for certain phrases:

- “Our contributions are 1. … 2. … 3. …”
- “The key insight of this work is …”
- “We show that *x* is better than *y* by *z*% under workload *w*”
- Upfront admissions of weakness in comparison to some related work

More implications of peer review

Things (top systems conference) reviewers don’t like:

- Incremental papers (“Least Publishable Units”)
- Poor experimental methodology
- Inflated claims
- No characterisation of the system’s limitations
- Engineering without ideas or insights

How to read a paper

Reference

How to Read a Paper (living document)
S.Keshav
ACM SIGCOMM Computer Communication Review, 37(3)
July 2007

See also:

- Murat Demirbas blog post

The approach

- Read in 3 passes, gradually diving deeper
  - Pass 1: the general idea (a “bird’s-eye view”)
  - Pass 2: content without details
  - Pass 3: the details, thoughtfully
- Leaving some time between passes can help to absorb ideas
- Depending on the purpose, 1 or 2 passes may be enough

The first pass

Should take 5-10 mins

Read the title, abstract, intro and conclusion.

- Sometimes the conclusion is the most informative (factual vs aspirational)

Look at the section headings, note how the material is organised

What to know after the first pass (“the five Cs”)

1. **Category**: type of paper (system, measurement study etc)
2. **Context**: related work
3. **Correctness**: valid assumptions?
4. **Contributions**
5. **Clarity**: well written?

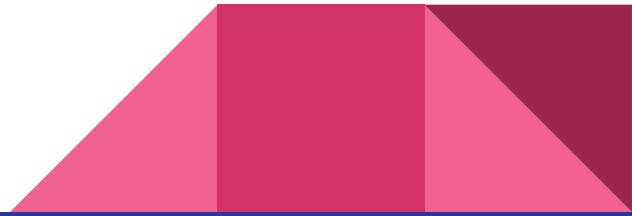

The second pass

Read the paper! Expect this to take at least 1 hour.

Look carefully at the diagrams, tables and graphs.

Check the references

- Look at venues and years published
- Have you read any of them? Are there some you haven't read but seem interesting?

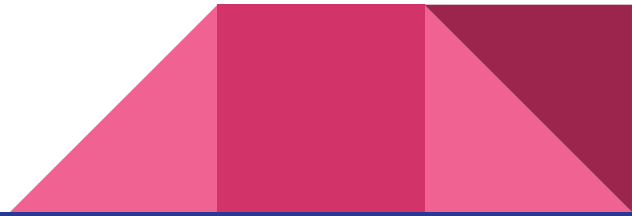

After the second pass...

… you should be able to describe what the paper is about to someone else who hasn’t read it.

This is a skill that improves with practice. Try it out on your friends and colleagues.

The third pass

Read again, in great detail, to “virtually re-implement” the paper (1-5 hours)

- Identify and challenge assumptions
- Think about how you would approach the problem, evaluate the solution, present the material, etc.
- Look for what the authors have not said or evaluated

The third pass takes the most time.

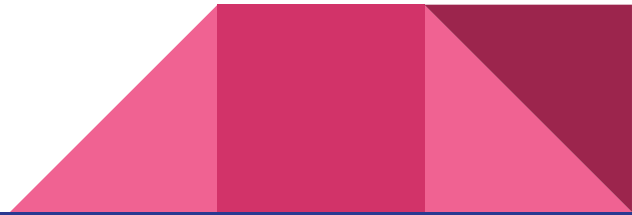

After the third pass...

You will have a pretty good idea of whether you think this is good work and a good paper, and you will be able to articulate why.

You may have a few ideas for future work yourself!

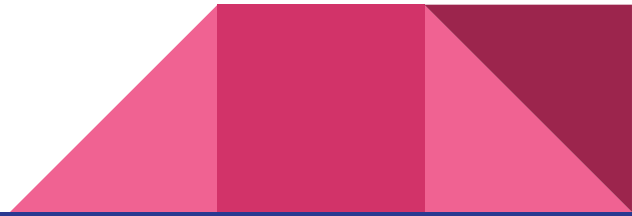

Dealing with the material

Papers can be dense and daunting. In addition to 3 passes, some other tricks:

- Don’t read in order. Jump around however you like.
- If the terminology is new, take one term at a time and try to work out its definition in isolation.
- Allow hours or days between passes.
- Watch a video of the conference presentation.
- Find somebody else’s write-up of the paper.

Discussing a paper

Discussing a paper: Peel the onion

Top level

- What’s the cool idea?
- Why does it matter?
- How, when, where and why does it work?

What questions do we have about the paper?

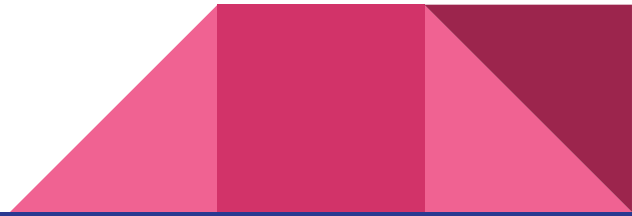

Discussing a paper (cont)

Mid level: execution

- What are the assumptions? How is the work scoped?
- What's the evaluation?
  - Setup, workload, choice of experiments
- Does the evaluation support the claims?
- Does the paper do a good job with related work?
  - Including existing (unpublished) systems
- Did you enjoy reading the paper? Are you excited about the ideas?

Discussing a paper (cont)

The heart of the discussion

- How does the work advance the state-of-the-art?
- What are the paper's strengths and limitations?
- Will it have a big impact? If so, how?
  - For older papers, does it stand the test of time?
- How does this paper stack up against reality?
  - Is the work applicable in "the real world"?
  - Do other systems solve the same problem differently?

Final thoughts

Writing and reading papers can be challenging, rewarding, and fun!

This advice is not intended to be prescriptive. Please feel free to cherry-pick what you like and ignore the rest.