

Systems and Papers

Kang Chen

Oct. 8, 2024

Disclaimers

- This talk should never have happened. This was entirely coerced by Prof. Cheng Li. I have not gone crazy thinking that I can do this topic.
- **The talk below is completely untrue and just lies, so DO NOT believe even a single word!** (excluding this page, and some materials from other guys.)
- **The content in this presentation reflects personal thoughts and should be evaluated based on its relevance to your context.**
- If you do not want to waste your time, QUIT now!
- If you accidentally got this manuscript, just throw it away! It's harmful!
- Some claims, even if they are true, in this manuscript might only work for very few universities and colleges in mainland China.

Systems Research Community

- **Operating Systems, Distributed Systems, Storage Systems, Database Systems, Network Systems, Security/Dependability, Programming Languages, Software Engineering, ML Systems.....**
- **Researchers in various fields of systems studies contribute to this community.**

Driving Forces of Systems Research

- **Application:**
 - Internet, Datacenters, Graph Computing, Serverless Computing, Data Mining, AI (LLM) !
 - Theory (erasure coding, consensus, formal)
- **Computer Architecture:**
 - Processing: Multi-core, SIMD, GPU, xPU
 - Storage: Cache, DRAM, NVM, CXL, PIM, Processor in SSD (Firmware)
 - Networking: Disaggregated Memory, SmartNIC, Smart Switch
 - Quantum Computing

What Systems Want to Achieve?

- **Correctness**
 - Specification, Invariants
 - Consistency: Atomicity, Crash Consistency, Consensus, Serializability
 - Security
- **Performance**
 - Throughput (OPS, TPS)
 - Latency
 - Tail Latency
- **Generality**
 - Adaptability
 - Fairness
 - Isolation

Messages from Professionals

- How (and How Not) to Write a Good Systems Paper
- <https://www.usenix.org/conferences/author-resources/how-and-how-not-write-good-systems-paper>
- Roy Levin and David D. Redell,
Ninth SOSP Program Committee Co-chairmen, 1983

- How to Write a Good Systems Paper
- <https://www.sslab.ics.keio.ac.jp/apsys2015/assets/howto-systems.pdf>
- By Gernot Heiser

What to Remember?

- While there is no guaranteed formula for paper acceptance, improving clarity and innovation increases your chances.
- There are **plenty of reasons** your paper gets rejected!
- Rejection is part of life, get used to it!
- **Reviewers look for reasons to reject** --- Don't give them any!
- There's no need to be too pessimistic about it either. I often see reviewers looking for reasons to accept a paper.
- **PC: At least I should learn something from this paper.**

Reasons to Reject Systems Paper

- Good papers get rejected. OSDI2023 rejected paper, SOSP2023 best paper! (See disclaimers!)
 - Get rejected for wrong reasons.
 - It is part of life, get used to it!
- Does this paper address a real problem? (motivation)
Choosing problems with real-world applications, industry relevance, or those that improve existing systems.
- Does the paper really solve the issues raised? (system design vs. system evaluation)

Reasons to Reject Systems Paper

- **Lack of sufficient innovation.** Engineering work is acceptable for deployed systems but should address real-world workloads. Engineering work is often considered a small contribution to a research prototype.
- **Novelty vs. Incremental Work:** Exploring whether your work solves an existing problem in a novel way or simply improves upon previous work.

Reasons to Reject Systems Paper

- Better fit another venue (ASPLOS, VLDB, PODC, DISC). Limited system level contribution.
- Not meet the bar of {SOSP, OSDI, EuroSys, NSDI...}
 - **Believe me, there is no such bar!**
 - **Not enough innovation/system level contribution. Incremental work does not fit top conferences**
- Not compare to SOTA. (I hate published papers (especially from universities) with no source code. But sometimes you have to rebuild previous systems.)
- Not an apple-to-apple comparison.
- No breakdown (ablation study). **Yes, this sometimes is one bar!**

Reasons to Reject Systems Paper

- No enough details / too many details --- What?
 - No enough details → It's unclear whether the system was built or whether your method solves the problem.
 - Too many details → You buried your main idea.
 - Guideline: At least, theoretically, it is feasible to reproduce the system after reading this paper.
- I don't understand. (Badly written)
- No enough related work. (**Do your homework! A firm reason to reject a paper.**)

Reasons to Reject Systems Paper

- **Your claim is misleading.**
 - **Think about alternatives. Why is your choice reasonable?**
 - **You claim your method solves the challenge. However, it is possible that the improvement comes from better engineering.**
- **Not general. (System paper should deal with a set of important applications, not a specific application. But what about transformers? It depends.....)**

Reasons to Reject Systems Paper

- The paper in its current form is not ready for publishing. (Writing issues, system immature, evaluation not enough)
- Assumption is not realistic. (not convincing)
- Important parts (design, implementation, evaluation, discussion, related work) are missing. (not convincing)
- **(Rare case)** Correctness: Is the system logically sound? Is it evaluated thoroughly and honestly? (If reviewers are sure something is wrong, a definite rejection.)
- Poor writing! (poor English.) Reviewers will try their best to understand your paper. But don't torture them. They are kind people. **Recent years, reviewers try to find reasons to accept, not to reject.**

Reasons to Reject Systems Paper

- **Lucky!**
- **Raise the possibility your paper gets accepted.**
 - **Good system + Good writing**
 - **Decent system + reasonable writing**

Writing a paper

But, Wait! System or Paper, Which One is More Important?

- Can a bad or even poor job become a decent paper?
- **Maybe yes, it's possible. But you'd better not roll your dice!**
- Sometimes, bad papers got accepted. That is also part of life. We are not doing research in math.

First Priority is to Carefully Consider Your Project

- **Sadly, there is no formula for you to get an idea.**
 - This is what research is.
- **Different people have different research styles. You have to find your own way to get ideas.**

Working on Your Idea

- Pick your area. Read 3~5 papers in your area in top conferences/research teams. (Plenty of tutorials/papers/videos on how to read papers.) (Assuming that you have taken enough advanced courses from **Bilibili (MIT, CMU, Stanford, Berkeley, ETH, EPFL), for example.**)
- Only read papers from recent years. If you can advance some recent works, **there is a good reason to publish a new paper.**
 - **Often, this is misleading.** Authors are in a more appropriate position to improve their own work. Try to review the papers from different perspectives.
 - There are other ways. For example, working with industry, solving your motivation (But, there are some dangers here. **Engineers do not read papers but are ridiculously confident.**)
- Create your initial idea. Don't be afraid of silly ideas.
- Brainstorming! Brainstorming! Brainstorming!

When to Start Running Your Project?

- Remember, you want to use your project to publish a paper, not a startup. Minimal efforts.
- Not until you are sure a paper can be written. Before that you can write some validation code to validate your idea. But not a serious system, not even a prototype.
- Weekly team meeting. Get feedback from your mentor/professor and teammates.
 - Read more related papers to validate that your idea is new.
- I had experiences forming a workable idea after a whole semester. (Two SOSP papers.)
- Some researchers are particularly good at coming up with good ideas magically/mysteriously. Find your own way.

How to Use Your Mentor/Professor?

- Present your idea to your mentor/professor. If she/he can understand, maybe others can, too.
- Ask your mentor/professor for resources. Maybe she/he can find a device in a big company.
- If your mentor/professor or other victims can help you revise/polish your paper, draft the first version. Remember, no one else can do this for you.
- **(MUST!)** Before submitting a paper, ask all co-authors (including your mentor/professor whose names mysteriously appear on the paper.) for **CONFLICTS!!! Take this seriously.**

What Materials are Ready Before writing?

- Motivation (why you do this?)
- Contributions (a clear improvement over existing works)
- Preliminary results (at least some improvements)

- Present the first draft to your mentor/professor/collaborator at least one month before the deadline! **(Don't create excuses for her/him.)**
- **Some students show me the draft two/three days before the deadline! You better QUIT your PhD immediately!**


Following Traditional/Classic Paper Format Strictly

- The paper's purpose is to reduce the reading burden. By organizing the paper in the same format, it can help to reduce the reading burden and better convey information to readers.
- Introduction, Background and Motivation, [Overview, Design, Implementation], Evaluation, [Discussion], Related Work, Conclusion.
- There is no need to be too creative here!

Writing a paper

Gernot Heiser, thank you!

With my own understanding. All faults are mine.



Keep Something in Mind

- One (main) message
 - Paper is not a tutorial/project report/manual/thesis.
 - Make sure that the reader gets it.
- Clear
 - Every sentence, paragraph, section has a clear purpose.
 - The purpose is clearly communicated.
- Concise (Brief but complete)
 - Don't waffle!!! (~~Use Jay's rule of thumb~~)
 - Be precise. ///mmmmmmm
 - ~~Make sure it's readable, lucid, enjoyable.~~
- Convincing
 - Be specific.
 - Everything you write must support the message.
 - The overall message is consistent.

Narrative

- It starts from zero and then works on transmitting the message.
- Maintain reader state! Be conscious of what the reader knows/remembers.
 - **help readers to read, rather than causing a reading barrier.**
 - define before use.
 - be aware of what the reader has learned.
 - recall/remind if necessary.

Paper Engineering

- **The best work is useless if you can't convince the reviewers.**
 - Reviewers are busy, may have to review 3020 papers.
 - They'll look for reasons to reject – don't give them any!
- **Important bits:**
 - Introduction: sell the idea, the significance and the approach
 - build tension, make reader interested
 - convincing argumentation
 - top-down, not bottom-up
 - maintain reader state
 - **convincing evaluation – thorough and honest**
 - **state assumption/limitations honestly**

Introduction: Most Important Part of the Paper!

- **The Overture:**
 - Explain the problem you're solving
 - Outline your approach
 - Indicate results/outcomes
 - State contributions
- **General hints for intro:**
 - Capture the reader's interest: sell your idea
 - Be concise: Stay within about one page! **//mmmmmm**
 - Make sure the paper delivers what you promise
 - – Reviewers kill for “bait and switch”

Other Parts

- **Background: set the scene in more detail**
 - cite related work as needed, don't discuss more than necessary
 - **Examples!!!!**
- **Describe the problem in detail**
- **Explain the solution in detail**
 - be honest and forthcoming with limitations and assumptions
- **Evaluation: often the largest part**
- **Related work. (Categorize related work. express the difference between each relevant work and your paper.)**
- **Conclusions (Can use terms in the paper. Not an option for abstract.)**
- **Abstract**
 - used to steer to the right reviewers
 - What, Why, Achievement, Implication
 - **IMPORTANT: Redo for camera-ready! ////mmmmmmm**

Style and Form

- **Write in an engaging style**, leading readers through the paper
 - avoid bottom-up structure, present ideas top-down
 - **follow style rules**
 - Use active voice!!!!
 - Avoid buzzwords (“novel”, “mobile social supercomputing in the cloud”)
- Be mindful of readers’ brain state (which is lossy)
 - maintain reader state
 - **don’t assume every reviewer is an expert in your narrow area**
 - but don’t think you can hide stuff from reviewers!

Style and Form

- Follow formatting rules
 - don't play with margin, baseline skip, etc.
 - don't use microscopic fonts, >40y olds have problems with
- Spell-check, proof-read, proof-read
 - get a native speaker to proof-read if you aren't ///your mentor?
 - get outsiders to read it – great way to spot holes before it's too late!

Mechanics

- **Use revision control**
 - especially (but not only) when it's a joint paper
- **Don't use MS Word !!!!!**
 - doesn't integrate well with revision control
 - requires coarse-grain locking
 - references are painful, formulae even more so
 - MSR people use LaTeX, so should you! (Donald Knuth at Stanford, Leslie Lamport at MSR)
- **Use BibTeX**
 - ... but use it correctly *///mmmmmmm*

Further Reading

- **Writing systems papers:**
 - Levin & Redell: An evaluation of the 9th SOSP submissions, or How (and how not) to write a good systems paper
 - Simon Peyton Jones (MSRC): How to write a great research paper
 - <http://research.microsoft.com/en-us/um/people/simonpj/papers/giving-atalk/giving-a-talk-slides.pdf>
 - Gernot Heiser's paper/thesis writing guide
 - <http://gernot-heiser.org/style-guide.html>
- **General writing/style, etc. (recommended by systems folks)**
 - Zobel: Writing for computer science, Springer
 - Strunk & White: The elements of style, Allyn & Bacon
 - Dupré: Bugs in writing: A guide to debugging your prose, AddisonWesley

Something More.....

Annoying

- Terms used arbitrarily.
 - **Terms should be defined before used.**
- I will thoroughly review all spelling errors in the manuscript during its final stages, and kindly request that they be ignored now.
 - **NO! Fix those errors right away!**
- Use short sentences.
 - **Admit that you are not a native English speaker!**
 - The collaborators may not be able to see the structure of sentences clearly.
 - Many PCs are not native English speakers. Thus, even authors of native English speakers try to use simple English to help reviewers.

When is It ready to Submit A Paper?

- Rule: do not roll your dice!
- Think about what you get after submission.
 - Accepted! Lucky! Everything is fine.
 - Rejected! At least you need to get some **useful feedback**.
- If you have already known the weak points of your paper and also know what reviewers would reject because of these points, why are you submitting it? Fix them first!
- Wasting reviewers' time. **Wasting your time!**

Worry about Writing?

- I am not a native English speaker. My paper would be rejected because of writing.
- This is the least thing to worry about.
- It depends on different reviewers. But reviewers' tolerance for writing is much higher than you think.
- Papers don't need writing skills. They just need a clear expression of what you want to express.
- In the age of LLM, we have plenty of tools to improve the writing.
- The only thing you need to worry about is that if you can not express your ideas clearly in Chinese, you can not express them in English.

Wait, what? I can not speak Chinese?

- No, you can not!
- In fact, it is not because of what language you use. It is because of how you think about the problem.
- Skill: ask more why, ask less how.
- Think about alternatives. Why is your design reasonable? Why your choice is the only choice?
- The key to articulating problems and writing papers lies in thinking more about why, rather than how.

PC's Criteria

- There is no universal standard that defines what kind of papers can be accepted.
- Different reviewers have varying perspectives and evaluation criteria when assessing a paper.
- These questions are meaningless. “My score was higher than theirs, so why was their paper accepted while mine was rejected?”
- Papers without a PC “champion” have a hard stand. Make sure there’s something which at least one reviewer will think cool.

Hot/Popular vs. Unpopular Problems

- The advantage of working on a hot topic is that it can conveniently address the motivation section of the paper.
- For topics that are not hot or trending, more effort needs to be put into the motivation section to explain why it is still an interesting/important problem.
- Problems are much more important than solutions.

What to Do after Receiving a Rejection?

- I have an experience of getting a paper accepted after 3 years!
 - This is an extreme case. Usually, it is not the case.
 - Lamport got his Paxos paper published after 9 years.
 - Suggestion: maybe you need to submit your first paper as early as possible during your PhD time if you want to get your degree early.
- Evaluate the feedback: Some feedback **might be** constructive and helpful, while others could be subjective or less relevant. Identify which comments you agree with and which you might disagree with, but always be open to improvement.
- Discussing the feedback with a colleague or mentor.
- **Do not blindly adopt all the reviewers' comments. Avoid compromising the integrity and coherence of the paper.**

Paper is Bad without A Single Line of Formula

- For a system paper, formulas are meant to assist in explanations, not to be used just for the sake of having formulas.
- Same for algorithms.
- The use of terminology in the paper should be kept to a minimum. Abbreviations should also be minimized to reduce reading difficulties.
- Using a table to list the notations of terms appearing in the paper will provide a poor reading experience.

Evaluation

- Different papers/systems have different metrics for evaluation.
- You have to define the right metrics for your evaluation. Remember that the evaluation part should support the one (main) message of your paper.
- Reviewers will first assess whether your evaluation supports your claims in the paper.
- typical evaluations needed (often must)
 - End-to-end evaluation
 - Comprehensive benchmark evaluation (why do you omit YCSB-C and YCSB-D while evaluating other test sets?)
 - Breakdown

Conclusion

- Not the conclusion of a paper.
- The conclusion of this talk.
- Scientific writing is not just a way to communicate findings with others—**it is an essential tool for researchers themselves:**
 - **Writing forces researchers to organize their ideas logically and coherently.** This process can reveal gaps in reasoning, improve critical thinking, and help clarify complex concepts. By articulating their research in writing, researchers refine their understanding and presentation of their findings.
 - (In today's world), research problems have grown so complex that it's often not until we commit our thoughts to paper that we fully grasp the scope and nuances of our work. Writing, therefore, is not merely a final step in the research process; **it is a powerful mechanism for discovery and intellectual development.**

Thank you!
Ask me anything!

However, do not expect me to answer all questions.