# Seminar in Computer Architecture

## Lecture 1a: Course Logistics

Dr. Mohammed Alser

Prof. Onur Mutlu

ETH Zürich

Fall 2021

23 September 2021

# The Role of This Course

# Seminar in Comp Arch

- We will cover **fundamental** and **cutting-edge** research papers in computer architecture

- Multiple components that are aimed at improving students'
  - technical skills in computer architecture
  - critical thinking and analysis
  - technical presentation of concepts and papers
    - in both spoken and written forms
  - familiarity with key research directions

# Key Goal

(Learn how to)

rigorously

**analyze, present, discuss**

papers and ideas

in computer architecture

# Steps to Achieve the Key Goal

- Steps for the Presenter
  - Read
  - Absorb, read more (other related works)
  - Critically analyze; think; synthesize
  - Prepare a clear and rigorous talk
  - Present
  - Answer questions
  - Analyze and synthesize (in meeting, after, and at course end)

- Steps for the Participants
  - Discuss
  - Ask questions
  - Analyze and synthesize (in meeting, after, and at course end)

# Topics of Papers and Discussion

- hardware security;
- architectural acceleration mechanisms for key applications like machine learning, graph processing, and bioinformatics;
- memory systems;
- interconnects;
- processing inside memory;
- various fundamental and emerging ideas/paradigms in computer architecture;
- hardware/software co-design and cooperation;
- fault tolerance;
- energy efficiency;
- heterogeneous and parallel systems;
- new execution models, etc.

# Recap: Some Goals of This Course

- Teach/enable/empower you to:
  - Think critically
  - Think broadly
  - Learn how to understand, analyze, and present papers and ideas
  - Get familiar with key first steps in research
  - Get familiar with key research directions

# The Virtuous Cycle of Scientific Progress

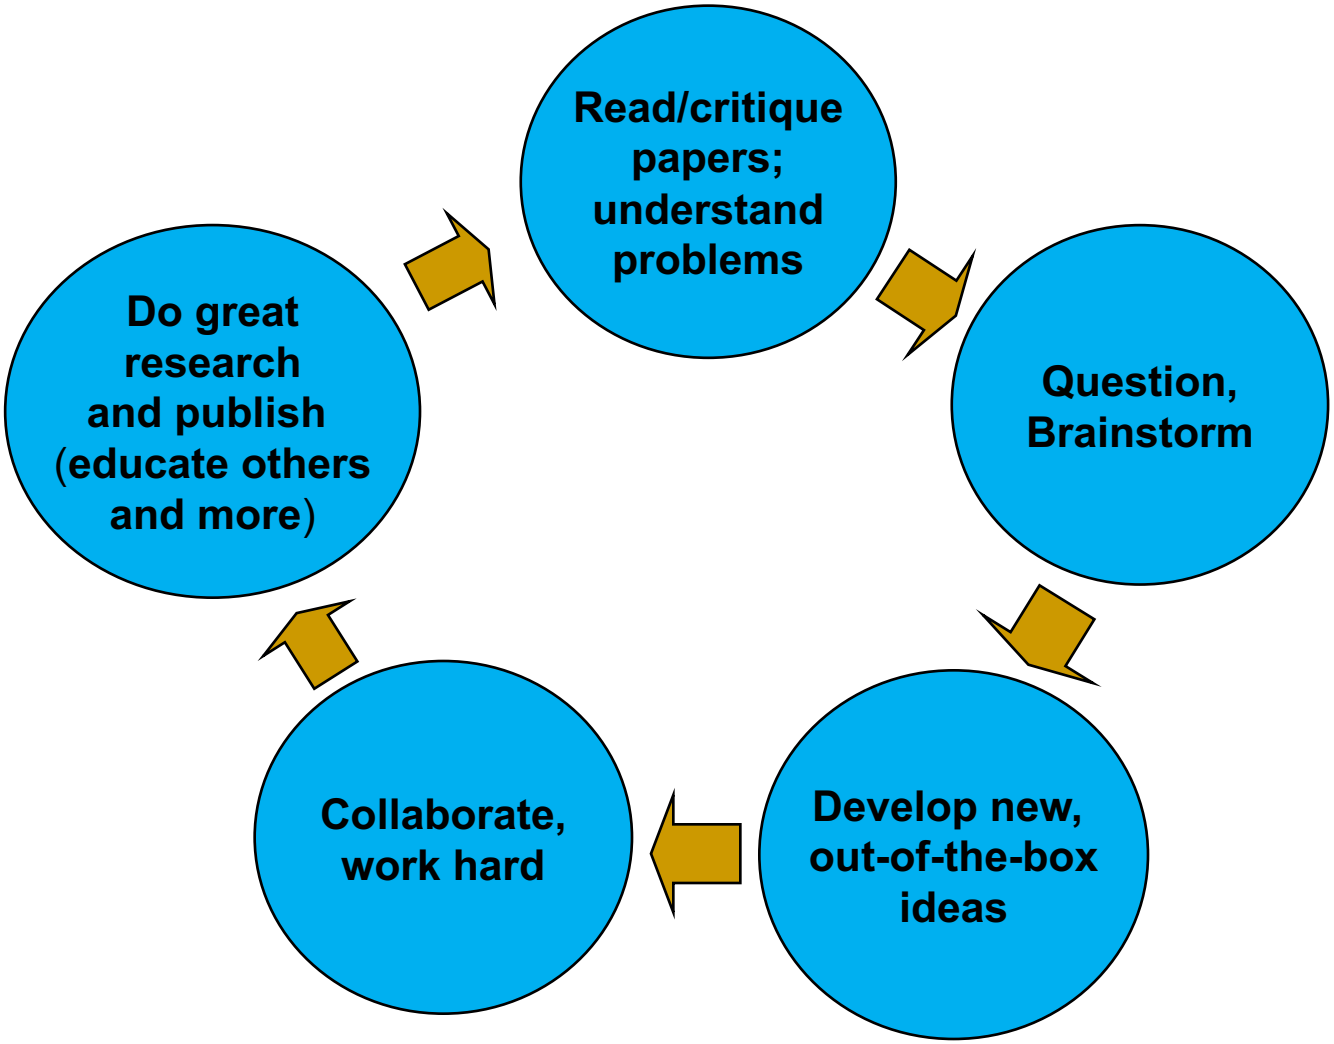

# Course Info and Logistics

# Course Info: Who Are We?

- **Onur Mutlu**
  - Full Professor @ ETH Zurich ITET (INFK), since September 2015
  - Strecker Professor @ Carnegie Mellon University ECE/CS, 2009-2016, 2016-…
  - PhD from UT-Austin, worked at Google, VMware, Microsoft Research, Intel, AMD
  - https://people.inf.ethz.ch/omutlu/
  - omutlu@gmail.com (Best way to reach me)
  - https://people.inf.ethz.ch/omutlu/projects.htm
- Research and Teaching in:
  - Computer architecture, computer systems, hardware security, bioinformatics
  - Memory and storage systems
  - Hardware security, safety, predictability
  - Fault tolerance
  - Hardware/software cooperation
  - Architectures for bioinformatics, health, medicine
  - …

# Lecturers

## Mohammed Alser

Senior Researcher and Lecturer

Bioinformatics | Computational Genomics | Hardware/Software Cooperation | Specialized/Heterogeneous Computing Systems | Processing-in-Memory

@mealser

## Juan Gómez Luna

Senior Researcher and Lecturer

Processing-In-Memory | Heterogeneous computing | Memory Systems | Bioinformatics | Medical imaging

# Course Info: Who Are We?

- **Teaching Assistants**
  - Dr. Jisung Park,
  - Dr. Lois Orosa Nogueira,
  - Dr. Gagandeep Singh,
  - Dr. Nour Almadhoun Alserr,
  - Dr. Haiyu Mao,
  - Dr. Behzad Salami,
  - Dr. Mohammad Sadr,
  - Hasan Hassan,
  - Can Firtina,
  - Geraldo Francisco De Oliveira Junior,
  - Abdullah Giray Yaglikci,
  - Rahul Bera,
  - Konstantinos Kanellopoulos,
  - Nika Mansouri Ghiasi,
  - Rakesh Nadig,
  - João Dinis Ferreira,
  - Roknoddin Azizibarzoki

- Get to know them and their research as they will be your mentors
  https://safari.ethz.ch/group-members/

# Course Requirements and Expectations

- Attendance required for all meetings
- Each student presents one paper
  - Prepare for presentation with engagement from the mentor
  - Full presentation + questions + discussion
- Non-presenters participate during the meeting
  - Ask questions, contribute thoughts/ideas
  - Better if you read/skim the paper beforehand
- Non-presenters take an online short quiz after each session
  - 5 MCQs for each presentation (4 hours to submit)
- Everyone comments on papers in the online review system
  - After presentation
- Write synthesis report at the end of semester
  - (sample synthesis report online)

# Course Website

- https://safari.ethz.ch/architecture_seminar/fall2021
- All course materials to be posted
- Plus other useful information for the course
- Check frequently for announcements and due dates

# Moodle

- https://moodle-app2.let.ethz.ch/course/view.php?id=15535

- Check frequently for:
  - Announcements and due dates
  - Quizzes
  - Paper discussion
  - Assignment submissions

# Homework 0

- Due September 30
  - https://safari.ethz.ch/architecture_seminar/fall2021
- Information about yourself
- All future grading is predicated on homework 0
- If it is not submitted on time, we cannot schedule you for a presentation.

# Paper Review Preferences

- Due TBD
- Check the website and Moodle for instructions
- If it is not submitted on time, we cannot schedule you for a presentation.

# How to Deliver a Good Talk

# Anatomy of a Good Paper Review (Talk)

- 0: Title, Authors, Venue
- 1: Summary
  - What is the problem the paper is trying to solve?
  - What are the key ideas of the paper? Key insights?
  - What are the key mechanisms? What is the implementation?
  - What are the key results? Key conclusions?
- 2: Strengths (most important ones)
  - Does the paper solve the problem well? Is it well written? …
- 3: Weaknesses (most important ones)
  - This is where you should think critically. Every paper/idea has a weakness. This does not mean the paper is necessarily bad. It means there is room for improvement and future research can accomplish this.
- 4: Thoughts/Ideas: Can you do better? Present your ideas.
- 5: Takeaways: What you learned/enjoyed/disliked? Why?
- 6: Discussion starters and questions.
- Review should be short and concise (20 minutes or < one page)

# Suggested Paper Discussion Format

- Problem & Goal
- Key Ideas/solution
- Novelty
- Mechanisms & Implementation
- Major Results
- Takeaways/Conclusions

**~20-25 minute Summary**

- Strengths
- Weaknesses
- Alternatives
- New ideas/problems
- Brainstorming and Discussion

**~10 min Critique plus ~15 min Discussion**

# More Advice on Paper Review Talk

- When doing the paper reviews and analyses, be very critical
- Always think about better ways of solving the problem or related problems
  - Question the problem as well
  - Read background papers (both past and future)
- This is how things progress in science and engineering (or anywhere), and how you can make big leaps
  - By critical analysis
- A few sample text reviews provided online

# Try to Avoid Rat Hole Discussions

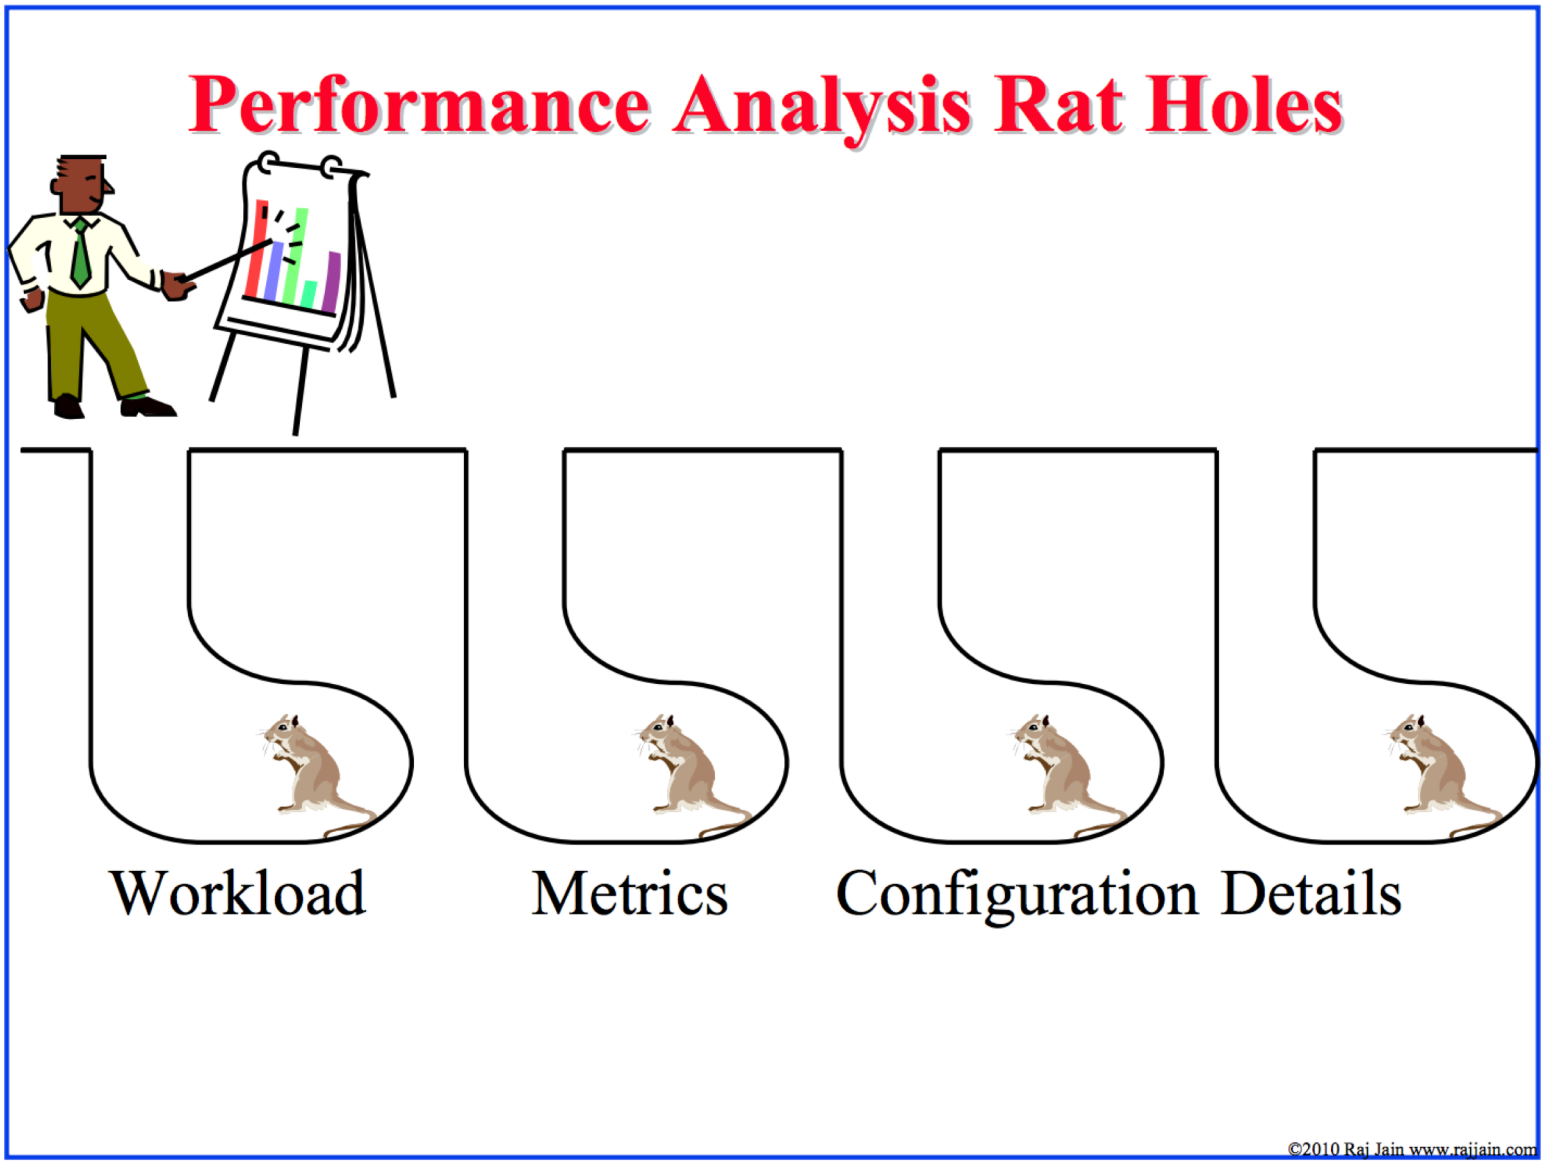


Source: https://www.cse.wustl.edu/~jain/iucee/ftp/k_10adp.pdf

# Aside: A Recommended Book

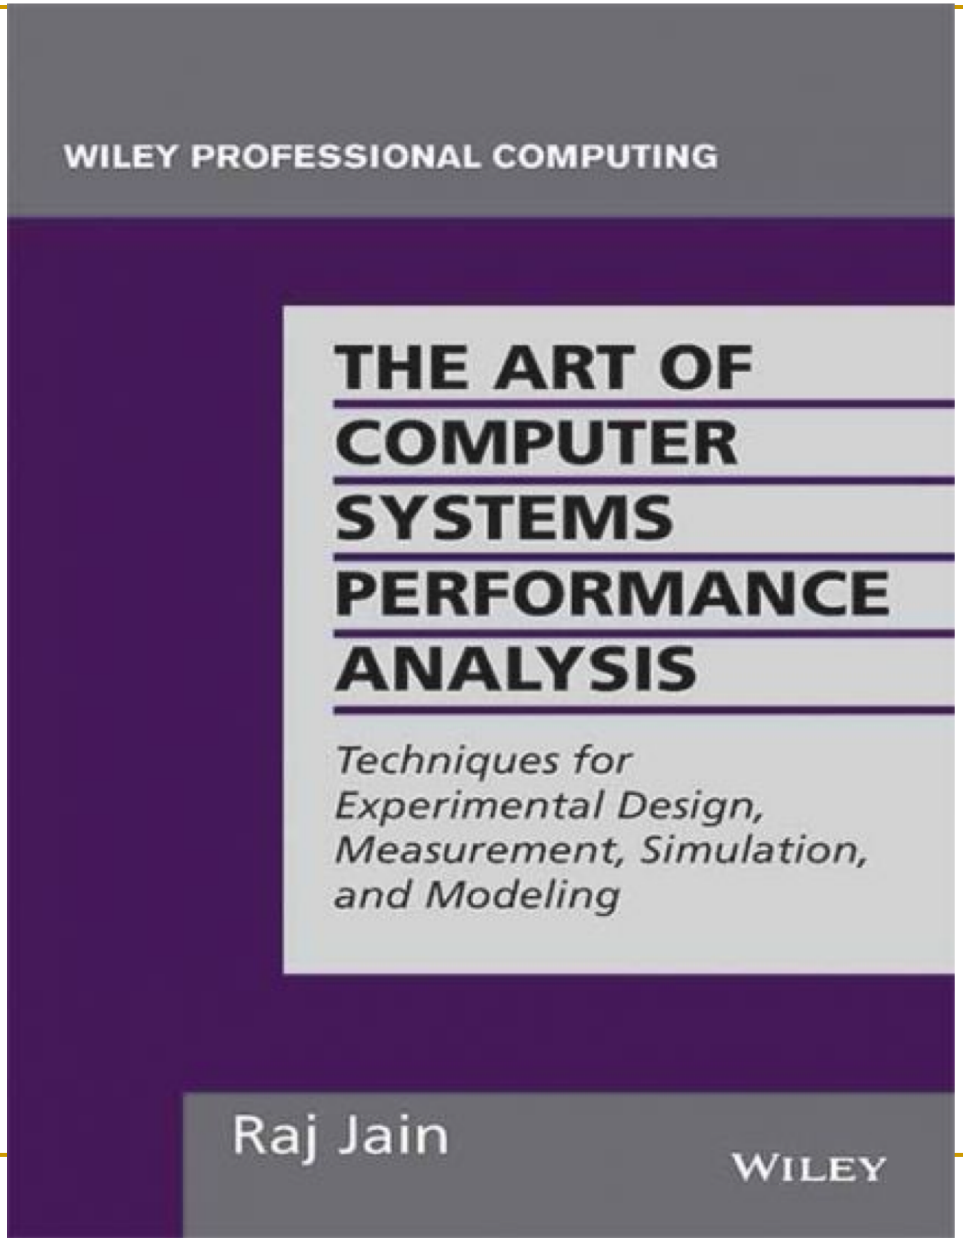


Raj Jain, "The Art of Computer Systems Performance Analysis," Wiley, 1991.

## 10.8 DECISION MAKER'S GAMES

Even if the performance analysis is correctly done and presented, it may not be enough to persuade your audience—the decision makers—to follow your recommendations. The list shown in Box 10.2 is a compilation of reasons for rejection heard at various performance analysis presentations. You can use the list by presenting it immediately and pointing out that the reason for rejection is not new and that the analysis deserves more consideration. Also, the list is helpful in getting the competing proposals rejected!

There is no clear end of an analysis. Any analysis can be rejected simply on the grounds that the problem needs more analysis. This is the first reason listed in Box 10.2. The second most common reason for rejection of an analysis and for endless debate is the workload. Since workloads are always based on the past measurements, their applicability to the current or future environment can always be questioned. Actually workload is one of the four areas of discussion that lead a performance presentation into an endless debate. These "rat holes" and their relative sizes in terms of time consumed are shown in Figure 10.26. Presenting this cartoon at the beginning of a presentation helps to avoid these areas.

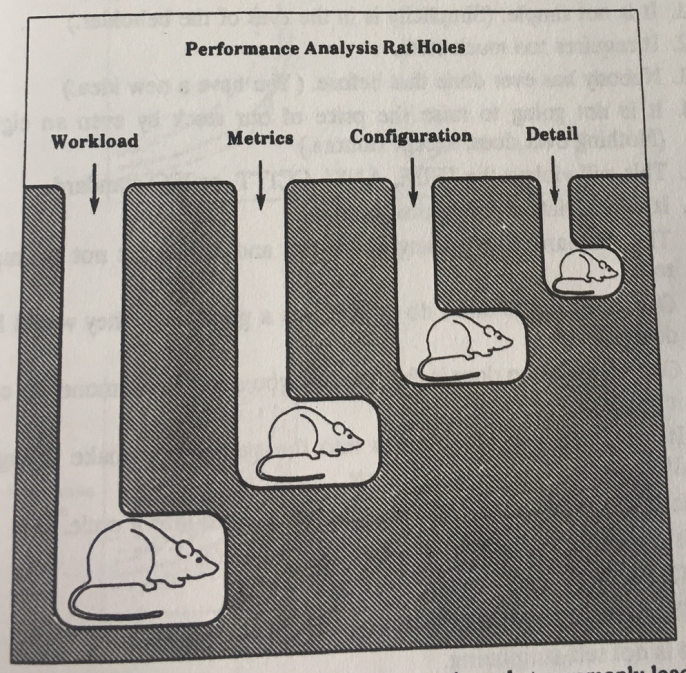


**FIGURE 10.26** Four issues in performance presentations that commonly lead to endless discussion.

Raj Jain, "The Art of Computer Systems Performance Analysis," Wiley, 1991.

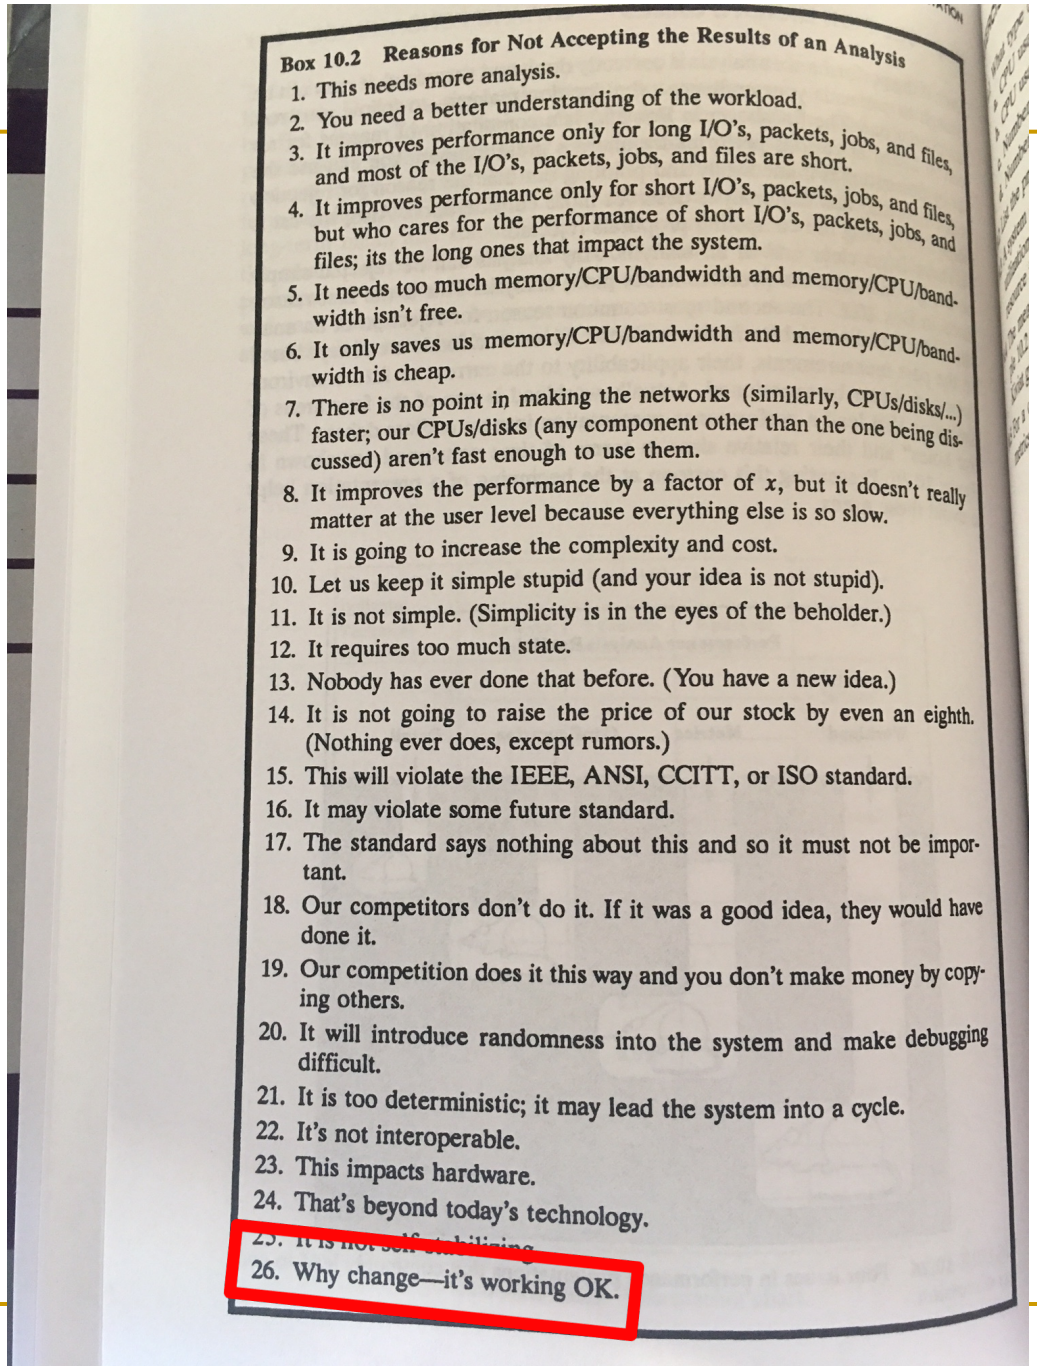

**Box 10.2 Reasons for Not Accepting the Results of an Analysis**

1. This needs more analysis.
2. You need a better understanding of the workload.
3. It improves performance only for long I/O's, packets, jobs, and files, and most of the I/O's, packets, jobs, and files are short.
4. It improves performance only for short I/O's, packets, jobs, and files, but who cares for the performance of short I/O's, packets, jobs, and files; its the long ones that impact the system.
5. It needs too much memory/CPU/bandwidth and memory/CPU/bandwidth isn't free.
6. It only saves us memory/CPU/bandwidth and memory/CPU/bandwidth is cheap.
7. There is no point in making the networks (similarly, CPUs/disks/...) faster; our CPUs/disks (any component other than the one being discussed) aren't fast enough to use them.
8. It improves the performance by a factor of $x$, but it doesn't really matter at the user level because everything else is so slow.
9. It is going to increase the complexity and cost.
10. Let us keep it simple stupid (and your idea is not stupid).
11. It is not simple. (Simplicity is in the eyes of the beholder.)
12. It requires too much state.
13. Nobody has ever done that before. (You have a new idea.)
14. It is not going to raise the price of our stock by even an eighth. (Nothing ever does, except rumors.)
15. This will violate the IEEE, ANSI, CCITT, or ISO standard.
16. It may violate some future standard.
17. The standard says nothing about this and so it must not be important.
18. Our competitors don't do it. If it was a good idea, they would have done it.
19. Our competition does it this way and you don't make money by copying others.
20. It will introduce randomness into the system and make debugging difficult.
21. It is too deterministic; it may lead the system into a cycle.
22. It's not interoperable.
23. This impacts hardware.
24. That's beyond today's technology.
25. It is not self-stabilizing.
26. Why change—it's working OK.

Raj Jain, "The Art of Computer Systems Performance Analysis," Wiley, 1991.

# More Advice on Talks

- Kayvon Fatahalian, "Tips for Giving Clear Talks"
  - http://graphics.stanford.edu/~kayvonf/misc/cleartalktips.pdf
  - Many useful and simple principles here

**"Every sentence matters"**

**"The audience prefers not to think" (about things you can just tell them)**

**"Surprises are bad": say why before what**
**(indicate why you are saying something before you say it)**

**Explain every figure, graph, or equation**

**When improving the talk, the audience is always right**

# Who Painted This Painting?

Salvador Dali @ 1924

Source: http://www.museoreinasofia.es/en/collection/artwork/retrato-luis-bunuel-portrait-luis-bunuel http://salvadordaliprints.org/biography/

# What About This?

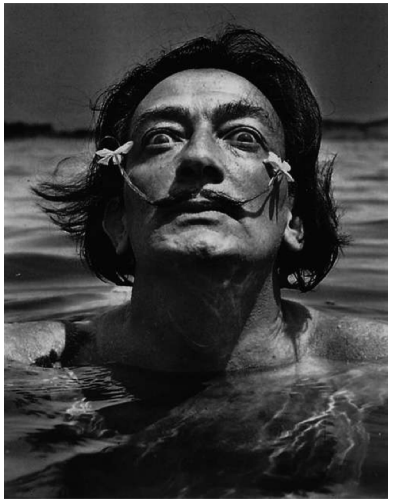

Salvador Dali @ 1937

Source: http://salvadordaliprints.org/the-burning-giraffe/ https://www.dalipaintings.com/

# Takeaway

Learn the basic principles

**before** you can

*consciously* choose to break them

# Seminar Talk Examples

# Past Seminar Presentation

**Sofie Daniels**
[BlockHammer: Preventing RowHammer at Low Cost by Blacklisting Rapidly-Accessed DRAM Rows, HPCA 2021,](#)
**Session 1.2: Seminar in Computer Architecture - Spring 2021,**
[Talk Video] (34 minutes excluding discussion)
[Slides (pptx) (pdf)]

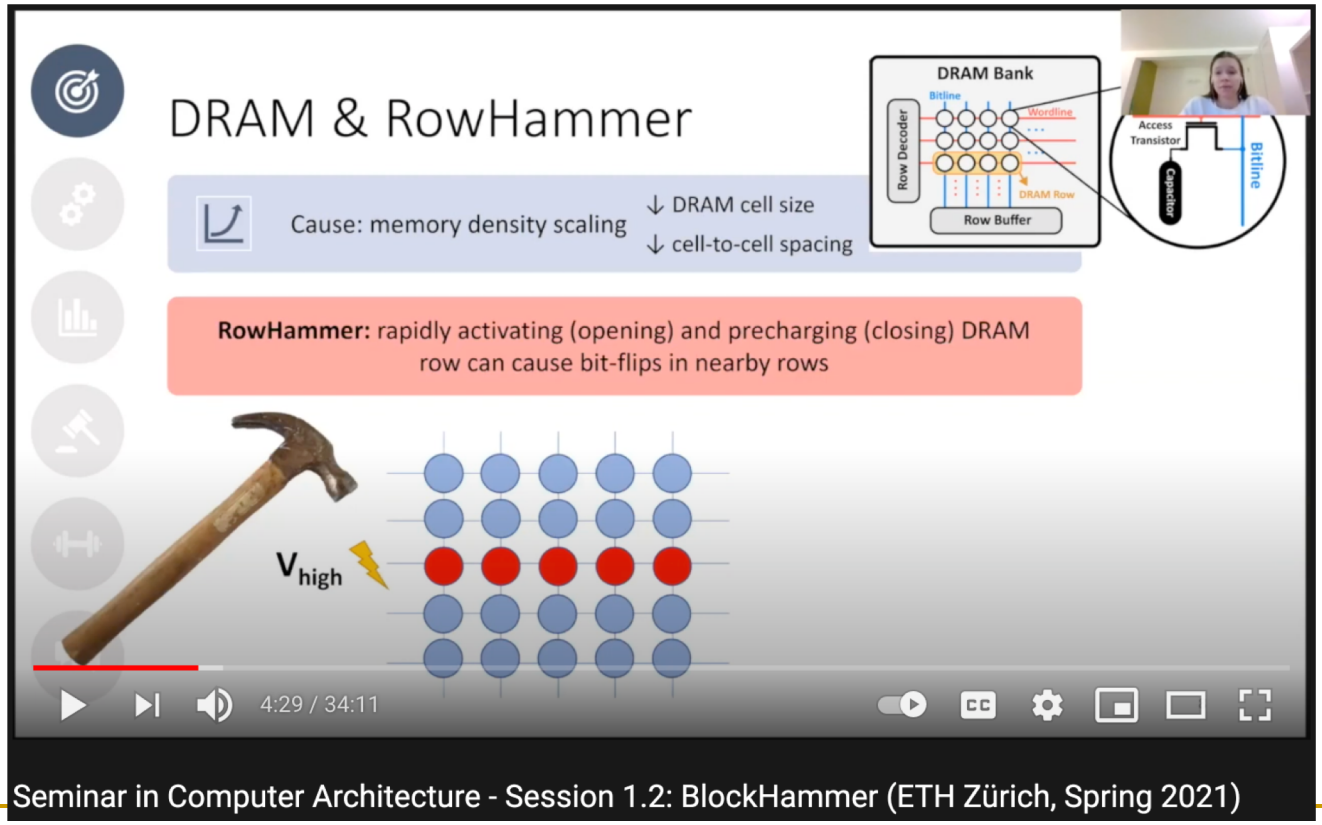

# Past Seminar Presentation and Discussion

**Lorenzo Rai**
ComputeDRAM: In-Memory Compute Using Off-the-Shelf DRAMs, MICRO 2019
**Session 1.2: Seminar in Computer Architecture - Spring 2021,**
[Talk Video] (62 minutes including discussion)
[Slides (pptx) (pdf)]

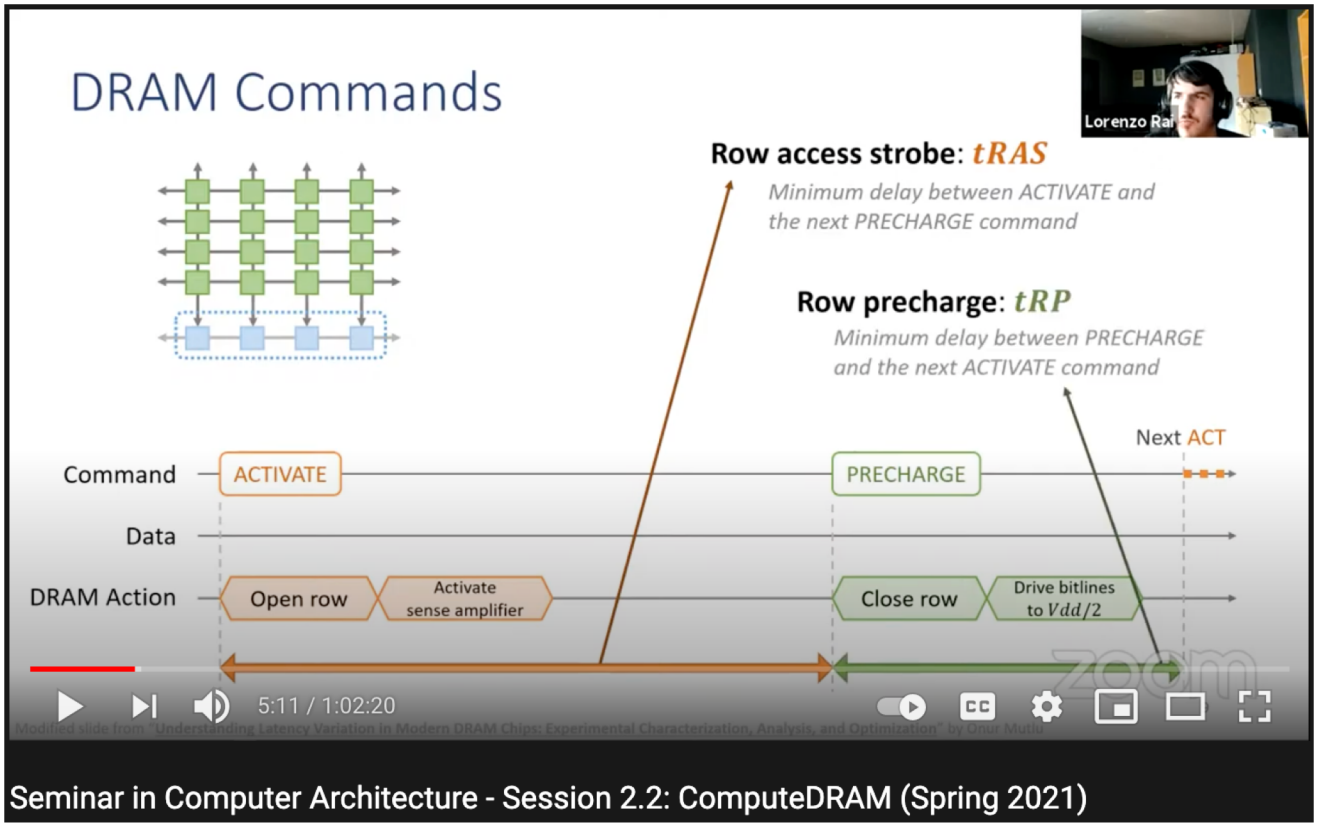

# Past Example Reviews

[List of example reviews]:
https://safari.ethz.ch/architecture_seminar/spring2021/doku.php?id=schedule
[List of YouTube videos]:
https://www.youtube.com/playlist?list=PL5Q2soXY2Zi_awYdjmWVIUegsbY7TPGW4

## Spring 2021 Lectures/Schedule

| Week | Date | Livestream | Lecture | Readings | Assignments |
|---|---|---|---|---|---|
| W2 | 04.03 Thu. | YouTube Live | L2: Example Review: RowClone (PDF) (PPT) | Suggested | |
| W3 | 11.03 Thu. | YouTube Live | L3: Example Review: Memory Channel Partitioning (PDF) (PPT) | Suggested | |
| W4 | 18.03 Thu. | YouTube Live | L4: Example Review: GateKeeper (PDF) (PPT) | Suggested | |

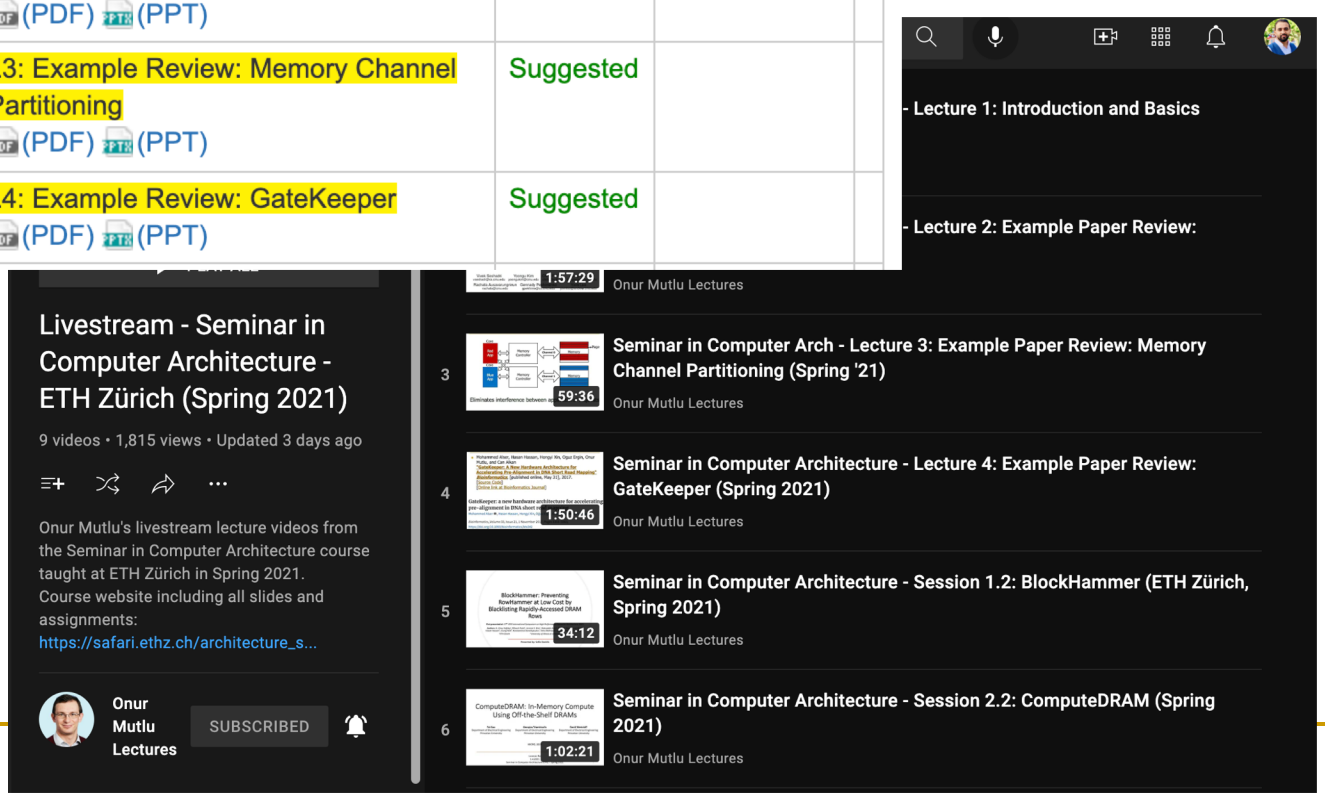

# Past Example Reviews

[List of ex
https://sa
[List of Yo
https://ww

## Spring 2

| Week | Date | |
|---|---|---|
| W2 | 04.03 Thu. | |
| W3 | 11.03 Thu. | |
| W4 | 18.03 Thu. | |

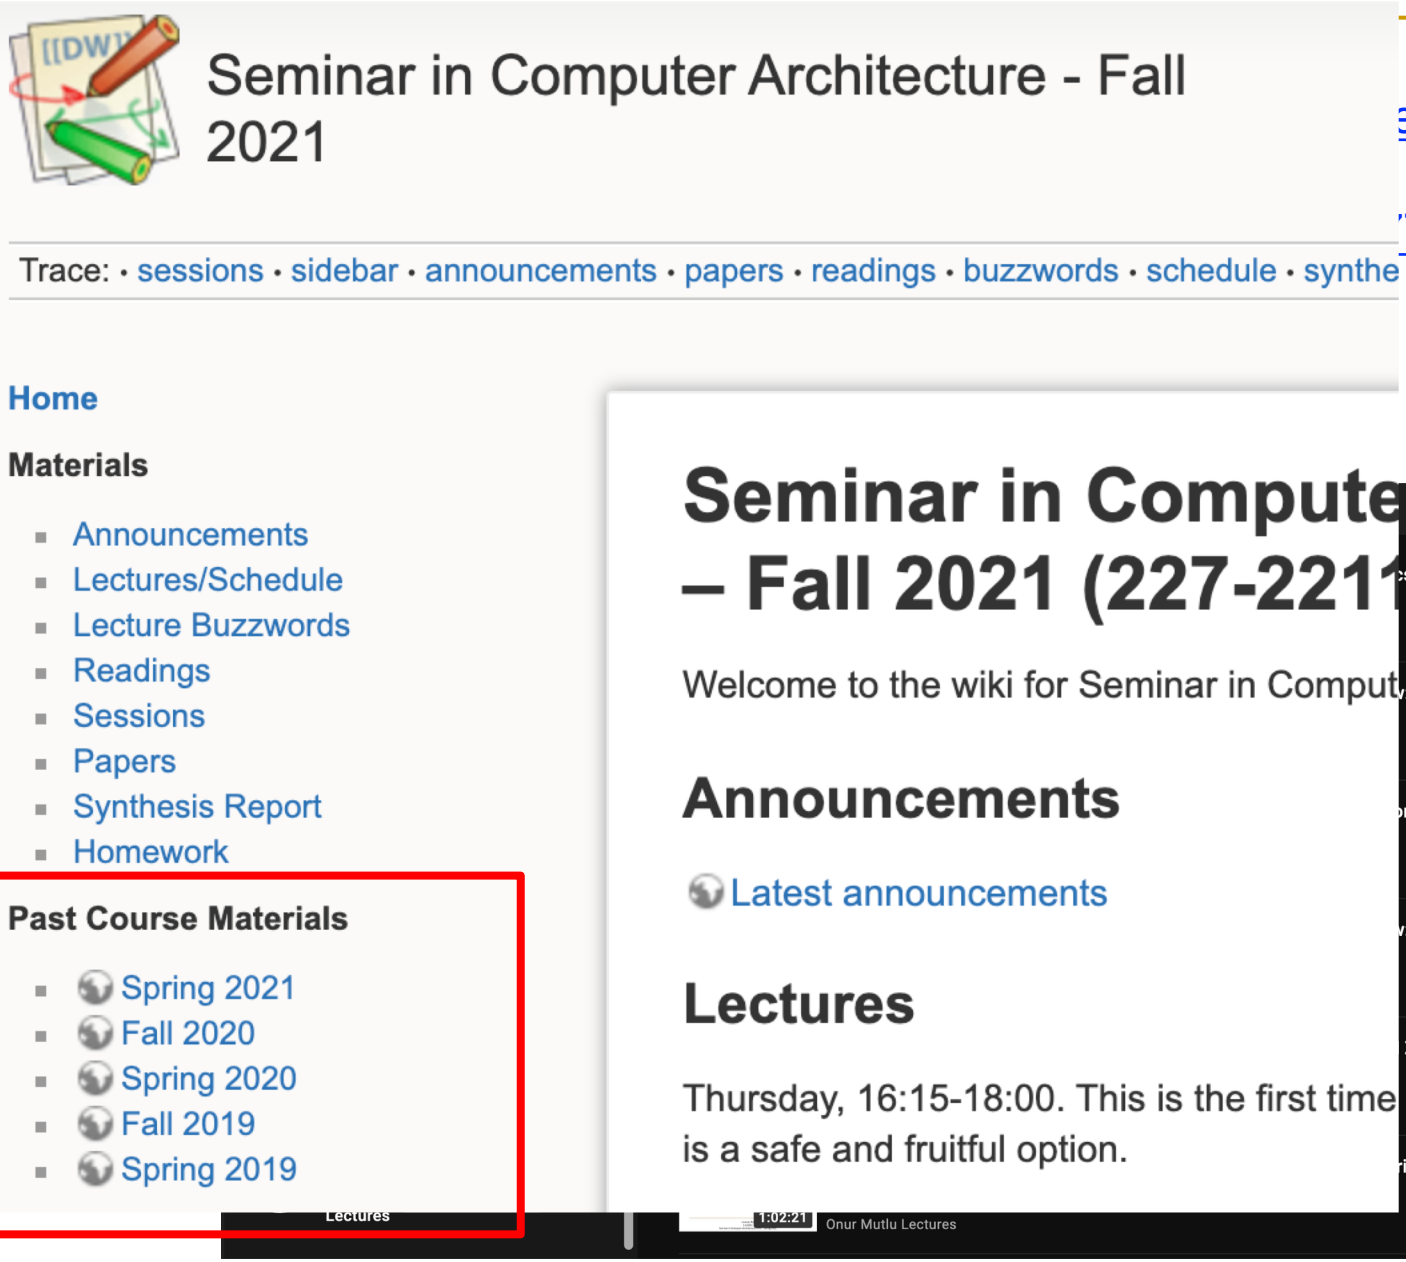

# How to Participate

# How to Make the Best Out of This?

- Come prepared → Read and critically evaluate the paper
- Think new ideas
- Bring discussion points and questions; read other papers
- Be critical
- Brainstorm – be open to new ideas
- Pay attention and discuss+contribute
- Participate online before and after each meeting

# Guided Talk Preparation

# Preparing a Talk

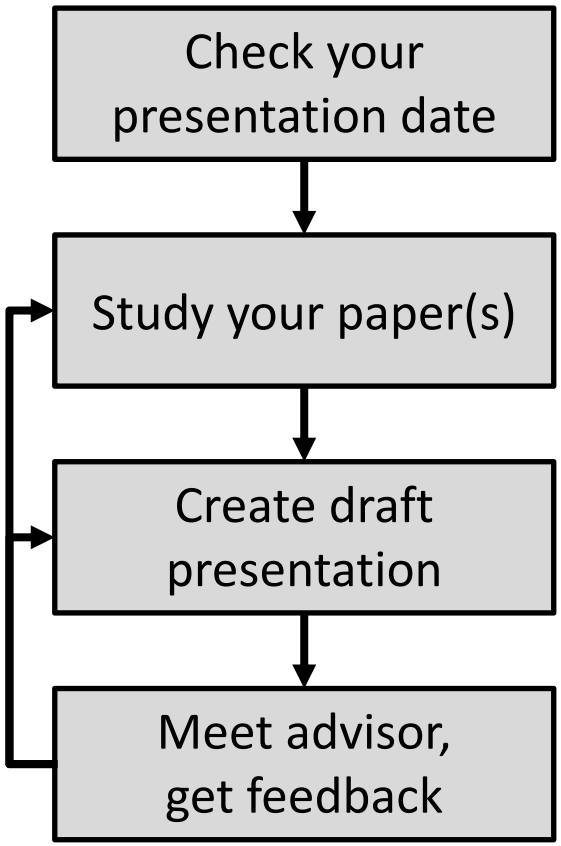

# Preparing a Talk: Start Early

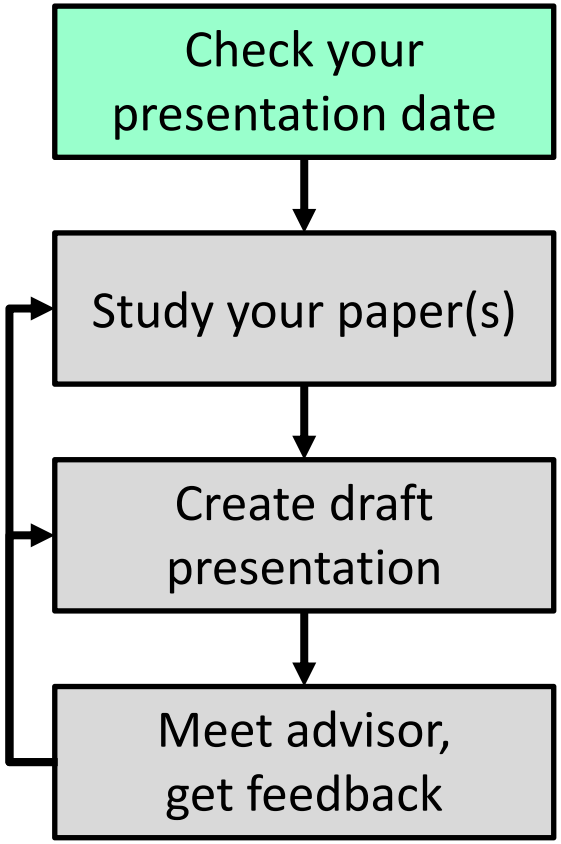


- Preparing a good presentation takes time
- Start early!

# Preparing a Talk: Study Paper

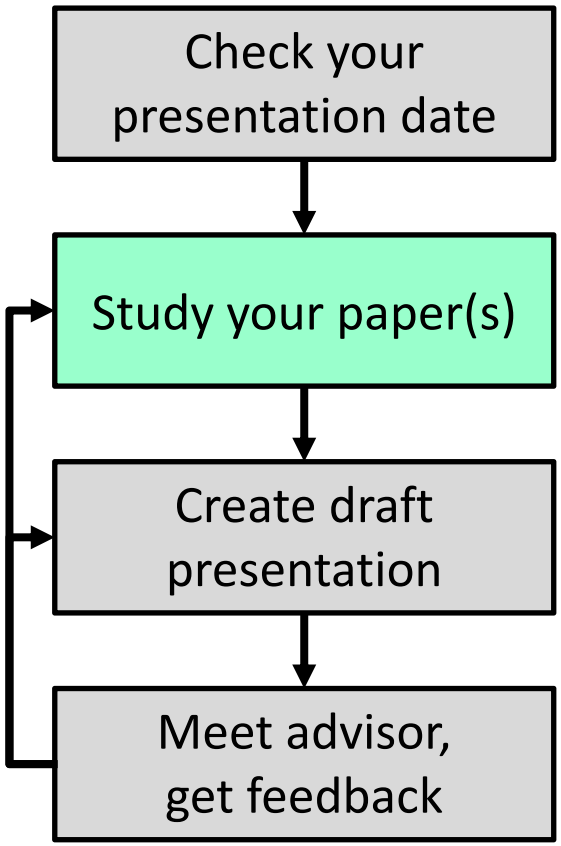


- 3 'C's of reading
  - *Carefully:* look up terms, possibly read cited papers
  - *Critically:* find limitations, flaws
  - *Creatively:* think of improvements
- Try examples by hand
- Try tools if available

- Consult with TA if questions

# Preparing a Talk: Create Draft

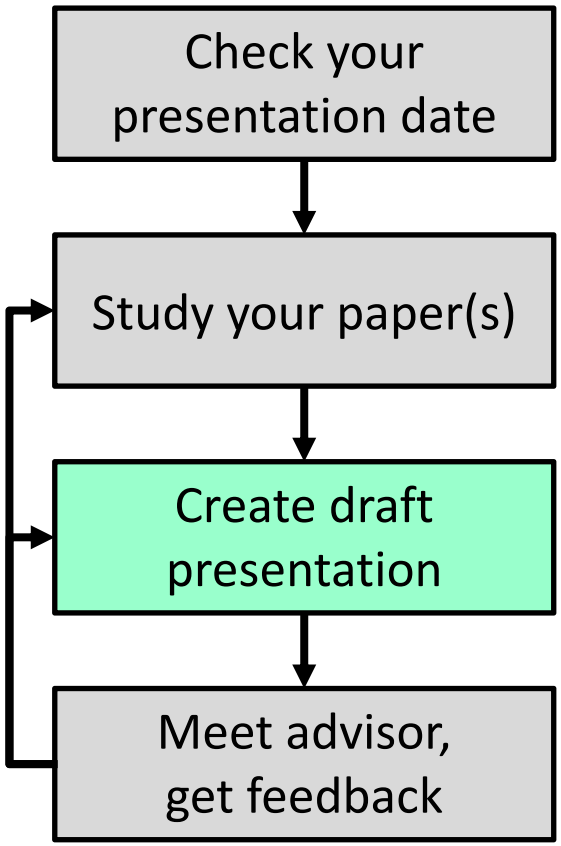


- Explain the motivation for the work
- Clearly present the technical solution and results
  - Include a demo if appropriate
- Outline limitations or improvements
- Focus on the key concepts
  - Do not present all of the details

# Preparing a Talk: Get Feedback

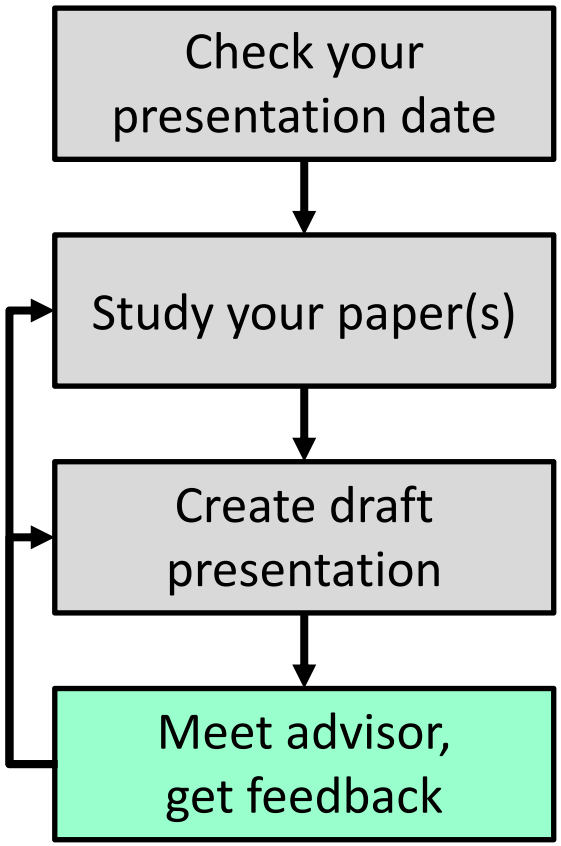


- Prepare for the meeting
  - Schedule early
  - Send slides in advance
  - Write down questions
- Make sure you address feedback
  - Take notes
- Meetings are mandatory!
  - At least one week before the talk
  - Two meetings

# Grading and Feedback

# Grading Rubric

- Quality of your presentation (60%)
  - How well did you understand the material?
  - How well did you present it?
  - How well did you answer the questions?
  - Be prepared to explain technical terms
  - We will take into account the difficulty of the paper and the time you had to prepare.
- Quality of the final synthesis paper (30%)
  - How well did you understand some of the papers presented during the seminar?
- Attendance & Quizzes (10%)
- Participation (during class and online) (BONUS 10%)
  - Did you ask good questions?
  - Did you participate and contribute to the discussion?

# Feedback

- We will try to (briefly) discuss strengths/weaknesses of your talk in class
  - Let us know upfront if you would prefer **not** to
- You can arrange a meeting with your TA to get feedback

# Expected Schedule

# Schedule

- We will meet once a week, with two presentations per session
  - *Next meeting next week*
  - *Your presentations start on 7 October*
  - 27 presentations in total
  - Each presentation 50 minutes including questions and discussion

- Paper assignment
  - Will be done online
  - Study the list of papers
  - Check your email and be responsive

# Homework 0

- Due September 30
  - https://safari.ethz.ch/architecture_seminar/fall2021
- Information about yourself
- All future grading is predicated on homework 0
- If it is not submitted on time, we cannot schedule you for a presentation.

# Paper Review Preferences

- Due TBD
- Check the website and Moodle for instructions
- If it is not submitted on time, we cannot schedule you for a presentation.

# Seminar in Computer Architecture

## Lecture 1a: Course Logistics

Dr. Mohammed Alser

Prof. Onur Mutlu

ETH Zürich

Fall 2021

23 September 2021