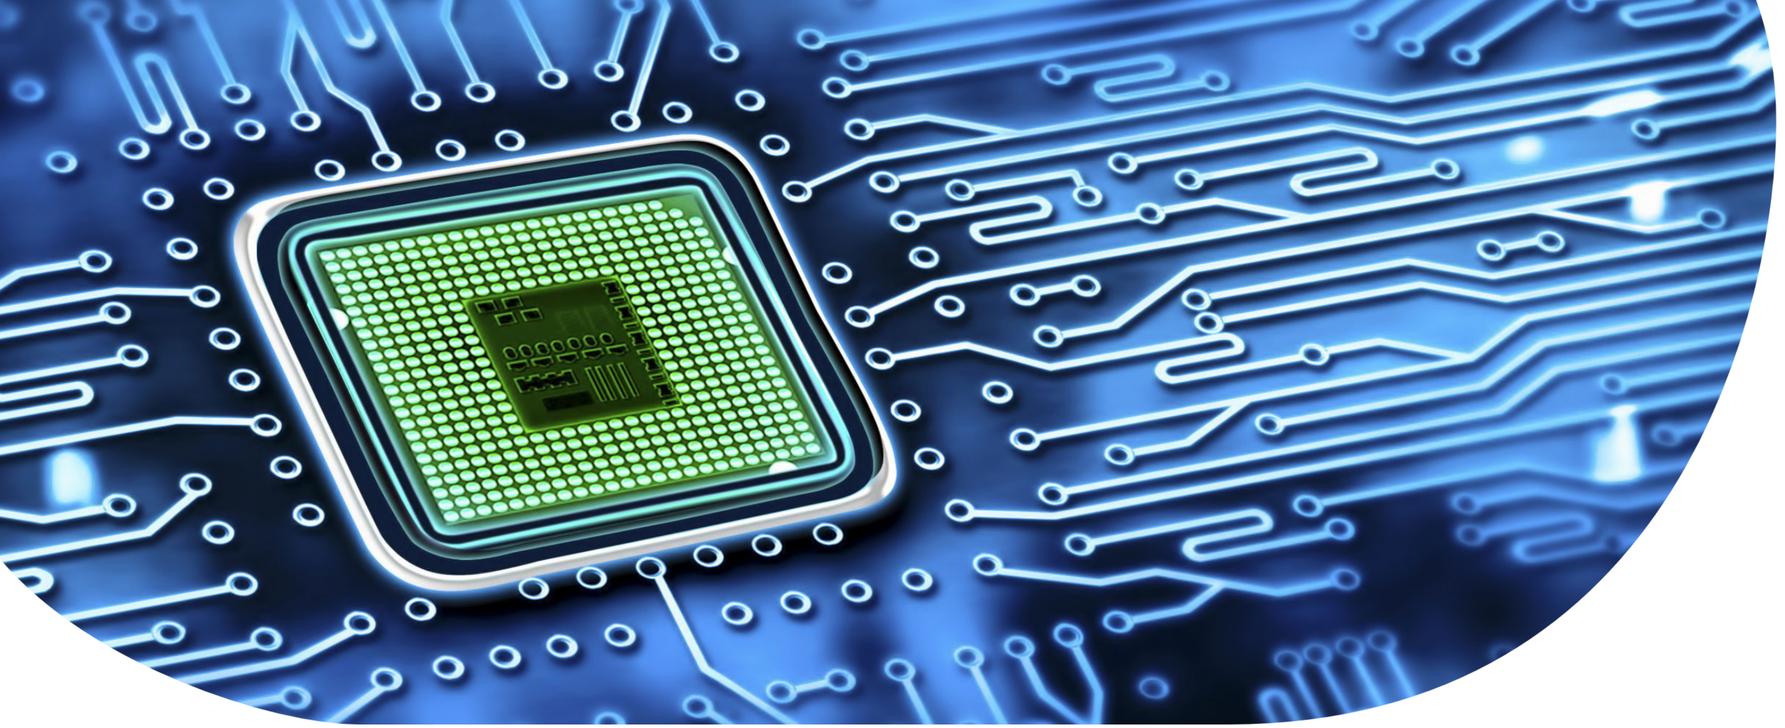

جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

CS356 Hardware Accelerator Architectures

Reading and Reviewing Papers

Suhaib Fahmy
suhaib.fahmy@kaust.edu.sa

Slides based on “How to Recognize a Great Idea” by Nick Feaster and Alex Gray

Discussion

- Provide one example of a great technical idea from your research area
- Explain why you think this idea is great

Questions about this idea

- Did it open up new avenues for research?
- Did it change your way of thinking?
- Did it frame a problem more clearly?
- Did it enable new techniques or capabilities?
- Does it affect practice?
- Has it been applied across disciplines?
- How old is the idea? (Did it stand the “test of time”?)

More Questions

- Is it original?
- Does it challenge or question assumptions?
- Does it resolve longstanding issues?
- Is it a large “leap” forward?

How Do We Measure “Impact”?

- In some ways, impact or contribution is subjective.
- As time passes, it is easier to see.
  - Sometimes we can predict these trends (“time travel”)
- To judge, must know the context, and what existed before.
- Occasionally, we look to the leaders of the field to declare impact.
  - …but sometimes they are the defenders of the old ideas.

Quantitative Metrics

- Number of citations
- Downloads of code/number of projects using the code
- Unfortunately, these are imperfect. Why?

The Problem and the Solution Both Matter

- Can take many forms:
  - solves a small problem, much better than before…
  - solves a major problem, slightly better…
- The size of the problem is at least as important as the goodness of the solution

How to Judge a Paper?

- Key question: Does the paper have a great idea?

- Other questions
  - Is it clearly presented?
  - Is the idea evaluated credibly?
  - Do the authors demonstrate command of the related literature and work?

Invariant Questions

- Is the problem important?
- What is the **“intellectual nugget”**?
- What is the main conclusion? Is it important?
- Does the content support the conclusion?
  - Are the methods sound and state-of-the-art?
  - Are the results likely to be affected by the method?
  - If the assumptions in the paper change, will the conclusions hold?

Simplicity and Elegance

- Is simplicity good or bad?
- What constitutes a “hack”?
- What does it mean for research to be “incremental”?
- Is math (in a paper) good or bad?

Why Read Research Papers?

- Read for a conference or a class
- Keep current in your own field
- Get “up to speed” in a new field
  - Learn about a sub-field (e.g., wireless)
  - Learn about another discipline that may offer solutions to a problem

Step 1: Deciding What to Read

- Purpose: Learn about “hot topics” of current research in an area. (searching for problems, etc.)
  - Approach: Scan papers in latest conference proceedings
- Purpose: Get up to speed on sub-field
  - Approach: Transitive closure of related work of papers in a top conference
- Purpose: Learn about an area that is further afield
  - Approach: Ask expert colleagues

Step 2: Deciding How to Read

- Always “top down”
  - First: Abstract, introduction, conclusion
  - Rest of paper if necessary
    - If you want to do follow-up research
    - If you want to better understand the methods/conclusions
- Next steps depend on specific purpose

Reading the News

- Conference proceedings
  - Goal: Grasp main idea of a collection of a large number of papers. Keep informed about problems and recent solutions
- Top-Down Method
  - Skim table of contents: Papers are clustered into “sessions” which typically identify the main areas
  - Consider authors
  - Prioritize by (1) area of interest (2) reputable authors

How to Conduct a Literature Survey

- Create the seed
  - Recent paper from top conference
  - Survey paper, if one exists
  - Seminal paper, if it is different from the above
- Perform transitive closure of cited work
  - Read related work sections of above papers

Keeping Notes

- One-sentence summaries are infinitely better than nothing at all
- Primitive approach: Single file of notes
- Better: Database with BibTeX
  - Use tools for bibliography management like Zotero, Mendeley
  - Will also help you more quickly construct related work sections for your papers
- Always sanitize your BibTeX!

How to Review a Paper

- What are the differences between reading and reviewing?
  - Reading: information gathering, typically for the benefit of your own research benefit (You are a scientist.)
  - Reviewing: goal is to (1) determine a paper's suitability for some conference (2) provide feedback to authors to improve paper (You are a teacher/evaluator.)

Consider the Audience

- Will this generate discussion?
- Is this a paper that's going to send people to the hallway?
- Will the people who commonly read these proceedings benefit from the contributions?
  - Would people who read other proceedings benefit more from the paper?

Consider the Standards

- Workshops are typically more permissive as far as accepting “vision” without completed, supported work
  - More emphasis on “fostering discussion”
- Conference: Depends on quality of papers in the reviewers’ piles and selectivity
- Journals often have the highest standards, especially since the review process is iterative

Consider the Purpose

- Survey
  - Is the overview complete?
- Tutorial
  - Is the description correct and clearly described?
- Proposal
  - Does the research agenda that is advocated make sense? Is it worthwhile?

How to Write the Review Itself

- Start with a summary
  - Demonstrates to the authors (and to you!) that you understand the main point of the paper
- Discuss how authors do or do not deliver on the claims/contributions of paper
- Discuss positive aspects (if any)…try to find something
- Provide high-level suggestions for improvement
- End with nits (spelling, punctuation, etc.)

Sample Categories from a Real Review Form

- Reviewer confidence
- Summary
- Novelty
- Clarity
- Relevance (“Scope”)
- Strengths / Weaknesses

Summarising

- Summarize the paper in a few sentences. Try to address these questions:
  - What type of paper is it?
  - What is the context for this paper?
  - Is it correct?
  - What are its contributions?
  - Is it comprehensible?

Detailed Comments

- Please try to be positive in your reviews: for instance, instead of saying "Your scheme is bad," consider saying "Your scheme would be stronger if it dealt with case X."
- If you think that the paper should cite prior work, please give a full citation to the work that should be cited. Please address any of the following that apply:
- Incorrect assumptions
- Insufficient evaluation
- Instances where the solution may not work correctly
- Portions of the paper that you found hard to read or understand

Detailed Comments

- Whether the focus of the work is too narrow, leading to incremental gains
- Whether the proofs are correct
- Whether the statistical analysis is correct
- Whether the claims made match the contributions
- Whether the authors use an appropriate data set
- Whether the system leaves out important components
- Whether the solution is deployable.

Scoring Criteria

- Impact
- Novelty
- Clarity
  - Problem stated clearly
  - Solution and evidence for its quality (e.g. experiments) stated clearly, ideally reproducible
  - Novelty stated clearly
- This is how your work will be “scored”, and how you’ll score others’ work.

First Step: Read and Re-Read

- Read the paper once to get the main ideas and contributions
  - Try to make a “one bit” decision here
- Read again and take notes (for your review)
- Start to organize a review

General Tips on Tone and Content

- Be polite and respectful
- Provide suggestions for how to improve the paper
  - You may see the paper again!
  - If the paper is accepted, the flaws should be fixed
- Be positive
- The point is **not** to shoot the paper down

Common Mistake: Being Too Critical

- Don’t miss forest for the trees!
  - Papers are never perfect
  - Your job is to determine whether a paper’s flaws invalidate the contributions (and whether the contributions are significant)
- Being too critical can prevent important research results from being published

Other mistakes and no-nos

- Insulting the authors
  - Criticize the paper, not the authors: “The paper did not address…”
- Revealing your own research agenda
- Distributing submitted papers
- Spending too much time reviewing a paper
  - Rule of thumb: Don’t spend more time reviewing a paper than the authors did writing it!
  - If a paper is sloppy or flawed, don’t waste your time

Build a Reading Habit!

- As a graduate student, you should be reading papers regularly
  - Create a reading list that you add you whenever you come across a relevant/interesting paper
  - Separate those directly relevant to your research and those that are a curiosity
- Assign some time in your week where you work through this reading list
- Regularly check new proceedings and journal issues