

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/>

Agenda

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


Agenda

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

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

Old School CS61C



Scientists from the RAND Corporation have created this model to illustrate how a "home computer" could look like in the year 2004. 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 50 years from now scientific progress is expected to solve these problems. With teletype interface and the Fortran language, the computer will be easy to use.

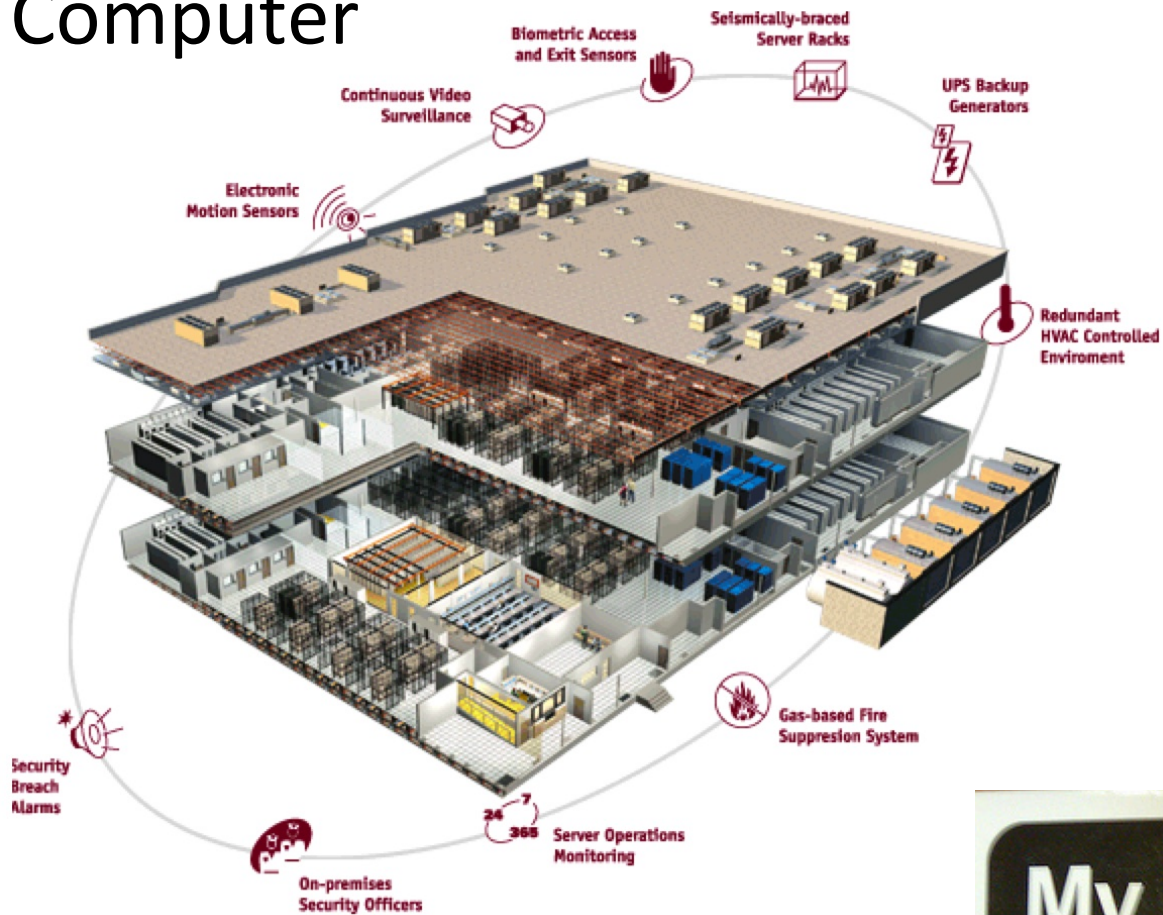
Personal
Mobile
Devices

New School CS61C (1/2)



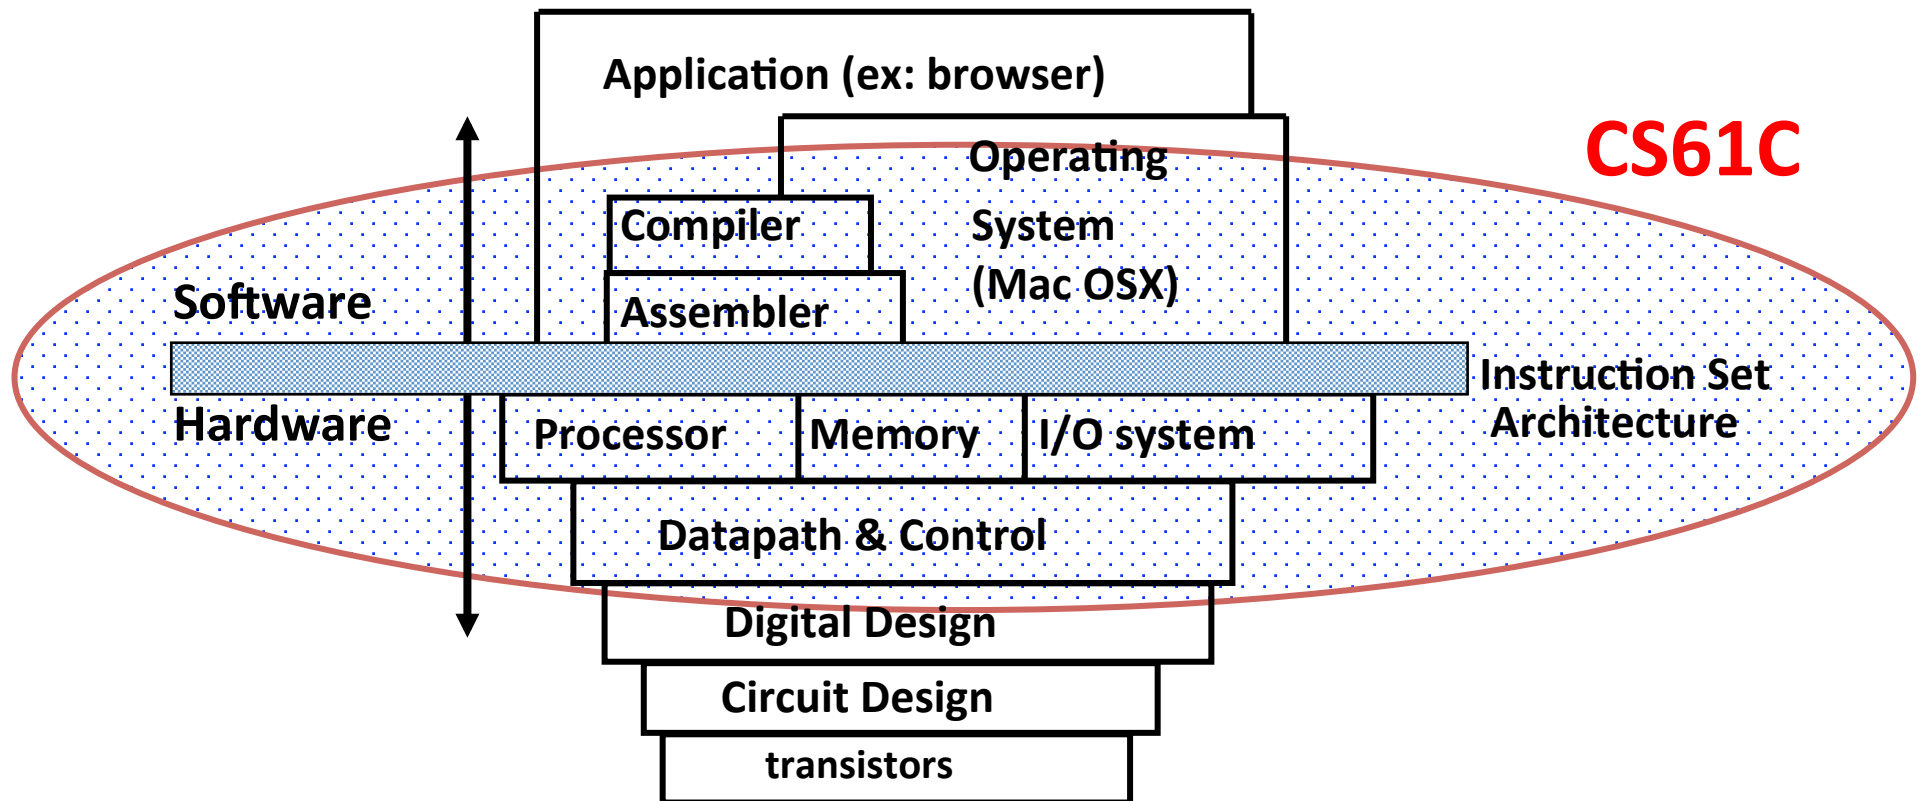
Warehouse Scale Computer

New School CS61C (2/2)



**My other computer
is a data center**

Old School Machine Structures



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

Project 2

Software

Hardware

- Parallel Requests
Assigned to computer
e.g., Search "Katz"

Warehouse
Scale
Computer



Smart
Phone



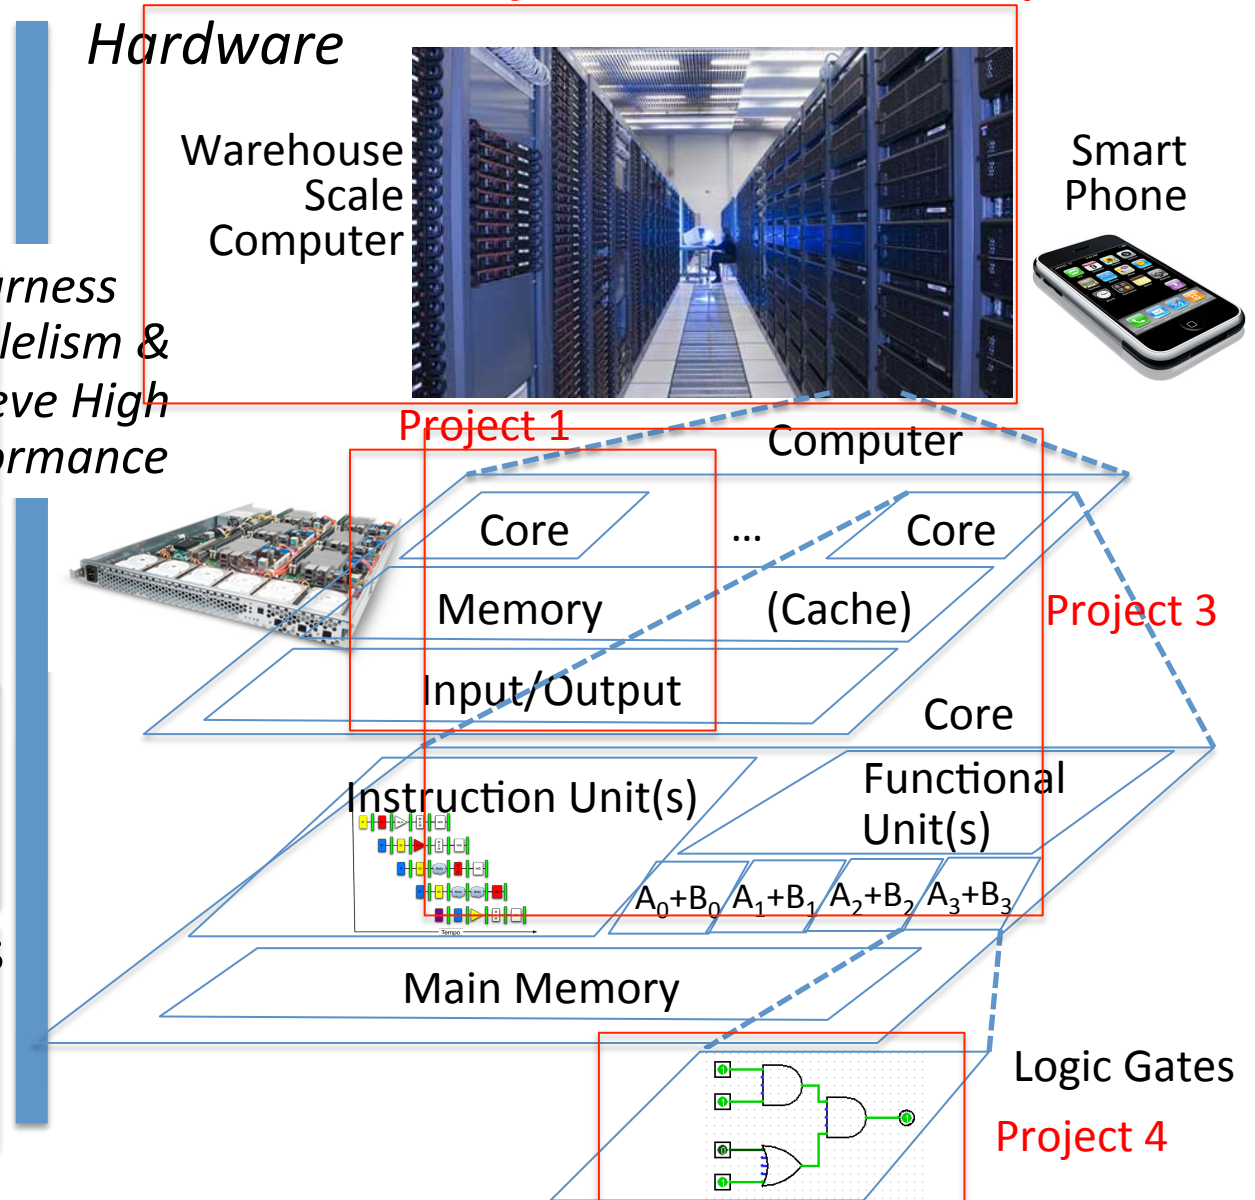
- Parallel Threads
Assigned to core
e.g., Lookup, Ads

*Harness
Parallelism &
Achieve High
Performance*

- 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



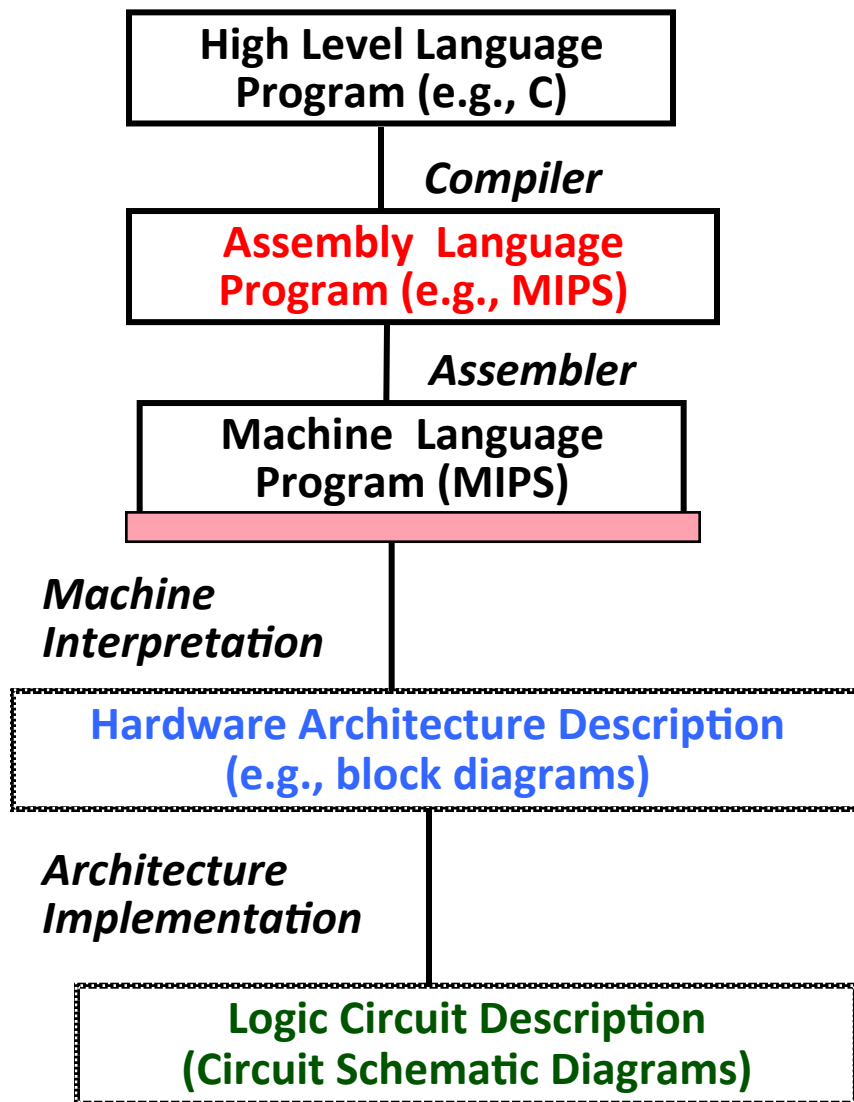
Agenda

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

6 Great Ideas in Computer Architecture

1. Abstraction
(Layers of Representation/Interpretation)
2. Moore's Law
3. Principle of Locality/Memory Hierarchy
4. Parallelism
5. Performance Measurement & Improvement
6. Dependability via Redundancy

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

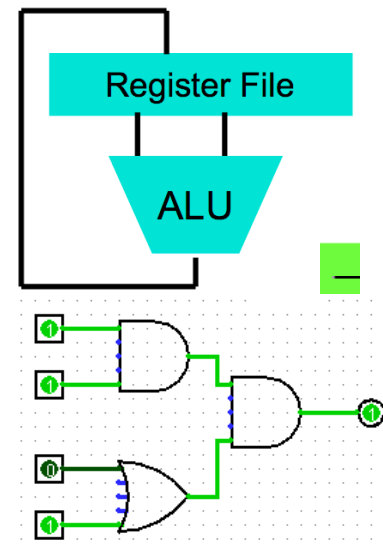


```
temp = v[k];
v[k] = v[k+1];
v[k+1] = temp;
```

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


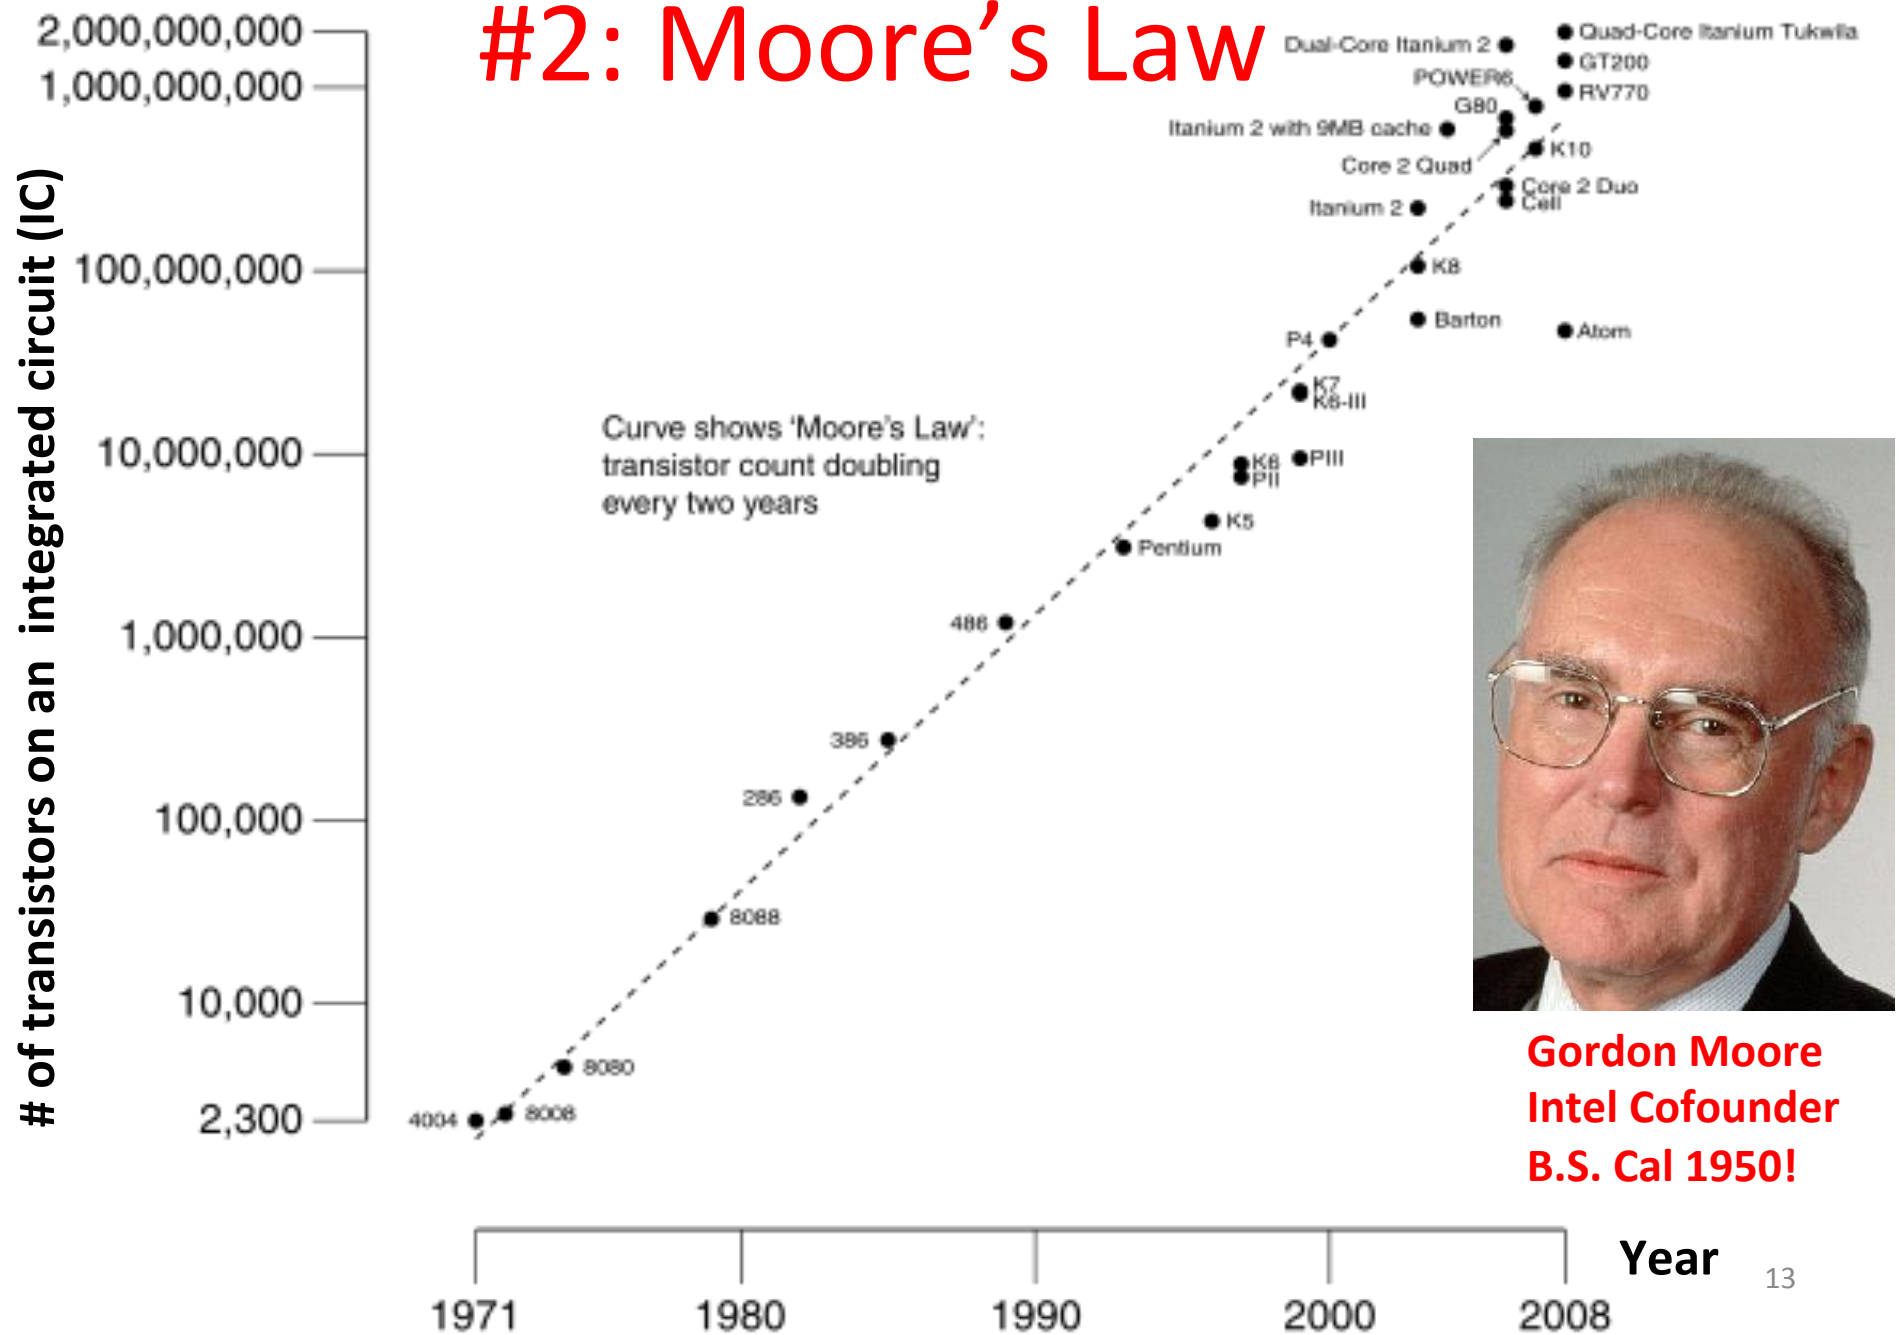
Anything can be represented
as a *number*,
i.e., data or instructions

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




Predicts: 2X Transistors / chip every 2 years

#2: Moore's Law

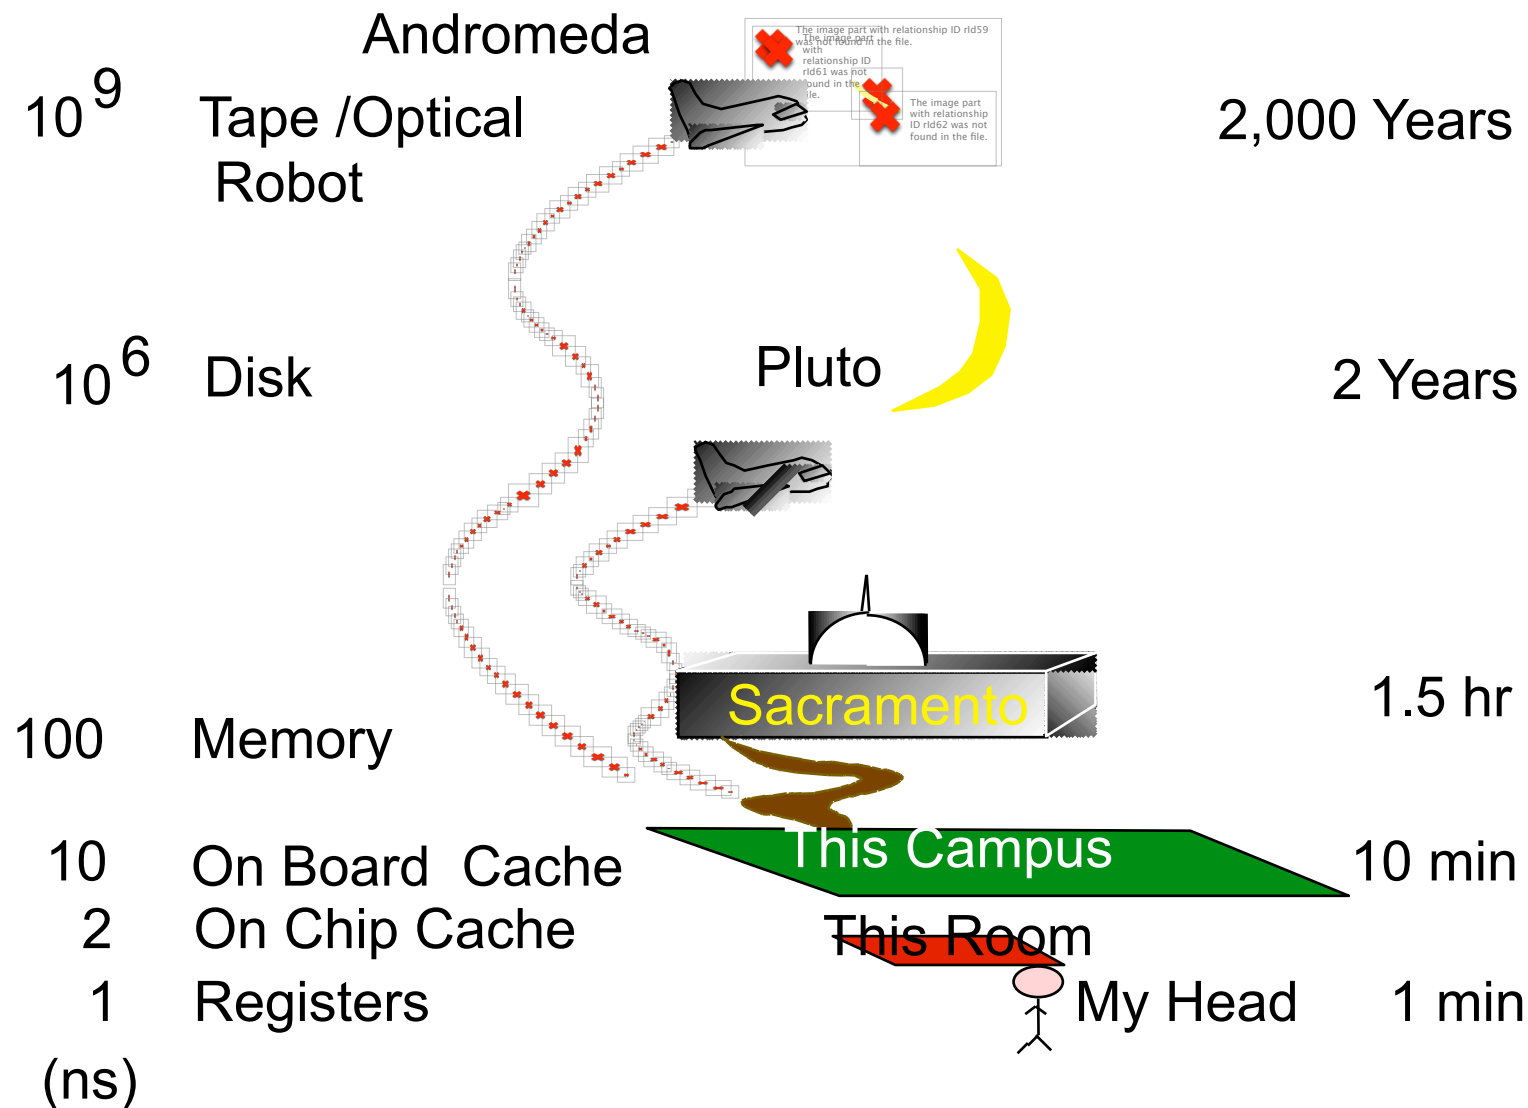


Gordon Moore
Intel Cofounder
B.S. Cal 1950!

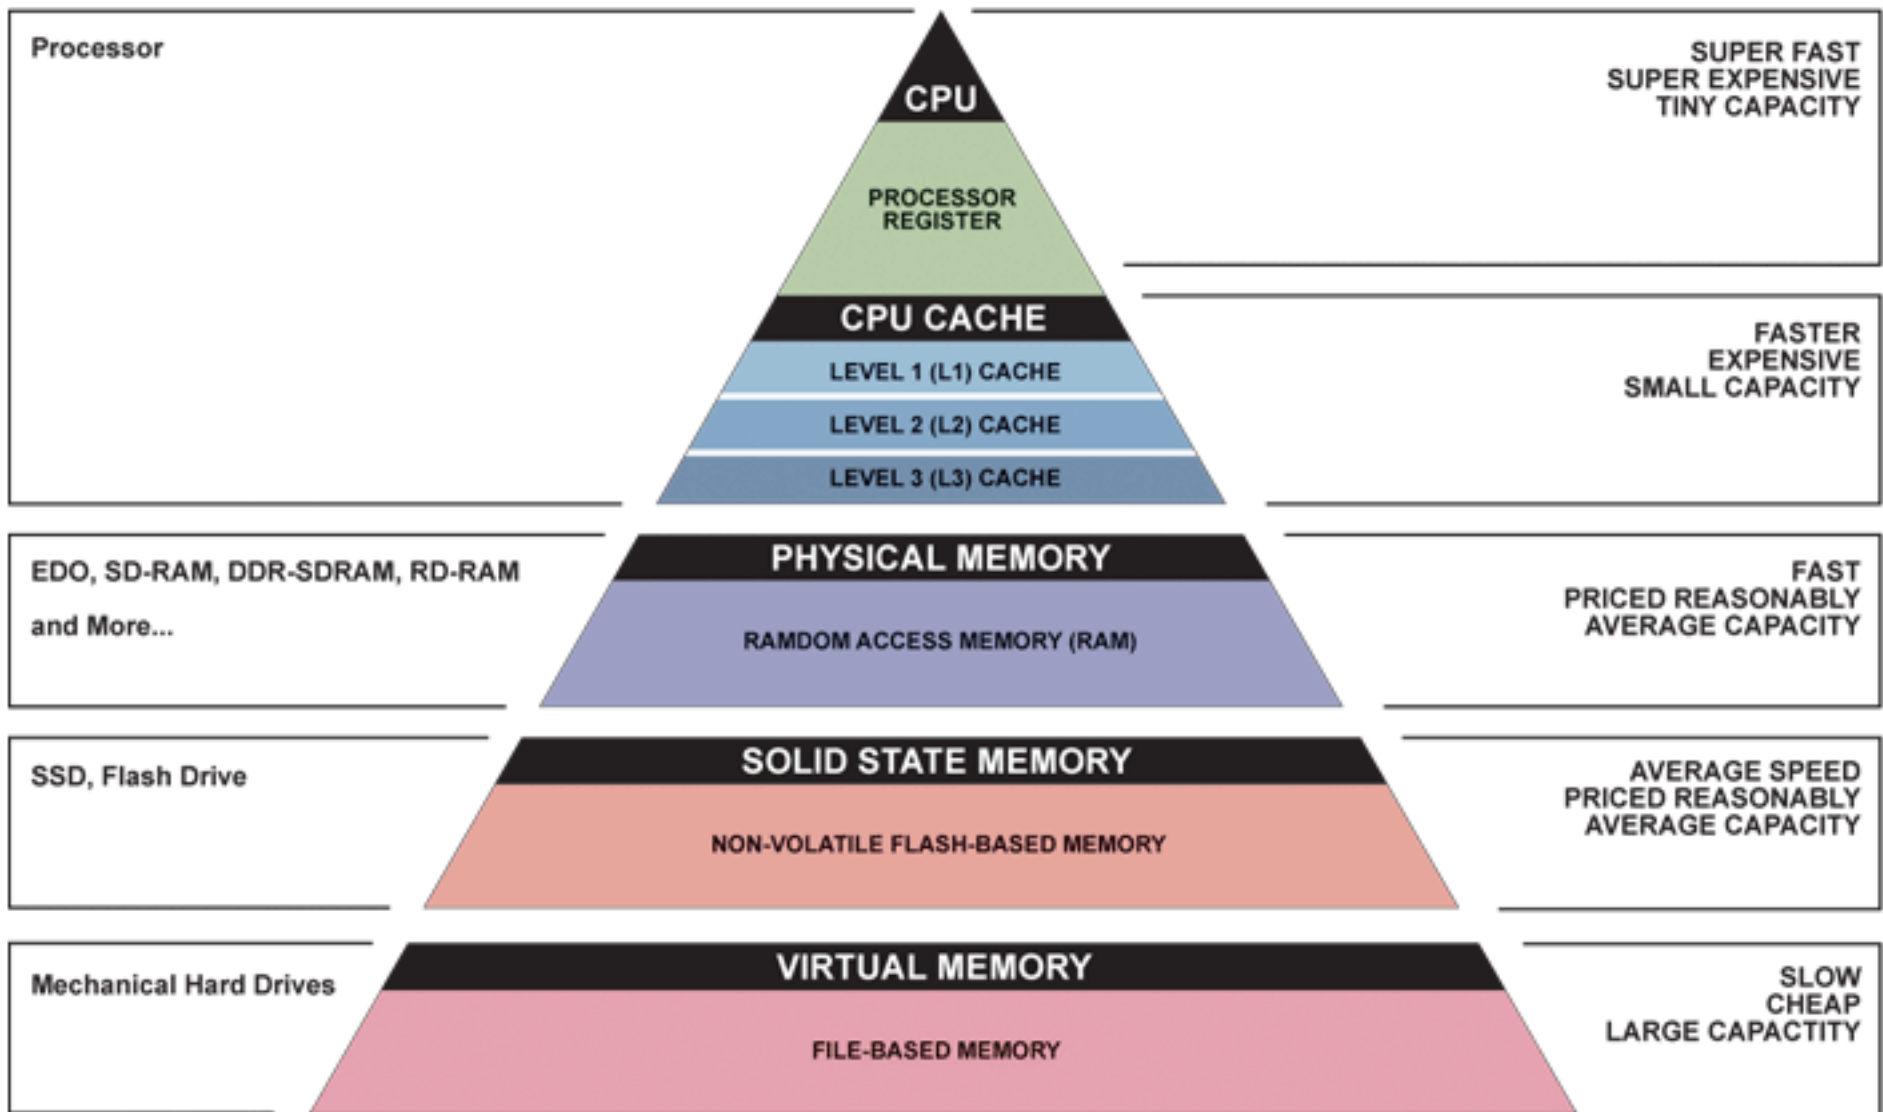
Jim Gray's Storage Latency Analogy: How Far Away is the Data?



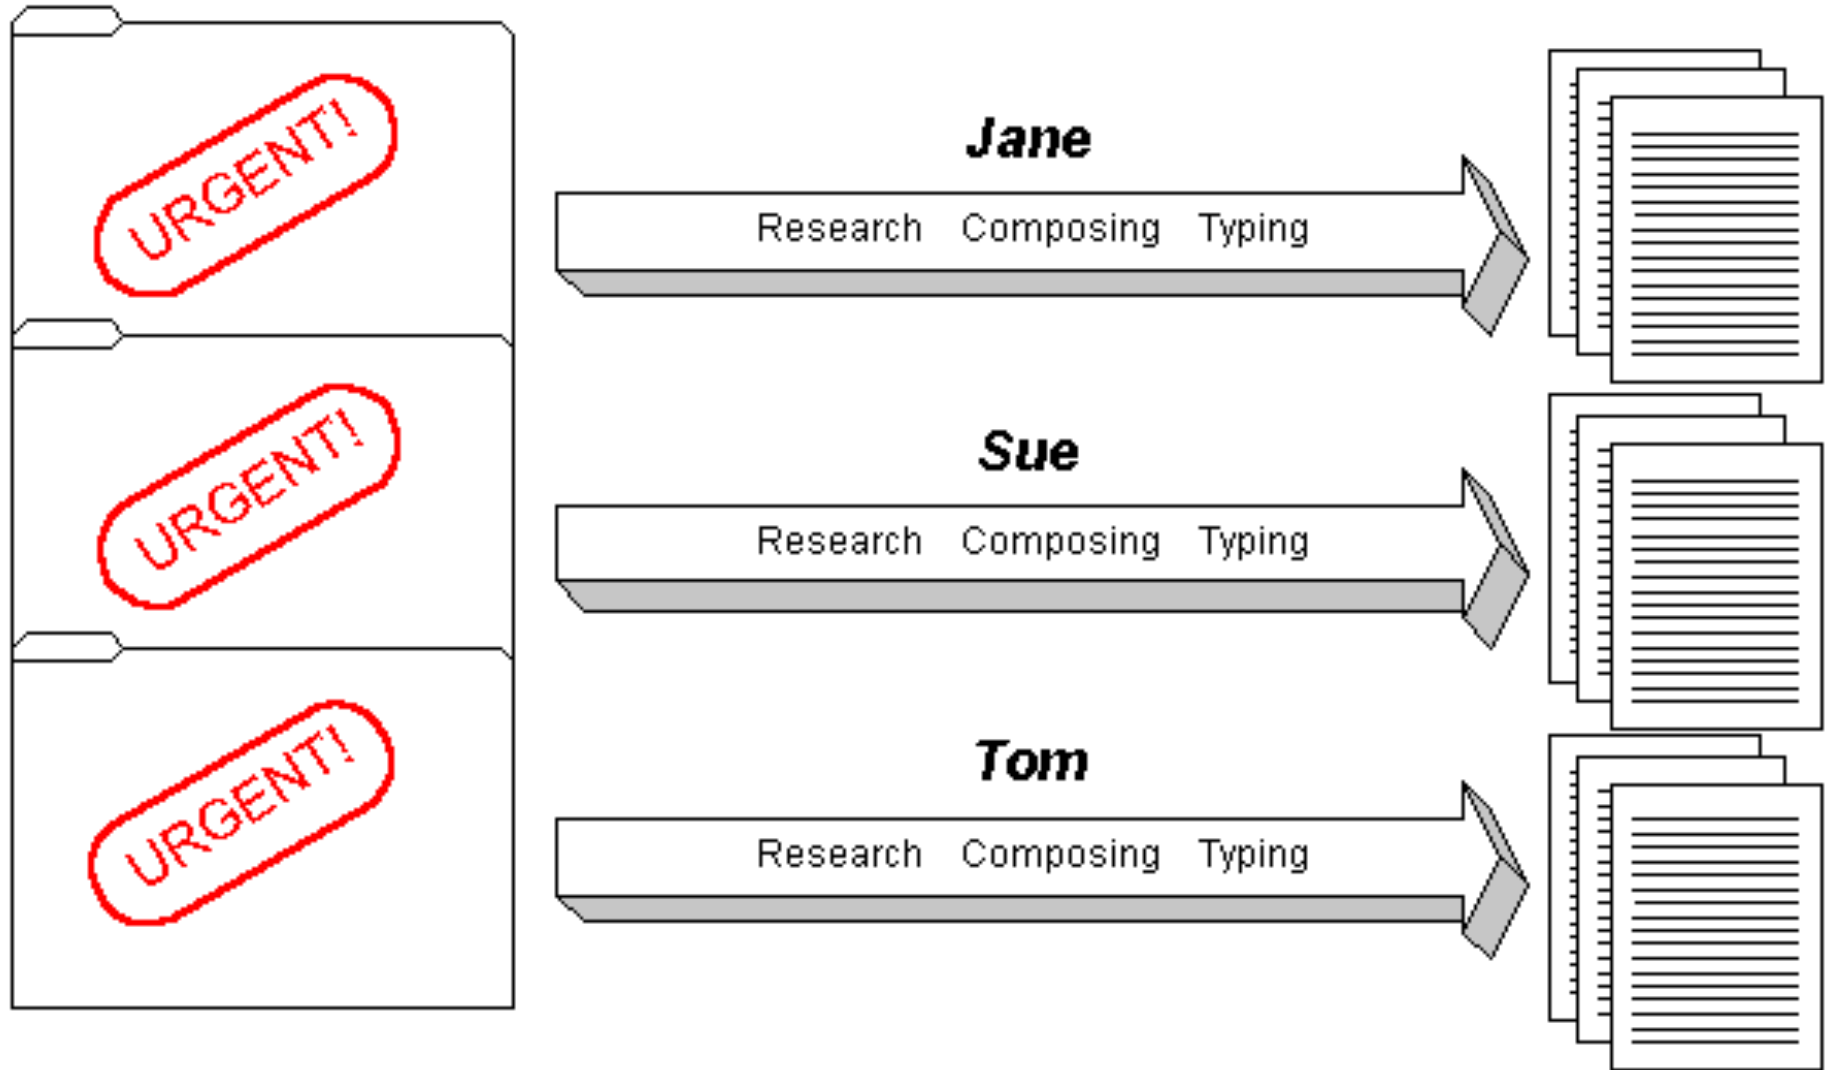
Jim Gray
Turing Award
B.S. Cal 1966
Ph.D. Cal 1969!



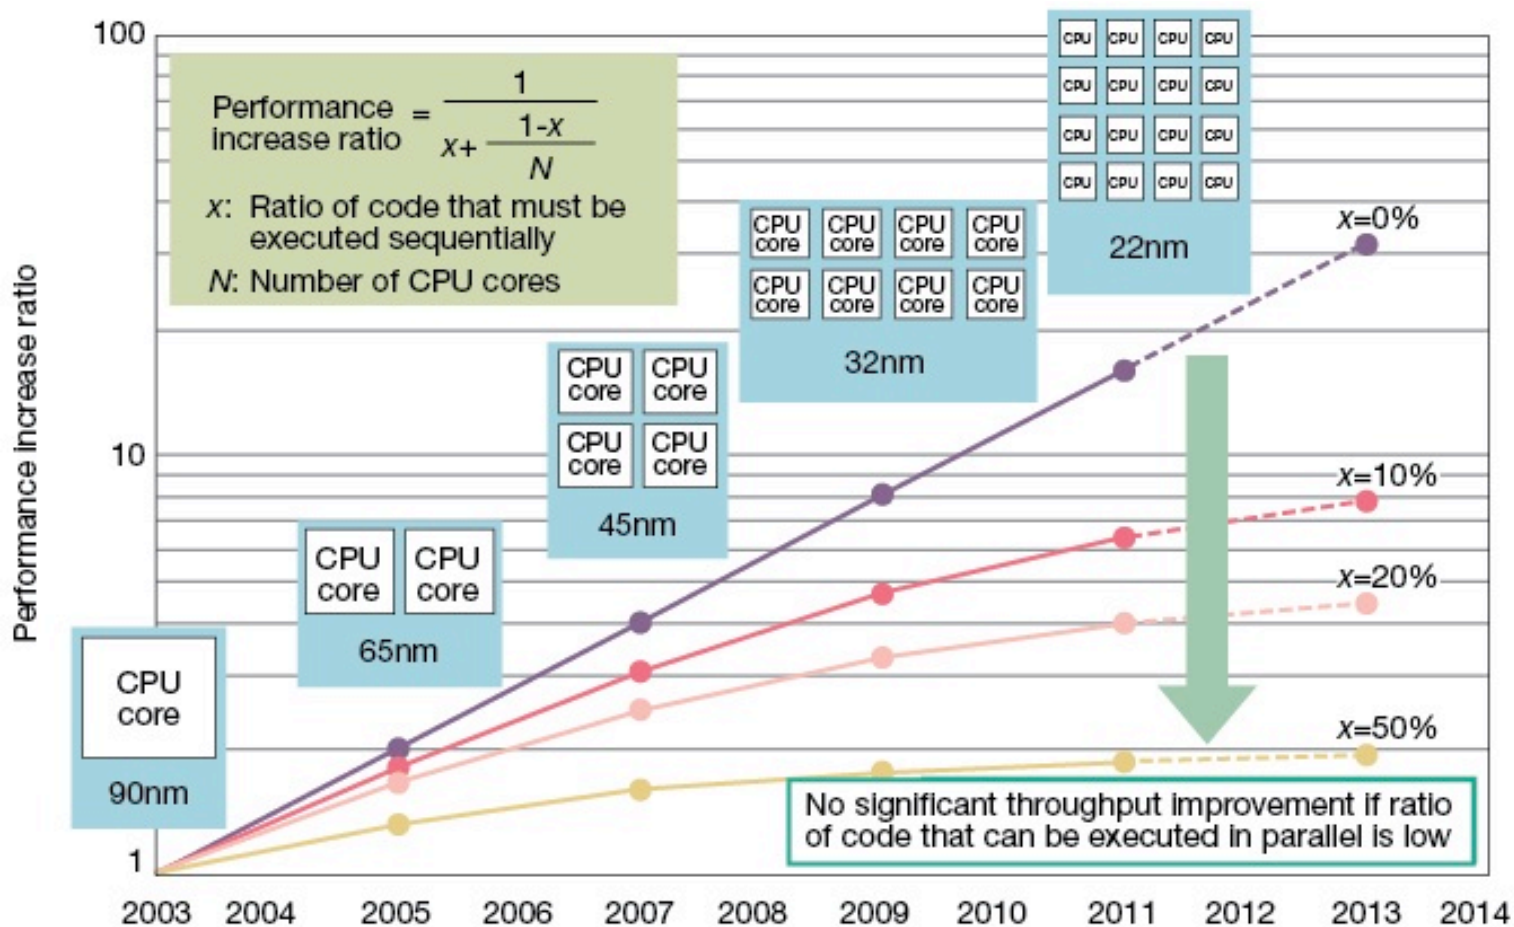
Great Idea #3: Principle of Locality/ Memory Hierarchy



Great Idea #4: Parallelism



Caveat: Amdahl's Law



Gene Amdahl
Computer Pioneer

Fig 3 Amdahl's Law an Obstacle to Improved Performance Performance will not rise in the same proportion as the increase in CPU cores. Performance gains are limited by the ratio of software processing that must be executed sequentially. Amdahl's Law is a major obstacle in boosting multicore microprocessor performance. Diagram assumes no overhead in parallel processing. Years shown for design rules based on Intel planned and actual technology. Core count assumed to double for each rule generation.

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

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

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

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 of so that can lose 1 bit but no data (Error Correcting Code/ECC Memory)



Agenda

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

Yoda says...

“Always in motion, the future is...”



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

Hot off the presses

- Everyone (on the waitlist), consider telling TeleBears you're moving to a more open section. We should be able to accommodate everyone, based on past experience.
- **Come to labs and discussion this week**
 - Switching Sections: if there's room (confirmed by TA in person), go ahead
 - Partners on ALL PROJECTS and LABS

Weekly Schedule

Weekly Schedule										
	Monday		Tuesday		Wednesday		Thursday		Friday	
8:00-9:00					LAB 22 Kevin L. 330 Soda					
9:00-10:00			LAB 15 Kevin L. 330 Soda		OH - Kevin Y. 611 Soda					
10:00-11:00	OH - William 411 Soda				LAB 23 William 330 Soda		DIS 115 Kevin L. 85 Evans		DIS 122 Kevin L. B51 Hildebrand	
11:00-12:00			LAB 16 Kevin Y. 330 Soda				OH - Kevin L. 651 Soda		DIS 116 Kevin Y. 310 Soda	
12:00-1:00	OH - Sagar 411 Soda				LAB 24 William 330 Soda				DIS 124 William 75 Evans	
1:00-2:00			LAB 17 Sagar 330 Soda						OH - Jeffrey 611 Soda	
2:00-3:00	LECTURE 1 Pimentel				LECTURE 1 Pimentel		DIS 117 Sagar 285 Cory		LECTURE 1 Pimentel	
3:00-4:00	LAB 11 Sung Roa 330 Soda	OH - Alan 611 Soda	LAB 18 Sagar 330 Soda		DIS 111 Sung Roa 71 Evans		DIS 118 Sagar 405 Soda			
4:00-5:00									OH - Sung Roa 611 Soda	
5:00-6:00	LAB 12 Sung Roa 330 Soda		LAB 19 Jeffrey 330 Soda		DIS 112 Sung Roa 3105 Etcheverry	DIS 114 Roger 24 Wheeler	DIS 119 Jeffrey 75 Evans	DIS 121 Shreyas 71 Evans		
6:00-7:00							DIS 113 Roger 75 Evans			
7:00-8:00	LAB 13 Roger 330 Soda		LAB 20 Alan 330 Soda				DIS 120 Alan 75 Evans			
8:00-9:00										
9:00-10:00	LAB 14 Roger 330 Soda		LAB 21 Shreyas 330 Soda							
10:00-11:00										

Course Information

- Course Web: <http://inst.eecs.Berkeley.edu/~cs61c/>
- Instructors:
 - Dan Garcia
- Teaching Assistants: (see webpage)
- Textbooks: Average 15 pages of reading/week (can rent!)
 - Patterson & Hennessey, *Computer Organization and Design*, 5/e (we'll try to provide 4th Ed pages, not Asian version 4th edition)
 - Kernighan & Ritchie, *The C Programming Language*, 2nd Edition
 - Barroso & Holzle, *The Datacenter as a Computer*, 1st Edition
- Piazza:
 - Every announcement, discussion, clarification happens there

Course Organization

- Grading
 - EPA: Effort, Participation and Altruism (5%)
 - Homework (10%)
 - Labs (5%)
 - Projects (20%)
 1. Non-Parallel Application (MIPS & C)
 2. Data Parallelism (Map-Reduce on Amazon EC2)
 3. Parallelize Project1, SIMD, MIMD
 4. Computer Processor Design (Logisim)
 - Performance Competition for honor (and EPA)
 - Midterm (25%): 2014-03-03, can be clobbered!
 - Final (35%): 2014-05-13 @ 11:30am-2:30pm

Tried-and-True Technique: Peer Instruction



- Increase real-time learning in lecture, test understanding of concepts vs. details
- As complete a “segment” ask multiple choice question
 - 1-2 minutes to decide yourself
 - 2 minutes in pairs/triples to reach consensus.
 - Teach others!
 - 2 minute discussion of answers, questions, clarifications
- You can get transmitters from the ASUC bookstore OR you can use i>clicker GO app for less!
 - We’ll start this on Friday



EECS Grading Policy

- <http://www.eecs.berkeley.edu/Policies/ugrad.grading.shtml>

“A typical GPA for courses in the lower division is 2.7. This GPA would result, for example, from 17% A's, 50% B's, 20% C's, 10% D's, and 3% F's. A class whose GPA falls outside the range 2.5 - 2.9 should be considered atypical.”

- Fall 2010: GPA 2.81
26% A's, 47% B's, 17% C's,
3% D's, 6% F's
- Job/Intern Interviews: They grill you with technical questions, so it's what you say, not your GPA (New 61C gives good stuff to say)

	Fall	Spring
2010	2.81	2.81
2009	2.71	2.81
2008	2.95	2.74
2007	2.67	2.76

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 tooooo 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 HW0, 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

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.



The Unplugged Challenge



Jim Wilson/The New York Times

Loren Frank, a professor of physiology, said downtime lets the brain go over experiences, "solidify them and turn them into permanent memory."


Just another day at the gym.

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

 FACEBOOK

 TWITTER

RECOMMEND


 COMMENTS
(206)

 SIGN IN TO E-MAIL

 PRINT

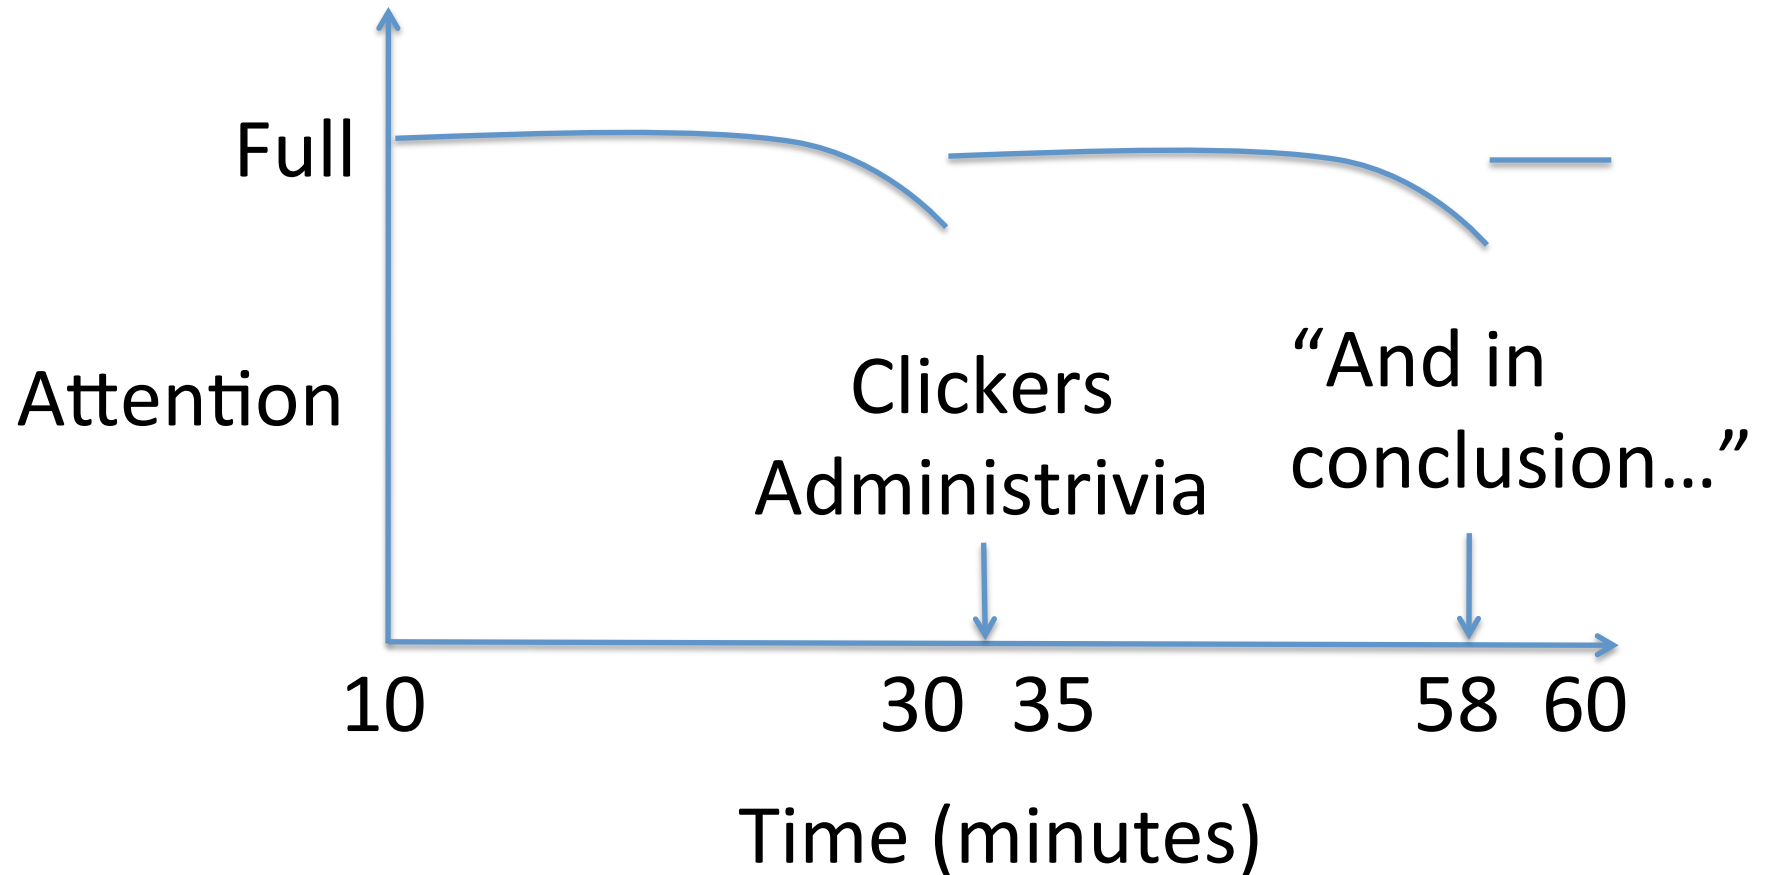
 SINGLE PAGE

 REPRINTS

 SHARE



Architecture of a typical Lecture



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