

CSE 548: Computer Systems Architecture

Class Overview

Spring 2017

Luis Ceze (Instructor)

Thierry Moreau (TA)

Luis Ceze




Joined UW-CSE 10 years ago...


Eating is more important than nutrition :)



Near zero-energy computer vision.



Accelerating ML training.



Realtime depth estimation/stitching
for 4K 360-3D video



Realtime, live!



OS

ML


PL

Architecture

Hardware

Biology

Global-scale DNA sequencing analysis.




sailipa

~Approximate computing

Efficiency and
performance

Output quality-of-
result

Silicon meets Biotech



TA: Thierry Moreau

4th year, advised by Luis

Can help you with ?s about:

- Computer architecture
- 548 Homework
- 548 Projects
 - (specifically how to come up with a good name)
- Systems, PL, compilers, ML...
- French lessons, how to choose french wine

Who are you? 😊

- 20+ students! Mix of seniorities & backgrounds/interests
 - HPC (2), Architecture (5), ML/NLP/Vision (5), PL (4), ICTD, Robotics, Networks(2)
 - 55% don't know what Amdahl's law is
 - *"Thou shalt not covet thy neighbor's processor architecture"* 😊
- Quick round of introductions
 - Research interests
 - Other classes/seminars you are taking this quarter
 - Anything else you want us to know about you
- This class will be *much* better if we have rich discussions. So don't be shy!

Why Learn Computer Architecture?

- All your programs **run in real machines**
 - you should understand well how they work
 - (ML needs machines 😊)
- Computer architecture is at a crossroads
 - **multi-cores, FGPAs, GPUs** everywhere
 - **power, dependability** are major concerns
 - **technology trends** motivate new applications

Class Outline

- Intro, Basic Principles, Technology and Metrics
- The Hardware/Software Interface (ISA)
- Superscalar/OoO execution
- Memory Hierarchy
- Parallelism (ILP, DLP, TLP)
 - Multiprocessors (Coherence/Consistency)
 - Multithreading/GPUs
- Specialized computing (FPGAs, ASICs)
 - Some emphasis on machine learning
- Warehouse-scale computing
 - Data-center visit confirmed for May 24th
- Silicon technology trends (e.g., 3D integration) and impact in IT
- Putting it all together with Real HW Systems dissections
 - Modern microprocessors, Xbox One, Raspberry Pi, Hololens
- Historical systems (and a special assignment)

Guest lectures/events

- Andrew Putnam – FPGA in the cloud
- Doug Carmean – Xeon Phi
- Brad Beckman – GPUs
- Gabe Loh – 3D chips
- Data-center visit
- Living Computer Museum visit
 - And historical machine assignment

Administrivia

- Readings:
 - Book (Hennessy and Patterson, 5th edition)
 - Readings for your own benefit, slides often enough
 - Papers: **Some are readings essay assignments**
 - write 1-2 (short 😊) **paragraphs** about the key points, submit to dropbox
- **Homework: 4 (light coding)**
- **Project: please form groups of (≤ 3) schedule meeting with me/Thierry**
 - We are aiming at 10-12 projects
 - Be encouraged to do a project that overlaps with your research and/or another (quals) class
- Grading:
 - 15% reading assignments
 - 35% homeworks
 - 50% project (report, poster)
- Check our website regularly
 - Shared calendar with classes, office hours, project meeting slots etc
 - <https://courses.cs.washington.edu/courses/cse548/17sp/>

Website up

Date	Topic & Readings	HW
March 29 (W)	No class.	
March 31 (F)	<div>Intro / Metrics</div> <ul style="list-style-type: none">• Lecture Slides• Chapter 1 of HP (that's <i>Hennessy & Patterson</i> not Harry Potter)• Moore's Law• Characterizing Computer Performance• Power: A First-Class Architectural Design Constraint• Producing Wrong Data Without Doing Anything Obviously Wrong!	<div>Commentary: Power: A First-Class Architectural De...</div> <div>Start thinking about a project</div> <div>Form a group</div> <div>Sign up for a meeting slot ⓘ</div>
April 5 (W)	<div>The ISA</div> <ul style="list-style-type: none">• Appendix A of HP• Transmeta• Instruction Sets and Beyond• ISA Power Struggles	

A *few* project ideas

- HW/SW codesign
 - find an app you care, profile, come up with HW extension ideas
 - e.g., ML, games, cryptocurrencies, wearables, vision, VR
 - verify aspects of HW+SW
- Using ML to predict program/system behavior
 - branch prediction, memory accesses, user behavior?
 - Using performance counters
 - security, performance optimization, malware detection, energy consumption estimation, etc.
- Security
 - HW support, side-channels, HW-based attacks, e.g. --- row-hammer
- Tools
 - Infer data type from core dump
 - Find useless updates to memory (dynamic redundancy detection)
 - User-focused design: Use mechanical turk to assess people's tolerance to slow-down? Output degradation?
 - Interesting visualization of system behavior (performance, energy, etc)
- Uses of new technologies
 - Eg., error correction of non-volatile memory, DNA, memristors, etc.
- In-network processing. Profile click modular router/OpenFlow?
- Profile IoT devices? Dissect architecture of an IoT device?

Project proposal “recipe”

- Clearly state what the problem/opportunity is
- State what your initial ideas are
- Short action plan (3-4 steps) and milestones
- Aim high! Computer architecture conferences are often fun and in interesting places 😊

Project tools

- Profilers
 - Gprof, Gperftools, Intel VTune, etc.
- Intel Pin (awesome!)
- LLVM compiler
- Simulators
 - Gem5, SESC, custom first-order, etc
- Performance counter libraries
- FPGA SoC