

**CS 61C**  
**Great Ideas in Computer Architecture**  
 (a.k.a. Machine Structures)  
**Lecture 1: Course Introduction**

Instructors:  
**Senior Lecturer SOE Dan Garcia** (call me "Dan")  
 (lots of help from TAs, esp **Head TA Alan Christopher**)  
<http://inst.eecs.berkeley.edu/~cs61c/>

1

**Agenda**

- Thinking about Machine Structures
- Great Ideas in Computer Architecture
- What you need to know about this class

2

**Agenda**

- Thinking about Machine Structures
- Great Ideas in Computer Architecture
- What you need to know about this class

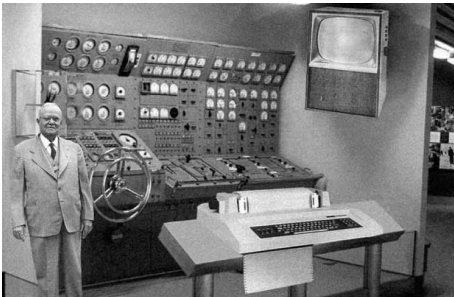
3

**CS61C is NOT really about C Programming**

- It is about the hardware-software interface
  - What does the programmer need to know to achieve the highest possible performance
- Languages like C are closer to the underlying hardware, unlike languages like Scheme!
  - Allows us to talk about key hardware features in higher level terms
  - Allows programmer to explicitly harness underlying hardware parallelism for high performance

4

**Old School CS61C**



*Scientists from the RAND Corporation have created this model to illustrate how a "beam computer" could look like in the near long. However the needed technology will not be economically feasible for the average home. Also the scientists readily admit that the computer will require not yet invented technology to actually work, but 30 years from now scientific progress is expected to solve these problems. With reshape interface and the Fortran language, the computer will be easy to use.*

5

Personal Mobile Devices

**New School CS61C (1/2)**

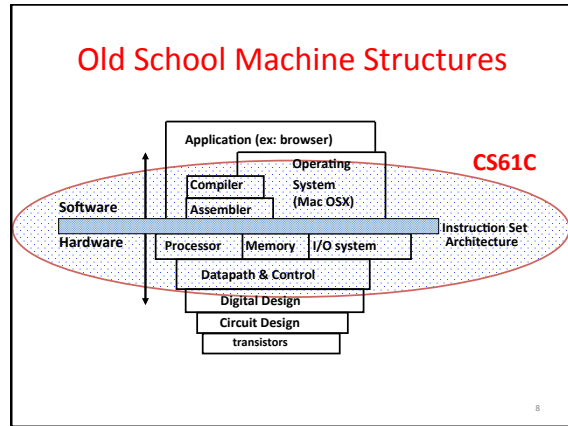


6

### Warehouse Scale Computer

## New School CS61C (2/2)

My other computer is a data center



### New-School Machine Structures (It's a bit more complicated!)

Project 2

**Software**

- Parallel Requests  
Assigned to computer  
e.g., Search "Katz"
- Parallel Threads  
Assigned to core  
e.g., Lookup, Ads
- Parallel Instructions  
>1 instruction @ one time  
e.g., 5 pipelined instructions
- Parallel Data  
>1 data item @ one time  
e.g., Add of 4 pairs of words
- Hardware descriptions  
All gates functioning in parallel at same time

**Hardware**

Warehouse Scale Computer

Smart Phone

Harness Parallelism & Achieve High Performance

### Agenda

- Thinking about Machine Structures
- Great Ideas in Computer Architecture
- What you need to know about this class

### 6 Great Ideas in Computer Architecture

- Abstraction (Layers of Representation/Interpretation)
- Moore's Law
- Principle of Locality/Memory Hierarchy
- Parallelism
- Performance Measurement & Improvement
- Dependability via Redundancy

### Great Idea #1: Abstraction (Levels of Representation/Interpretation)

High Level Language Program (e.g., C)

Compiler

Assembly Language Program (e.g., MIPS)

Assembler

Machine Language Program (MIPS)

Machine Interpretation

Hardware Architecture Description (e.g., block diagrams)

Architecture Implementation

Logic Circuit Description (Circuit Schematic Diagrams)


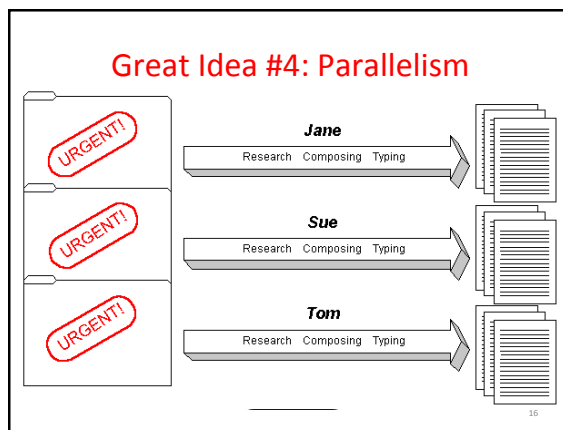
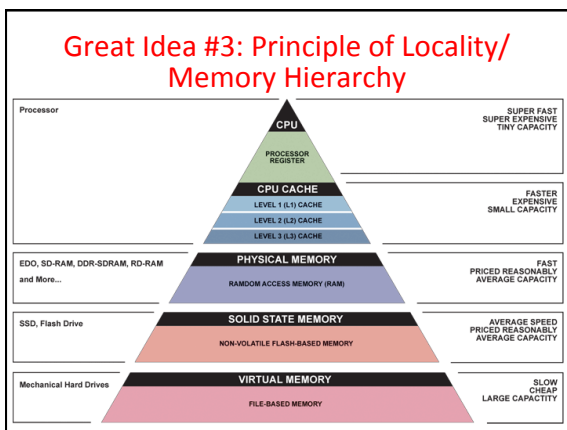
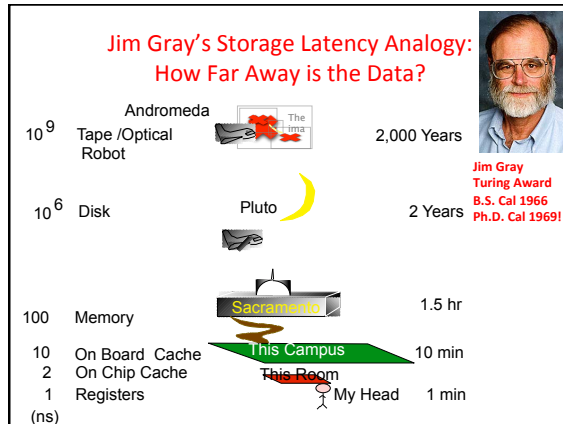

temp = v[k];  
v[k] = v[k+1];  
v[k+1] = temp;

lw \$t0, 0(\$2)  
lw \$t1, 4(\$2)  
sw \$t1, 0(\$2)  
sw \$t0, 4(\$2)

Anything can be represented as a number, i.e., data or instructions

```

0000 1001 1100 0110 1010 1111 0101 1000
1010 1111 0101 1000 0000 1001 1100 0110
1100 0110 1010 1111 0101 1100 1100 1100
0101 1000 0000 1001 1100
                                1
                                1
                                1
                                1
                    
```



- ### Great Idea #5: Performance Measurement and Improvement
- Matching application to underlying hardware to exploit:
    - Locality
    - Parallelism
    - Special hardware features, like specialized instructions (e.g., matrix manipulation)
  - Latency
    - How long to set the problem up
    - How much faster does it execute once it gets going
    - It is all about *time to finish*

### Coping with Failures

- 4 disks/server, 50,000 servers
- Failure rate of disks: 2% to 10% / year
  - Assume 4% annual failure rate
- On average, how often does a disk fail?
  - a) 1 / month
  - b) 1 / week
  - c) 1 / day
  - d) 1 / hour

19

### Coping with Failures

- 4 disks/server, 50,000 servers
- Failure rate of disks: 2% to 10% / year
  - Assume 4% annual failure rate
- On average, how often does a disk fail?
  - a) 1 / month
  - b) 1 / week
  - c) 1 / day
  - d) 1 / hour**

$50,000 \times 4 = 200,000$  disks  
 $200,000 \times 4\% = 8000$  disks fail  
 $365 \text{ days} \times 24 \text{ hours} = 8760$  hours

20

### Great Idea #6: Dependability via Redundancy

- Redundancy so that a failing piece doesn't make the whole system fail

Increasing transistor density reduces the cost of redundancy

21

### Great Idea #6: Dependability via Redundancy

- Applies to everything from datacenters to storage to memory to instructors
  - Redundant datacenters so that can lose 1 datacenter but Internet service stays online
  - Redundant disks so that can lose 1 disk but not lose data (Redundant Arrays of Independent Disks/RAID)
  - Redundant memory bits so that can lose 1 bit but no data (Error Correcting Code/ECC Memory)

22

### Agenda

- Thinking about Machine Structures
- Great Ideas in Computer Architecture
- What you need to know about this class

23

### Yoda says...


*“Always in motion, the future is...”*

Our schedule may change slightly depending on some factors.  
This includes lectures, assignments & labs...



### My goal as an instructor

- To make your experience in CS61C as enjoyable & informative as possible
  - Humor, enthusiasm, graphics & technology-in-the-news in lecture
  - Fun, challenging projects & HW
  - Pro-student policies (exam clobbering)
- To maintain Cal & EECS standards of excellence
  - Your projects & exams will be just as rigorous as every year.
- To be an HKN "7.0" man
  - I **know** I speak fast when I get excited about material. I'm told every semester. Help me slow down when I go toooo fast.
  - Please give me feedback so I improve! Why am I not 7.0 for you? I will listen!!



### Extra Credit: EPA!

- Effort**
  - Attending prof and TA office hours, completing all assignments, turning in HWO, doing reading quizzes
- Participation**
  - Attending lecture and voting using the clickers
  - Asking great questions in discussion and lecture and making it more interactive
- Altruism**
  - Helping others in lab or on Piazza

**EPA! extra credit points have the potential to bump students up to the next grade level! (but actual EPA! scores are internal)**

### Late Policy ... Slip Days!

- Assignments due at 11:59:59 PM
- You have 3 slip day tokens (NOT hour or min)
- Every day your project or homework is late (even by a minute) we deduct a token
- After you've used up all tokens, it's 33% deducted per day.
  - No credit if more than 3 days late
  - Save your tokens for projects, worth more!!
- No need for sob stories, just use a slip day!

### Policy on Assignments and Independent Work

- ALL PROJECTS WILL BE DONE WITH A PARTNER**
- With the exception of laboratories and assignments that explicitly permit you to work in groups, all homework and projects are to be YOUR work and your work ALONE.
- PARTNER TEAMS MAY NOT WORK WITH OTHER PARTNER TEAMS
- You are encouraged to discuss your assignments with other students, and extra credit will be assigned to students who help others, particularly by answering questions on Piazza, but we expect that what you hand in is yours.
  - It is NOT acceptable to copy solutions from other students.
  - It is NOT acceptable to copy (or start your) solutions from the Web.
  - It is NOT acceptable to use PUBLIC github archives (giving your answers away)**
- We have tools and methods, developed over many years, for detecting this. You WILL be caught, and the penalties WILL be severe.
- At the minimum F in the course**, and a letter to your university record documenting the incidence of cheating.
  - (We've caught people in recent semesters!)
- Both Giver and Receiver are equally culpable and suffer equal penalties**



SAN FRANCISCO — It's 1 p.m. on a Thursday and Dianne Bates, 40, juggles three screens. She listens to a few songs on her iPod, then taps out a quick e-mail on her iPhone and turns her attention to the high-definition television.

**Your Brain on Computers** Just another day at the gym.

At the University of California, San Francisco, scientists have found that when rats have a new experience, like exploring an unfamiliar area, their brains show new patterns of activity. But only when the rats take a break from their exploration do they process those patterns in a way that seems to create a persistent memory of the experience.

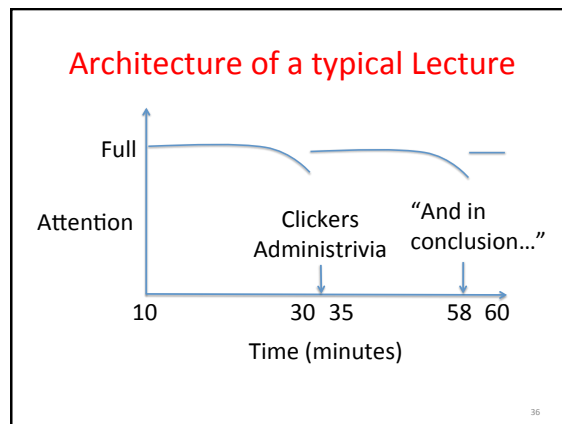
tasks, she is also in fast loops on an in a downtown is in good and elsewhere, and other to get work done — indite to boredom.

which in the last few years have become full-fledged with high-speed Internet connections, let people relieve the tedium of exercising, the grocery store line, stoplights in the dinner conversation.

The technology makes the tiniest windows of time entertaining, and potentially productive. But scientists point to an unanticipated side effect: when people keep their brains busy with digital input, they are forfeiting downtime that could allow them to better learn and remember information, or come up with new ideas.

Ms. Bates, for example, might be clearer-headed if she went for a run outside, away from her devices, research suggests.

### Architecture of a typical Lecture



The graph plots Attention (y-axis) against Time in minutes (x-axis). The attention level starts at a high level (Full) at 10 minutes and remains high until 30 minutes. At 30 minutes, there is a sharp drop in attention, labeled 'Clickers Administrivia'. The attention level then recovers slightly but remains lower than the initial level. At 58 minutes, there is another sharp drop, labeled '"And in conclusion..."', and the attention level falls to its lowest point at 60 minutes.

## Summary

- CS61C: Learn 6 great ideas in computer architecture to enable high performance programming via parallelism, not just learn C
  1. Abstraction  
(Layers of Representation/Interpretation)
  2. Moore's Law
  3. Principle of Locality/Memory Hierarchy
  4. Parallelism
  5. Performance Measurement and Improvement
  6. Dependability via Redundancy

37