

18-740: Computer Architecture

Recitation 1: Introduction, Logistics, and Jumping Into Research

Prof. Onur Mutlu

Carnegie Mellon University

Fall 2015

September 1, 2015

447 Strikes Back

- Who has taken 447?
- Do you remember the first assignment?
- Have you done it?

Let's Recall

Question: What Is This?




Answer: Masterpiece of A Famous Architect

Fallingwater

From Wikipedia, the free encyclopedia

Fallingwater or **Kaufmann Residence** is a house designed by architect [Frank Lloyd Wright](#) in 1935 in rural [southwestern Pennsylvania](#), 43 miles (69 km) southeast of [Pittsburgh](#).^[4] The home was built partly over a waterfall on [Bear Run](#) in the Mill Run section of [Stewart Township, Fayette County, Pennsylvania](#), in the [Laurel Highlands](#) of the [Allegheny Mountains](#).

Time cited it after its completion as Wright's "most beautiful job";^[5] it is listed among *Smithsonian's* Life List of 28 places "to visit before you die."^[6] It was designated a [National Historic Landmark](#) in 1966.^[3] In 1991, members of the [American Institute of Architects](#) named the house the "best all-time work of American architecture" and in 2007, it was ranked twenty-ninth on the [list of America's Favorite Architecture](#) according to the [AIA](#).



Recall: Your First 447 Assignment

- Go and visit Fallingwater
- Appreciate the importance of out-of-the-box and creative thinking
- Think about tradeoffs in the design of the building
 - Strengths, weaknesses
- Derive principles on your own for good design and innovation
- Due date: **After passing this course**
 - Apply what you have learned in this course
 - Think out-of-the-box

Recall: But First, Today's First Assignment

- Find The Differences Of This and That

Find Differences Of This and That



Many Tradeoffs Between Two Designs

- You can list them after you complete the first assignment...

Recall: A Key Question

- How Was Wright Able To Design Fallingwater?
- Can have many guesses
 - (Ultra) hard work, perseverance, dedication (over decades)
 - Experience of decades
 - Creativity
 - Out-of-the-box thinking
 - Principled design
 - A good understanding of past designs
 - Good judgment and intuition
 - Strong combination of skills (math, architecture, art, ...)
 - ...
- (You will be exposed to and hopefully develop/enhance many of these skills in this course)

A Quote from The Architect Himself

- “architecture [...] based upon **principle**, and not upon **precedent**”



A Principled Design

Organic architecture

From Wikipedia, the free encyclopedia

Organic architecture is a [philosophy](#) of [architecture](#) which promotes harmony between human habitation and the natural world through design approaches so sympathetic and well integrated with its site, that buildings, furnishings, and surroundings become part of a unified, interrelated composition.

A well-known example of organic architecture is [Fallingwater](#), the residence Frank Lloyd Wright designed for the Kaufmann family in rural Pennsylvania. Wright had many choices to locate a home on this large site, but chose to place the home directly over the waterfall and creek creating a close, yet noisy dialog with the rushing water and the steep site. The horizontal striations of stone masonry with daring [cantilevers](#) of colored beige concrete blend with native rock outcroppings and the wooded environment.



A Quote from The Architect Himself

- “architecture [...] based upon **principle**, and not upon **precedent**”




Recall: A Key Question

- How Was Wright Able To Design Fallingwater?
- Can have many guesses
 - (Ultra) hard work, perseverance, dedication (over decades)
 - Experience of decades
 - Creativity
 - Out-of-the-box thinking
 - Principled design
 - A good understanding of past designs
 - Good judgment and intuition
 - Strong combination of skills (math, architecture, art, ...)
 - ...
- (You will be exposed to and hopefully develop/enhance many of these skills in this course)

Recall: Role of the (Computer) Architect

Role of the Architect

- Look Backward (Examine old code)***
- Look forward (Listen to the dreamers)***
- Look Up (Nature of the problems)***
- Look Down (Predict the future of technology)***



from Yale Patt's lecture notes

Recall: Role of The (Computer) Architect

- Look backward (to the past)
 - Understand tradeoffs and designs, upsides/downsides, past workloads. Analyze and evaluate the past.
- Look forward (to the future)
 - Be the dreamer and create new designs. Listen to dreamers.
 - Push the state of the art. Evaluate new design choices.
- Look up (towards problems in the computing stack)
 - Understand important problems and their nature.
 - Develop architectures and ideas to solve important problems.
- Look down (towards device/circuit technology)
 - Understand the capabilities of the underlying technology.
 - Predict and adapt to the future of technology (you are designing for N years ahead). Enable the future technology.

Takeaways

- Being an architect is not easy
- You need to consider **many** things in designing a new system + have good intuition/insight into ideas/tradeoffs

- But, it is fun and can be very technically rewarding
- And, enables a great future
 - E.g., many scientific and everyday-life innovations would not have been possible without architectural innovation that enabled very high performance systems
 - E.g., your mobile phones

- This course will teach you how to become a good computer architect + advance the state of the art

Agenda

- Course Overview
- Assumed Background (Videos and Lectures to Study)
- How to Jump Into Research
- How to Do the Paper Reviews
- Some Promising Research Topics in Computer Architecture
- First Assignments (for Next Week)

18-740 Instructor: Onur Mutlu

- Associate Professor @ Carnegie Mellon University ECE/CS
- PhD from UT-Austin 2006, BS from Michigan 2000
- Past experience @ Microsoft Research, Intel, AMD
- omutlu@gmail.com (Best way to reach me)
- <http://www.ece.cmu.edu/~omutlu>
- <http://users.ece.cmu.edu/~omutlu/projects.htm>
- **Research and Education** in
 - Computer architecture and systems, bioinformatics
 - Memory and storage systems, emerging technologies
 - Many-core systems, heterogeneous systems, core design
 - Interconnects
 - Hardware/software interaction and co-design (PL, OS, Architecture)
 - Predictable and QoS-aware systems
 - Hardware fault tolerance and security
 - Algorithms and architectures for genome analysis
 - ...



18-740 Teaching Assistant: Nandita

- Nandita Vijaykumar

- PhD student with Onur Mutlu
- BE from PES Inst. Of Technology 2011
- Past experience @ AMD
- Research in Comp Arch (GPGPUs, ...)
- nandita@cmu.edu



- Office hours and locations will be posted online

- <http://www.ece.cmu.edu/~ece740/f15>

- Reach both of us at

- 740-official@ece.cmu.edu

Where to Get Up-to-date Course Info?

- Website: <http://www.ece.cmu.edu/~ece740>
- Piazza
- Your email
- Email to us: 740-official@ece.cmu.edu

Lecture and Recitation Locations, Times

■ Lectures:

- ❑ MW 7:30-9:30pm (Pitt), 4:30-6:20pm (SV)
- ❑ Hamerschlag Hall 1107 (Pitt), B23 118 (SV)
- ❑ Attendance is for your benefit and is therefore important
- ❑ Some days, we may have guest lectures

■ Recitations:

- ❑ T 7:30-9:20pm (Pitt), 4:30-6:20pm (SV)
- ❑ Doherty Hall A302 (Pitt), B19 1065 (SV)
- ❑ Very important to attend
- ❑ Most critical research related discussions will take place during recitation sessions

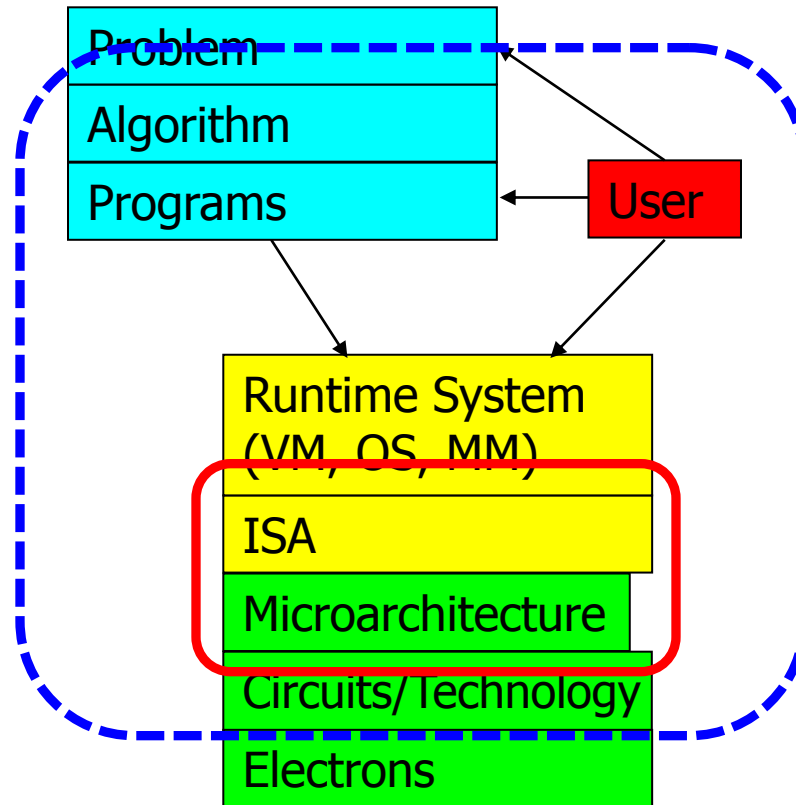
What Will You Learn

- **Computer Architecture:** The science and art of designing, selecting, and interconnecting hardware components and designing the hardware/software interface to create a computing system that meets functional, performance, energy consumption, cost, and other specific goals.
- **Traditional definition:** “The term *architecture* is used here to describe the attributes of a programmer, i.e., the conceptual structure and behavior as distinct from the organization and controls, the logic design, and the implementation.” *Gene Amdahl, IBM Systems Journal*, 1964



Dr. Amdahl holding a 100gate LSI air-cooled chip. On his desk is a circuit board with the chips on it. This circuit board was for an Amdahl 470 V/6 (photograph dated March 1973).

Levels of Transformation



- Read: Patt, "Requirements, Bottlenecks, and Good Fortune: Agents for Microprocessor Evolution," Proceedings of the IEEE 2001.

Scope of the Course

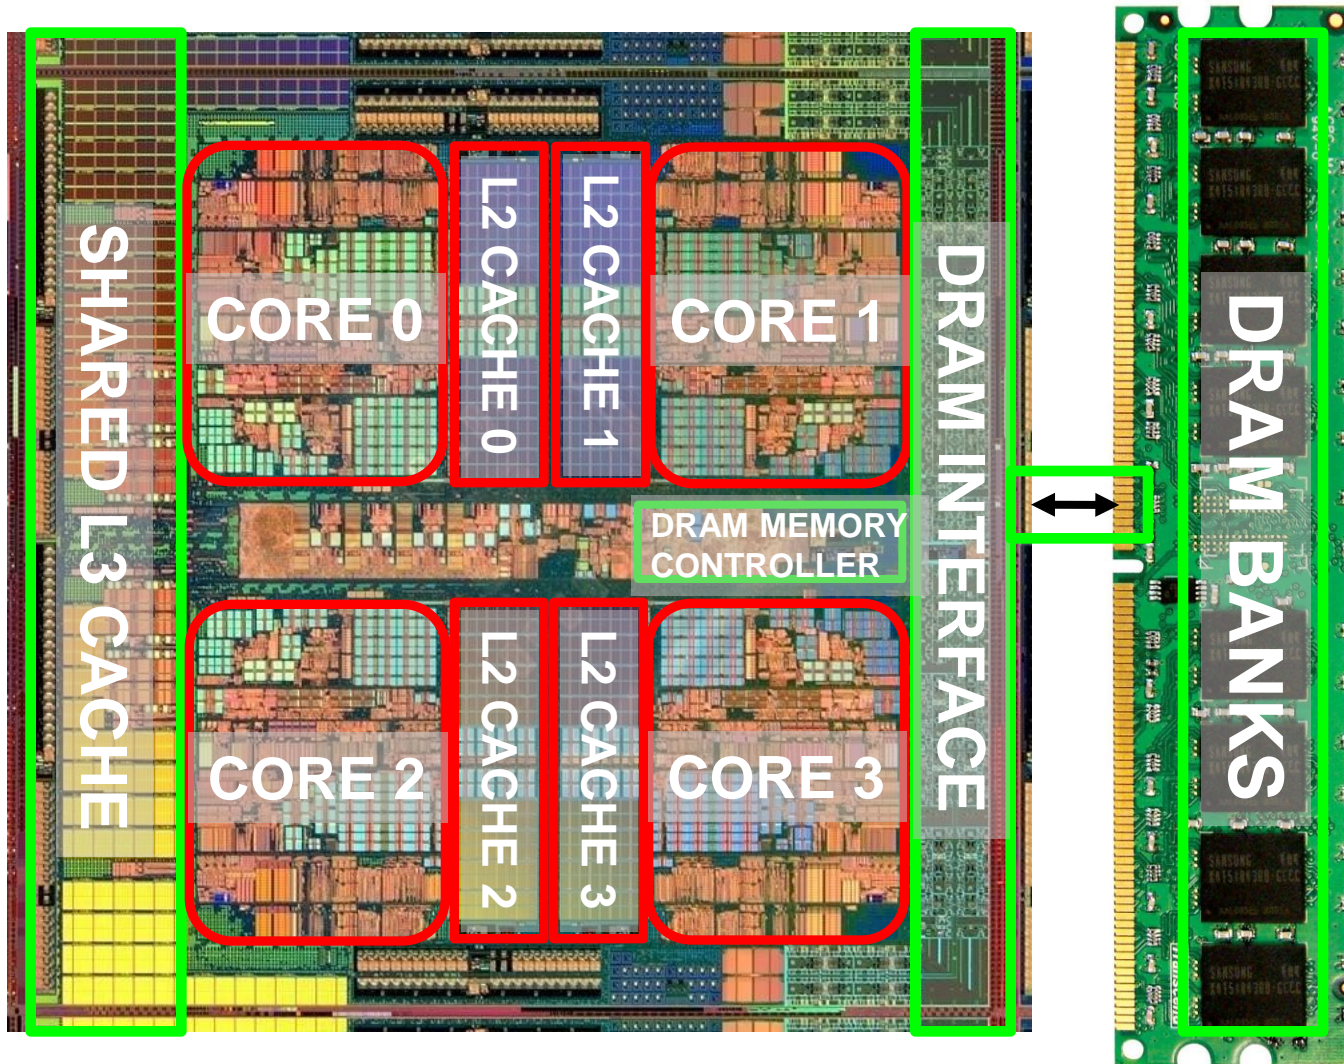
- **This course will take a broad view of architecture**
 - Beyond the ISA+microarchitecture levels
 - E.g., system-architecture interfaces and interactions
 - E.g., application-architecture interfaces and interactions
- **Out-of-the-box thinking is greatly encouraged**
 - E.g., research projects and readings on architectures that challenge the current dominant paradigms
 - processing in memory, approximate systems, asymmetry everywhere, ...
 - E.g., readings on topics that are traditionally covered less in computer architecture courses

What Will You Learn?

- Hardware/software interface, major components, and programming models of a modern computing platform
 - State-of-the-art as well as research proposals (lots of them)
 - Tradeoffs and how to make them
 - **Emphasis on cutting-edge (research & state-of-the-art)**
- Hands-on research in a computer architecture topic
 - **Semester-long research project**
 - Focus: How to design better architectures (not an intro course)
- How to dig out information
 - No textbook really required
 - But, see the syllabus


An Example: Multi-Core Systems

Multi-Core Chip

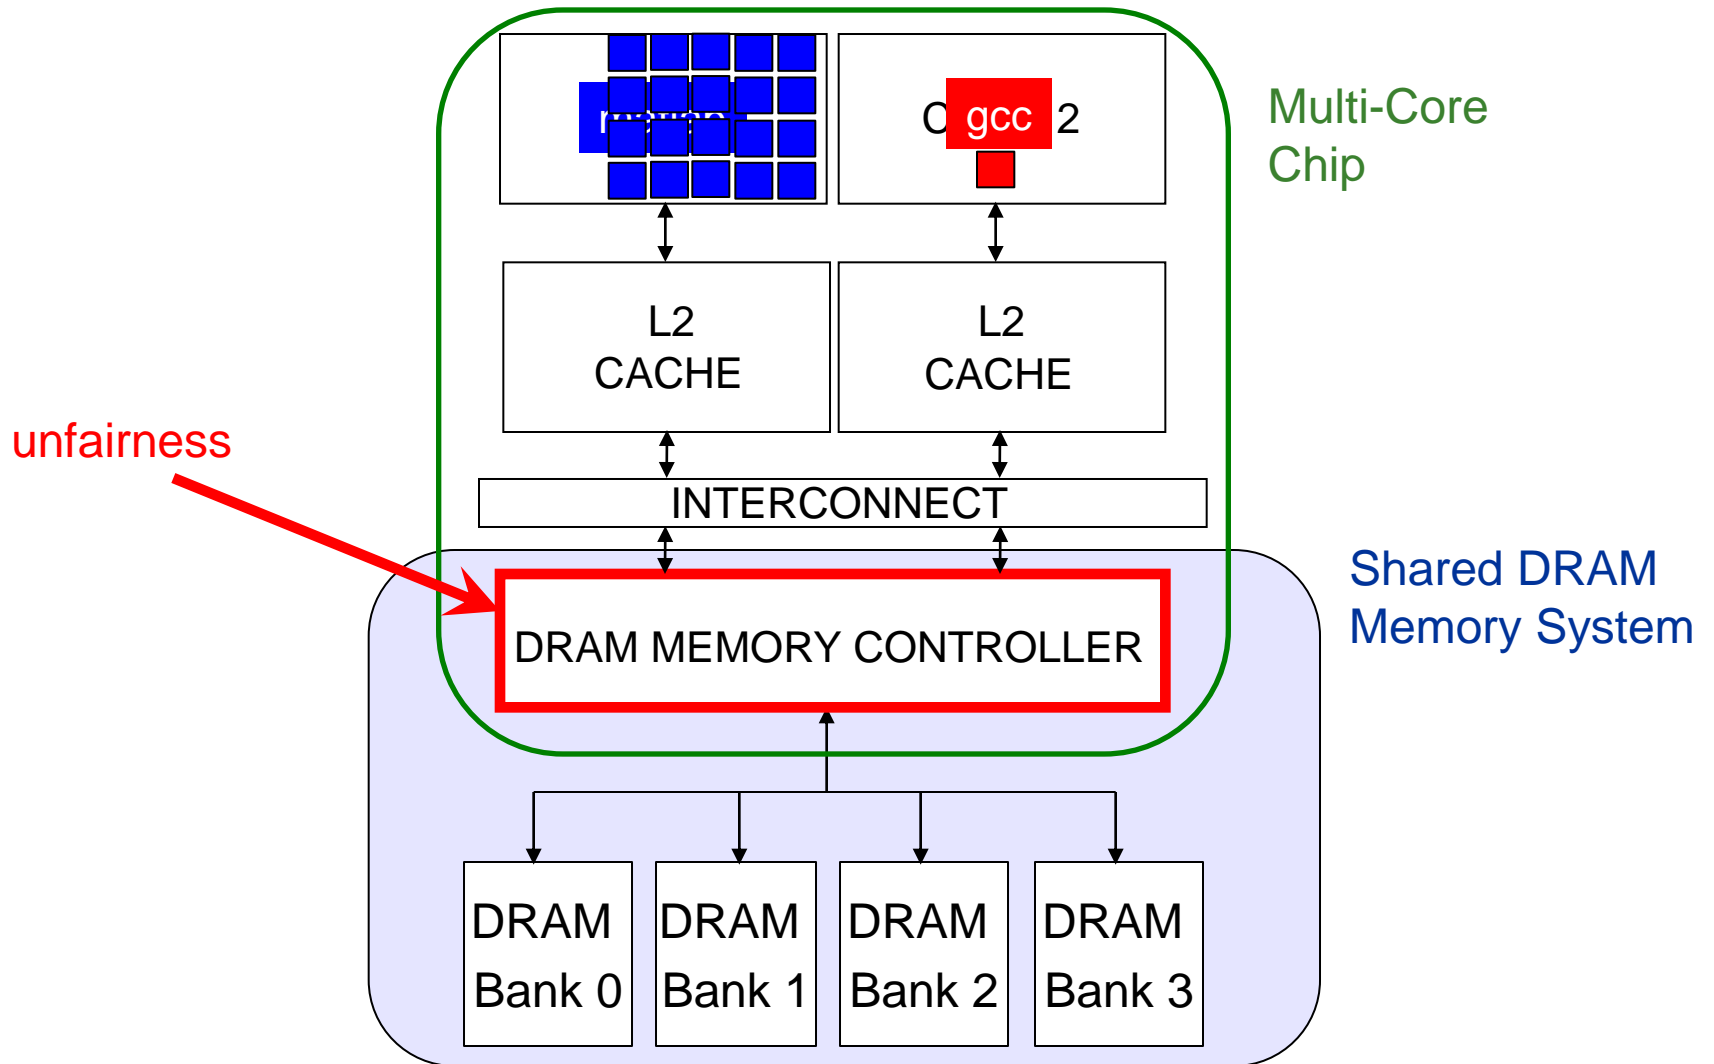


*Die photo credit: AMD Barcelona

Unexpected Slowdowns in Multi-Core



Why the Disparity in Slowdowns?



For More Information, Read

- Thomas Moscibroda and Onur Mutlu,
"Memory Performance Attacks: Denial of Memory Service in Multi-Core Systems"
*Proceedings of the 16th USENIX Security Symposium (**USENIX SECURITY**),*
pages 257-274, Boston, MA, August 2007. [Slides \(ppt\)](#)

Course Goals

- Goal 1: To familiarize you with both **fundamental design tradeoffs and recent research issues/trends** in **processor, memory, and platform architectures** in today's and future systems.
 - Strong emphasis on fundamentals and design tradeoffs.
- Goal 2: To provide the necessary **background and experience** to advance the state-of-the-art in computer architecture by **performing cutting-edge research**.
 - Strong emphasis on
 - **Critical analysis of research papers** (through reading and literature review assignments)
 - **Developing new mechanisms that advance the state of the art** (through the course research project).

This is a Graduate-Level Class

- Required background:
 - basic architecture (18-447)
 - basic compilers
 - basic OS
 - programming skills
 - **spirit, excitement, and dedication** for deep exploration of a topic in computer architecture

What Do I Expect From You?

- Learn the material & dig deeper
- **Work hard**
- **Ask questions, take notes, participate in discussion**
- **Critically review the assigned research papers & readings**
 - Discuss/critique them online with peers and us
- Use Piazza and Review Site frequently...
- **Start the research project early and focus**
- Remember “**Chance favors the prepared mind.**” (Pasteur)



How Will You Be Evaluated?

- Research Project: 50%
- Critical Reviews, Presentations, Participation: 35%
- Exam(s): 15%
- My evaluation of your performance: 5%

- Participation+discussion is very important

- Grading will be back-end heavy. Most of your grade will be determined in December
 - How you prepare and manage your time is important
 - But grades should not be the reason for taking this course

Research Project

- Your chance to explore in depth a computer architecture topic that interests you
- Perhaps even publish your innovation in a top computer architecture conference.
- **Start thinking about your project topic from now!**
- Interact with me and Nandita
- Read the project topics handout well
- Groups of 2-3 students (will finalize this later)
- Proposal due: within ~3-4 weeks of first recitation session
 - Exact date will be announced

Policies

- Late policy: Maximum five late days total on any assignment
- **Absolutely no tolerance on cheating or academic dishonesty**
 - See syllabus, CMU Policy, and ECE Academic Integrity Policy
 - Cheating → Failing grade (no exceptions)

Syllabus and Course Website

- See the course website for up-to-date information:
 - <http://www.ece.cmu.edu/~ece740>

Agenda

- Course Overview
- Assumed Background (Videos and Lectures to Study)
- How to Jump Into Research
- How to Do the Paper Reviews
- Some Promising Research Topics in Computer Architecture
- First Assignments (for Next Week)

Recommended Background Videos and Lectures (I)

- All 447 lecture videos and notes are at:
- **Undergraduate Computer Architecture Course Lecture Videos (2013, 2014, 2015)**
- **Undergraduate Computer Architecture Course Materials (2013, 2014, 2015)**

- Please watch as many as you can, to brush up on background material → enables research project ideas

- Everyone should watch
 - Lecture 1: Basics of Computer Architecture
 - https://www.youtube.com/watch?v=zLP_X4wyHbY

Recommended Background Videos and Lectures (II)

- Lectures 2-3: Fundamental Concepts and ISA, ISA Tradeoffs
 - <http://www.youtube.com/watch?v=BqJgYN6S6Qw&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=2>
 - <http://www.youtube.com/watch?v=BqJgYN6S6Qw&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=3>
- Lecture 8: Pipelining
 - http://www.youtube.com/watch?v=5E_W7EeNs8U&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=8
- Lecture 9: Data Dependence Handling
 - <http://www.youtube.com/watch?v=Gpz1I47LfDo&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=9>
- Lecture 10-11: Branch Prediction
 - <http://www.youtube.com/watch?v=XkerLktFtJg&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=11>
- Lecture 16: Virtual Memory
 - <http://www.youtube.com/watch?v=ppPq-ntaAWU&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=16>

Recommended Background Videos and Lectures (II)

■ Lecture 22: Memory Hierarchy

- <http://www.youtube.com/watch?v=JBdfZ5i21cs&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=22>

■ Lecture 23-24: Caches

- <http://www.youtube.com/watch?v=TpMdBrM1hVc&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=23>
- <http://www.youtube.com/watch?v=TboaFbjTd-E&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=24>

■ Lecture 30B: Multiprocessors

- http://www.youtube.com/watch?v=7ozCK_Mgxfk&list=PL5PHm2jkkXmidJOd59REog9jDnPDTG6IJ&index=31

You Can Find Many Research Project Ideas

- By examining past 447 lectures
- By doing recommended readings from past 447s

Agenda

- Course Overview
- Assumed Background (Videos and Lectures to Study)
- How to Jump Into Research
- How to Do the Paper Reviews
- Some Promising Research Topics in Computer Architecture
- First Assignments (for Next Week)

How To Do Research & Advanced Dev.

- We will talk a lot about this in this course
- Learning **by example**
 - Reading and evaluating strong and seminal papers
- Learning **by doing**
 - Semester-long research project
- Learning **by open, critical discussions**
 - Recitation sessions, online discussion of papers & ideas on Piazza and the Paper Review Site

What Is The Goal of Research?

- To generate new insight
 - that can enable what previously did not exist

- Research (in engineering) is a hunt for insight that can eventually impact the world

Some Basic Advice for Good Research

- Choose great problems to solve: Have great taste
 - Difficult
 - Important
 - High impact
- Read heavily and critically
- Think big (out of the box)
 - Do not restrain yourself to tweaks
- Aim high
- Write and present really well



Looking here for lost keys

Lost keys here




Looking here







Current Architecture Practice

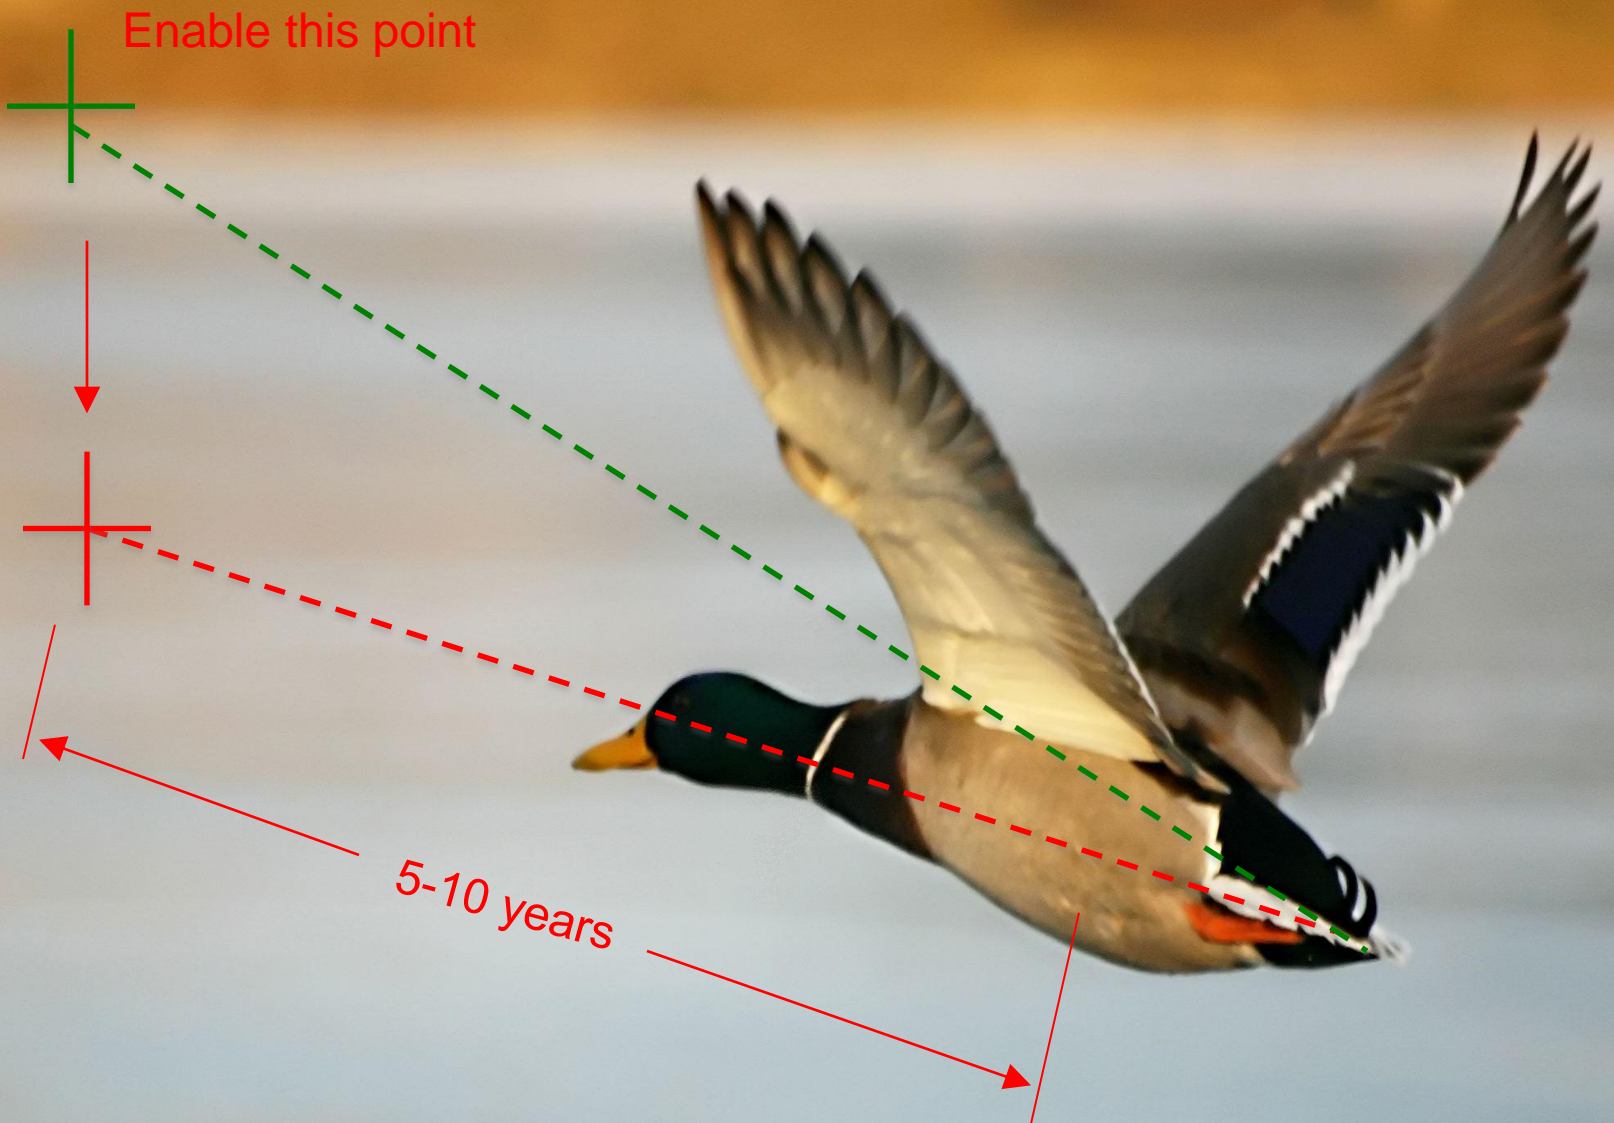


Aim Here



5-10 years





Enable this point

5-10 years

The Research Formula



$$ROI = \frac{\textit{reward}}{\textit{risk} \times \textit{effort}}$$

Reward

If you are wildly successful, what difference will it make?

$$ROI = \frac{\textit{reward}}{\textit{risk} \times \textit{effort}}$$

Effort

Learn as much as possible with as little work as possible

$$ROI = \frac{\textit{reward}}{\textit{risk} \times \textit{effort}}$$

Effort

Do the minimum analysis and experimentation necessary to make a point

$$ROI = \frac{\textit{reward}}{\textit{risk} \times \textit{effort}}$$

Research is a
hunt for insight

Need to get off the beaten
path to find new insights



Recommended Talk

- Bill Dally, Moving the needle: Effective Computer Architecture Research in Academy and Industry
ISCA 2010 Keynote Talk.
- Acknowledgment: Past few slides are from this talk


What transfers is *insight*

Not academic design

Not performance numbers



More Good Advice



“The purpose of computing is
insight, not numbers”

Richard Hamming

Agenda

- Course Overview
- Assumed Background (Videos and Lectures to Study)
- How to Jump Into Research
- How to Do the Paper Reviews
- Some Promising Research Topics in Computer Architecture
- First Assignments (for Next Week)

How to Do the Paper Reviews

- 1: Brief summary
 - What is the problem the paper is trying to solve?
 - What are the key ideas of the paper? Key insights?
 - What is the key contribution to literature at the time it was written?
 - What are the most important things you take out from it?
- 2: Strengths (most important ones)
 - Does the paper solve the problem well?
- 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: Can you do (much) better? Present your thoughts/ideas.
- 5: What have you learned/enjoyed/disliked in the paper? Why?

- Review should be short and concise (~a page)

Advice on Paper Reviews

- When doing the reviews, be very critical
- Always think about better ways of solving the problem or related problems
- Do background reading
 - Reviewing a paper/talk is the best way of learning about a research problem/topic
- Think about forming a literature survey topic or a research proposal based on the paper (for future studies)

Video on How to Do Paper Reviews

- Please watch this short video before you do the reviews:
 - <http://www.youtube.com/watch?v=tOL6FANAj8c>
 - <http://www.ece.cmu.edu/~ece740/f13/lib/exe/fetch.php?media=onur-740-fall13-lecture0-3-how-to-do-the-paper-reviews.pdf>

- Example reviews are provided here:
 - My 447 class website from Spring 2015
 - <http://www.ece.cmu.edu/~ece447/s15/lib/exe/fetch.php?media=review-sms.pdf>
 - <http://www.ece.cmu.edu/~ece447/s15/lib/exe/fetch.php?media=review-chapter.pdf>
 - <http://www.ece.cmu.edu/~ece447/s15/lib/exe/fetch.php?media=review-chapter-2.pdf>

Reading(s) on Refereeing CS Papers

- Smith, “The Task of the Referee,” IEEE Computer 1990.
 - Provides an idea of the publication process
 - Provides guidance on how to perform technical reviews

- Also see:
 - Hill and McKinley, “Notes on Constructive and Positive Reviewing”
 - <http://www.cs.utexas.edu/users/mckinley/notes/reviewing.html>
 - Levin and Redell, “How (and how not) to write a good systems paper,” OSR 1983.
 - Jones, “How to Write a Great Research Paper”

Literature Survey

- More information to come... In the meantime:
- Read a lot of papers; find focused problem areas to survey papers on
- We will provide a list of project ideas and papers associated with them
- A good way of finding topics to survey or do projects on is:
 - Examining the provided project ideas and papers
 - Reading assigned papers in lectures
 - Examining papers from recent conferences (ISCA, MICRO, HPCA, ASPLOS, ...)

Agenda

- Course Overview
- Assumed Background (Videos and Lectures to Study)
- How to Jump Into Research
- How to Do the Paper Reviews
- Some Promising Research Topics in Computer Architecture
- First Assignments (for Next Week)

Required Reviews

- Due Tuesday Sep 8 @ 3pm
- Enter your reviews on the review website
- Start discussing ideas and thoughts on Piazza

Review Paper 1 (Required)

- Onur Mutlu and Lavanya Subramanian,
"Research Problems and Opportunities in Memory Systems"
Invited Article in Supercomputing Frontiers and Innovations (SUPERFRI), 2015.
- Talk Video and Slides
 - Onur Mutlu,
"Rethinking the Systems We Design"
*Technical talk at A Computer Architecture Workshop: Visions for the Future (Celebrating Yale@75) (YALE@75), Austin, TX, September 2014. Slides (pptx) (pdf)
Video*
 - https://www.youtube.com/watch?v=QPPaHi_rp5E

Review Talk 1 (Required)

- Onur Mutlu,
"Rethinking the Systems We Design"
Technical talk at A Computer Architecture Workshop: Visions for the Future (Celebrating Yale@75) (YALE@75), Austin, TX, September 2014. Slides (pptx) (pdf)
- Video
https://www.youtube.com/watch?v=QPPaHi_rp5E

Review Paper 2 (Optional For Now)

- Junwhan Ahn, Sungpack Hong, Sungjoo Yoo, Onur Mutlu, and Kiyong Choi,
"A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing"
Proceedings of the 42nd International Symposium on Computer Architecture (ISCA), Portland, OR, June 2015.
[\[Slides \(pdf\)\]](#) [\[Lightning Session Slides \(pdf\)\]](#)

Review Paper 3 (Required)

- Vivek Seshadri, Kevin Hsieh, Amirali Boroumand, Donghyuk Lee, Michael A. Kozuch, Onur Mutlu, Phillip B. Gibbons, and Todd C. Mowry,
"Fast Bulk Bitwise AND and OR in DRAM"
IEEE Computer Architecture Letters (**CAL**), April 2015.

Agenda

- Course Overview
- Assumed Background (Videos and Lectures to Study)
- How to Jump Into Research
- How to Do the Paper Reviews
- Some Promising Research Topics in Computer Architecture
- First Assignments (for Next Week)

18-740: Computer Architecture

Recitation 1: Introduction, Logistics, and Jumping Into Research

Prof. Onur Mutlu

Carnegie Mellon University

Fall 2015

September 1, 2015