

# CS 110

## Computer Architecture

### Review for Midterm I

Instructor:  
Sören Schwertfeger

<http://shitech.org/courses/ca/>

School of Information Science and Technology SIST

ShanghaiTech University

Slides based on UC Berkley's CS61C

# Midterm I

- Date: Thursday, Apr. 19
- Time: 10:15 - 12:15 (normal lecture slot)
  - Be punctual – we start 10:15 sharp!
- Venue: Teaching Center 301 + 302
- One empty seat between students
- Closed book:
  - You can bring one A4 page with notes (both sides; English preferred; Chinese is OK): Write your Chinese and **Pinyin** name on the top! **Handwritten** by you!
  - You will be provided with the MIPS “green sheet”
  - No other material allowed!

# Midterm I

- Switch cell phones **off**! (not silent mode – off!)
  - Put them in your bags.
- Bags under the table. Nothing except paper, pen, 1 drink, 1 snack on the table!
- No other electronic devices are allowed!
  - No ear plugs, music, smartwatch...
- Anybody touching any electronic device will **FAIL** the course!
- Anybody found cheating (copy your neighbors answers, additional material, ...) will **FAIL** the course!









# COMPUTER ORGANIZATION AND DESIGN

THE HARDWARE/SOFTWARE INTERFACE

FIFTH EDITION

DAVID A. PATTERSON  
JOHN L. HENNESSY









# SIFT REFERENCE GUIDE (V.1.1) – CREATING TIMELINES WITH THE SIFT WORKSTATION

1. VISIT: <http://computer-forensics11.sans.org/community/downloads>

2. BOOT SIFT VM

3. ELEVATE PRIVS

4. CONNECT IMAGE TO SIFT

5

**SANS COMPUTER FORENSICS and INCIDENT RESPONSE**

THE PURPOSE OF THIS REFERENCE GUIDE IS TO WALK THROUGH THE PROCESS OF BOOTING THE SIFT WORKSTATION, CREATING A TIMELINE ("SUPER" OR "MICRO") AND REVIEWING IT.

## HOW TO CALCULATE THE OFFSET FOR MOUNTING

1. Run `mmls` to query partition layout
2. Identify partition and byte offset
3. (Partition byte offset) x (bytes per sector) = offset ##### to use!

Example: 63 X 512 = 32256

Note: If needed, repeat for each partition. Make new mount point: `# mkdir /mnt/windows_mount2/`

6. log2timeline default timezone is set to examiner local host. To change use `-z [TIMEZONE]` option. To list all available timezones: `# log2timeline -z list`

**log2timeline PARSING PLUGINS**

apache2\_error - Apache2 error log file

chrome - Chrome history file

encase\_dirlisting - CSV file that is exported from encase

evtx - Windows 2k/XP/2k3 Event Log

evtx - Windows Event Log File (EVTX)

exif - Metadata information from files using ExifTool

ff\_bookmark - Firefox bookmark file

firefox2 - Firefox 2 browser history

firefox3 - Firefox 3 history file

ftk\_dirlisting - CSV file that is exported from FTK Imager (dirlisting)

generic\_linux - Generic Linux logs that start with MMM DD HH:MM:SS

iehistory - index.dat file containing IE history

iis - IIS W3C log file

isatxt - ISA text export log file

jp\_ntfs\_change - CSV output file from JP (NTFS Change log)

mactime - Body file in the mactime format

mcafee - Log file

mft - NTFS MFT file

mssql\_errorlog - ERRORLOG file produced by MS SQL server

ntuser - NTUSER.DAT registry file

opera - Opera's global history file

oxml - OpenXML document pcap

pcap - PCAP file

pdf - Available PDF document metadata

prefetch - Prefetch directory

recycler - Recycle bin directory

restore 0.9 - Restore point directory

safari - Safari History.plist file

sam - SAM registry file

security - SECURITY registry file

setupapi - SetupAPI log file in Windows XP

skype\_sql - Skype database

software - SOFTWARE registry file

sol - .sol (LSO) or a Flash cookie file

squid - Squid access log (http emulate off)

syslog - Linux Syslog log file

system - SYSTEM registry file

tlf - Body file in the TLF format

volatility - Volatility output files (psscan2, socks2, ...)

win\_link - Windows shortcut file (or a link file)

wmiprov - WMI log file

xpfirewall - XP Firewall log

List plugins `# log2timeline -f list`  
...HELP EXPAND THIS LIST. BUILD PLUGINS!!!

BY DAVID NIDES (12/16/2011)  
TWITTER: @DAVNADS  
BLOG: DAVNADS.BLOGSPOT.COM  
EMAIL: DNIDES@KPMG.COM  
CREDITS TO: ED GOINGS, ROB LE, KRISTINN GUDBJONSSON, KPMG & QUESTIONS/FEEDBACK-CONTACT US

**KEY**

Red text – image/source

Blue text – mount point

Purple text – output file

Green text – log2timeline plugins

Brown text – TimeZone

## 5. HARD DRIVE MOUNTING (if you are using log2timeline-sift and Single DD you can skip to 7-A)

### SINGLE OR SPLIT IMAGE (2 options):

`# mount -t ntfs -o ro,loop,show_sys_files,streams_interface=windows,offset=##### /mnt/ewf/<image> /mnt/windows_mount/`

`# mount -t ntfs -o ro,loop,show_sys_files,streams_interface=windows,offset=##### image.dd /mnt/windows_mount/`

Not Needed For 7-A

### MOUNT TO MOUNT POINT

### SINGLE IMAGE

`# mount -t ntfs -o ro,loop,show_sys_files,streams_interface=windows,offset=##### image.dd /mnt/windows_mount/`

### SPLIT IMAGE (2 step process)

`# affuse image.001 /mnt/aff`

`# mount -t ntfs-3g -o loop,ro,show_sys_files,streams_interface=windows,offset=##### /mnt/aff/<image> /mnt/windows_mount/`

## 7-A: AUTOMATED SUPER TIMELINE CREATION

`log2timeline-sift -o -z [TIMEZONE] -p [PARTITION #] -i [IMAGE FILE]`

### DISK IMAGE (prompt for partition, mount, and run):

XP `# log2timeline-sift -z EST5EDT -i image`

WIN7 `# log2timeline-sift -win7 -z EST5EDT -i image`

### FOR PARTITION (mount and run using all applicable plugins)

XP `# log2timeline-sift -z EST5EDT -p 0 -i partition`

WIN7 `# log2timeline-sift -win7 -z EST5EDT -i partition`

### OTHER USAGE EXAMPLES:

Display list of available plugins: `# log2timeline -f list`

Run log2timeline using only specific plugins: `# log2timeline-sift -p prefetch -z EST5EDT -i image.dd`

Help (man page): `# log2timeline -h`

## 8. CSV FILE OUTPUT (/cases/timeline-output-folder)

-date: Date of the event, in the format of MM/DD/YYYY

-time: Time of day, expressed in a 24h format, HH:MM:SS

-timezone: The timezone that was used to call the tool with.

-source: MACB meaning of the fields, comp w/ mactime format.

-source: Source short name (i.e. registry entries are REG)

-sourcetype: Desc of the source ("Internet Explorer" instead of WEBHIST)

-type: Timestamp type (i.e. "Last Accessed", "Last Written")

-user: Username associated with the entry, if one is available.

-host: Hostname associated with the entry, if one is available.

-short: Contains less text than the full description field.

-desc: where majority info is stored, the actual parsed desc of the entry.

-version: Version number of the timestamp object.

-filename: Filename with the full path that contained the entry

-inode: inode number of the file being parsed.

-notes: Some input modules insert additional information in the form of a note, which comes here. Or it can be used during the review.

-format: Input module name used to parse the file.

-extra: Additional information parsed is joined together and put here.

## 7-B: MANUAL "MICRO" TIMELINE CREATION

`log2timeline-sift -o -z [TIMEZONE] [-f FORMAT] [-z TIMEZONE] [-o OUTPUT MODULE] [-w LOG_FILE/LOG_DIR [-] [FORMAT FILE OPTIONS]]`

### EXTRACT METADATA (using log2timeline or fls)

Extract system data w/ log2timeline from mounted file system:

`# log2timeline -f mft -o mactime -r -z EST5EDT -w`

OR Extract metadata using Sleuthkit:

`# fls -m "" -o offset -d > fls.body`

Convert body file format to mactime format w/ mactime:

`# mactime -b fls.body -d`

### ARTIFACTS (run l2l on mounted file system with plugins recursively)

Extract artifacts w/ log2timeline and run on mounted file system:

`# log2timeline -f firefox3,chrome -o mactime -r -z EST5EDT -w`

Convert body file format to CSV format w/ mactime:

`# mactime -b log2timeline.body -d > log2timeline.csv`

## 9. FILTER TIMELINE

Filter timeline with date range to include only:

`l2t_process -b timeline.csv MM-DD-YYYY..MM-DD-YYYY > filtered.csv`

Filter timeline with keyword list (one term per line in keywords.txt):

`l2t_process -b timeline.csv -k keywords.txt > filtered.csv`

What sources are in your timeline?

`awk -F, '{print $6;}' timeline.csv | grep -v sourcetype | sort | uniq`

Find all LNK files that reference E Drive

`grep "Shortcut LNK" timeline.csv | grep "E:"`

Find MountPoints2 entries that reference E Drive

`grep "MountPoints2 key" timeline.csv | grep "E drive"`

`grep USB timeline.csv | grep "SetupAPILog"`

File System	M	A	C	B
Ext2/3	Modified	Accessed	Changed	N/A
FAT	Written	Accessed	N/A	Created
NTFS	File Modified	Accessed	MFT Modified	Created
UFS	Modified	Accessed	Changed	N/A

## 10. CONNECT TO SIFT

1. SIFT Desktop -> SETTINGS -> OPTIONS -> Shared Folders -> Always Enabled (Check)
2. SIFT Desktop -> VMware-Shared-Drive
3. Access from a Win Machine \\SIFTWORKSTATION

## 11. REVIEW TIMELINE

- Review timelines using:
- Open, Soft, Filter with Excel
  - Import into SPLUNK
  - SIMILE
  - Tapestry

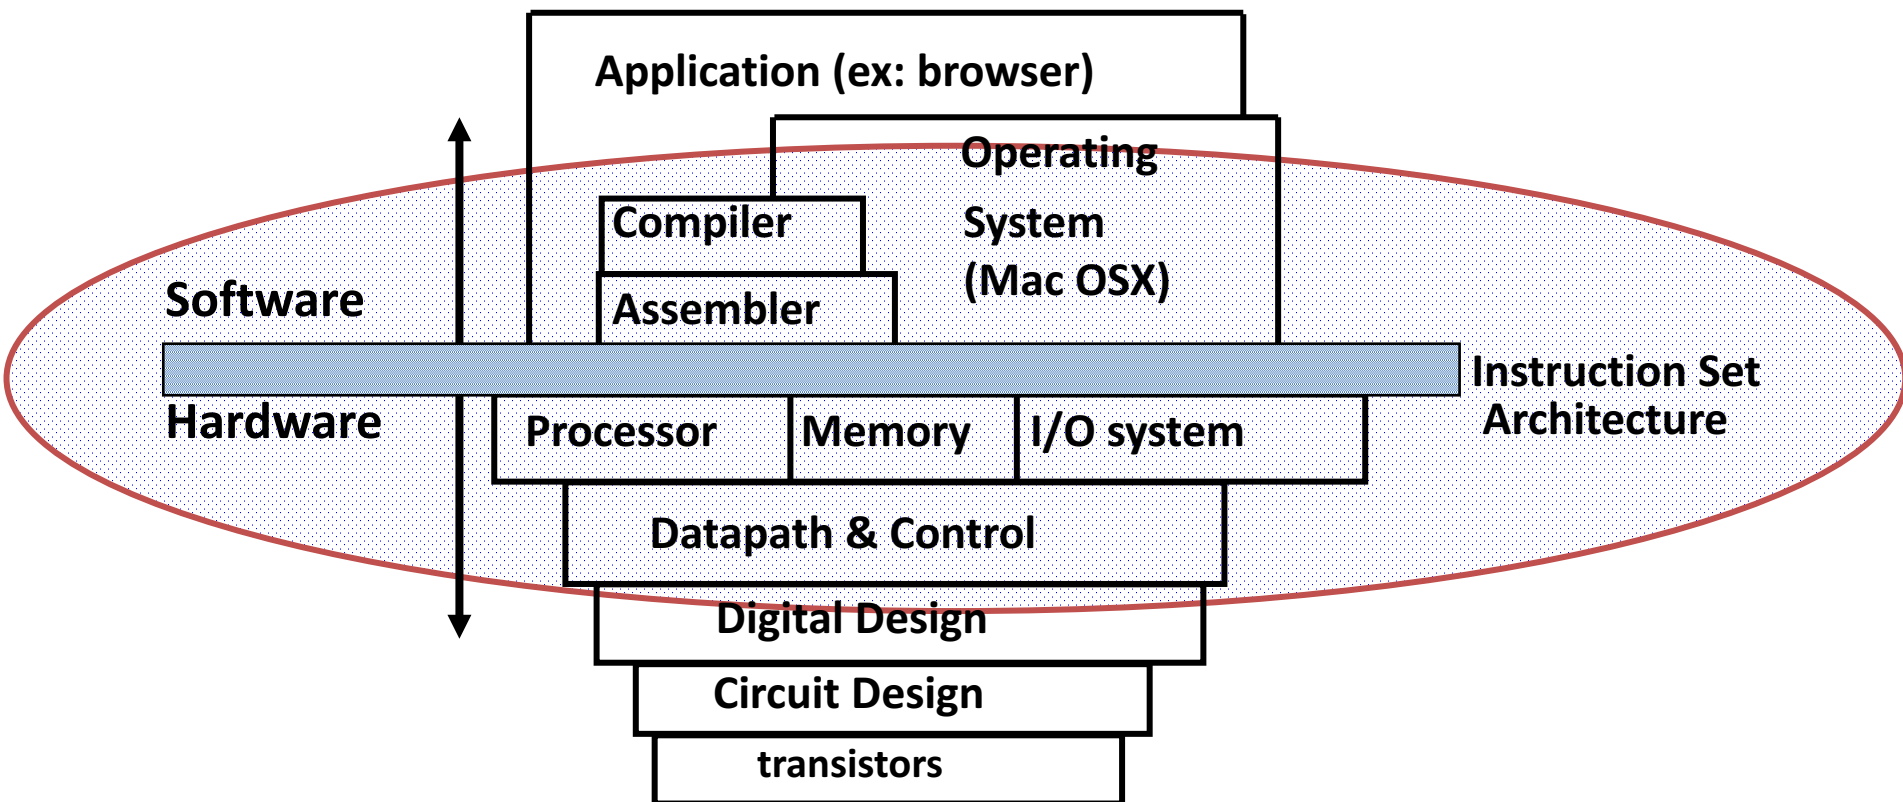
# Midterm I

- Ask questions today!
- Discussion is Q&A session
  - Suggest topics for review in piazza!
- This review session does not/ can not cover all possible topics!

# Content

- Main topics
  - Number representation
  - C
  - MIPS
- Plus general "Computer Architecture" knowledge
- Everything till lecture 8 CALL – including lecture 8

# Old School Machine Structures





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

*Software*

*Hardware*

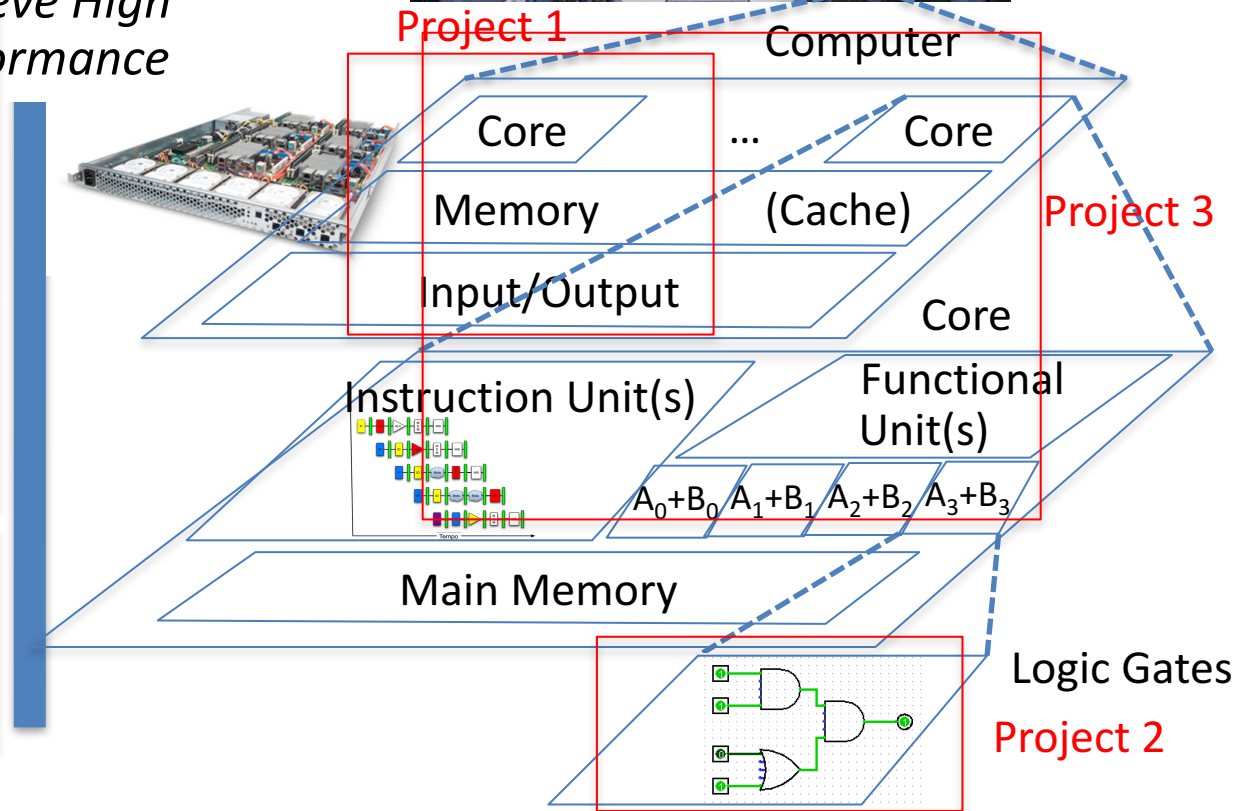
Warehouse  
-Scale  
Computer

Smart  
Phone



*Harness  
Parallelism &  
Achieve High  
Performance*

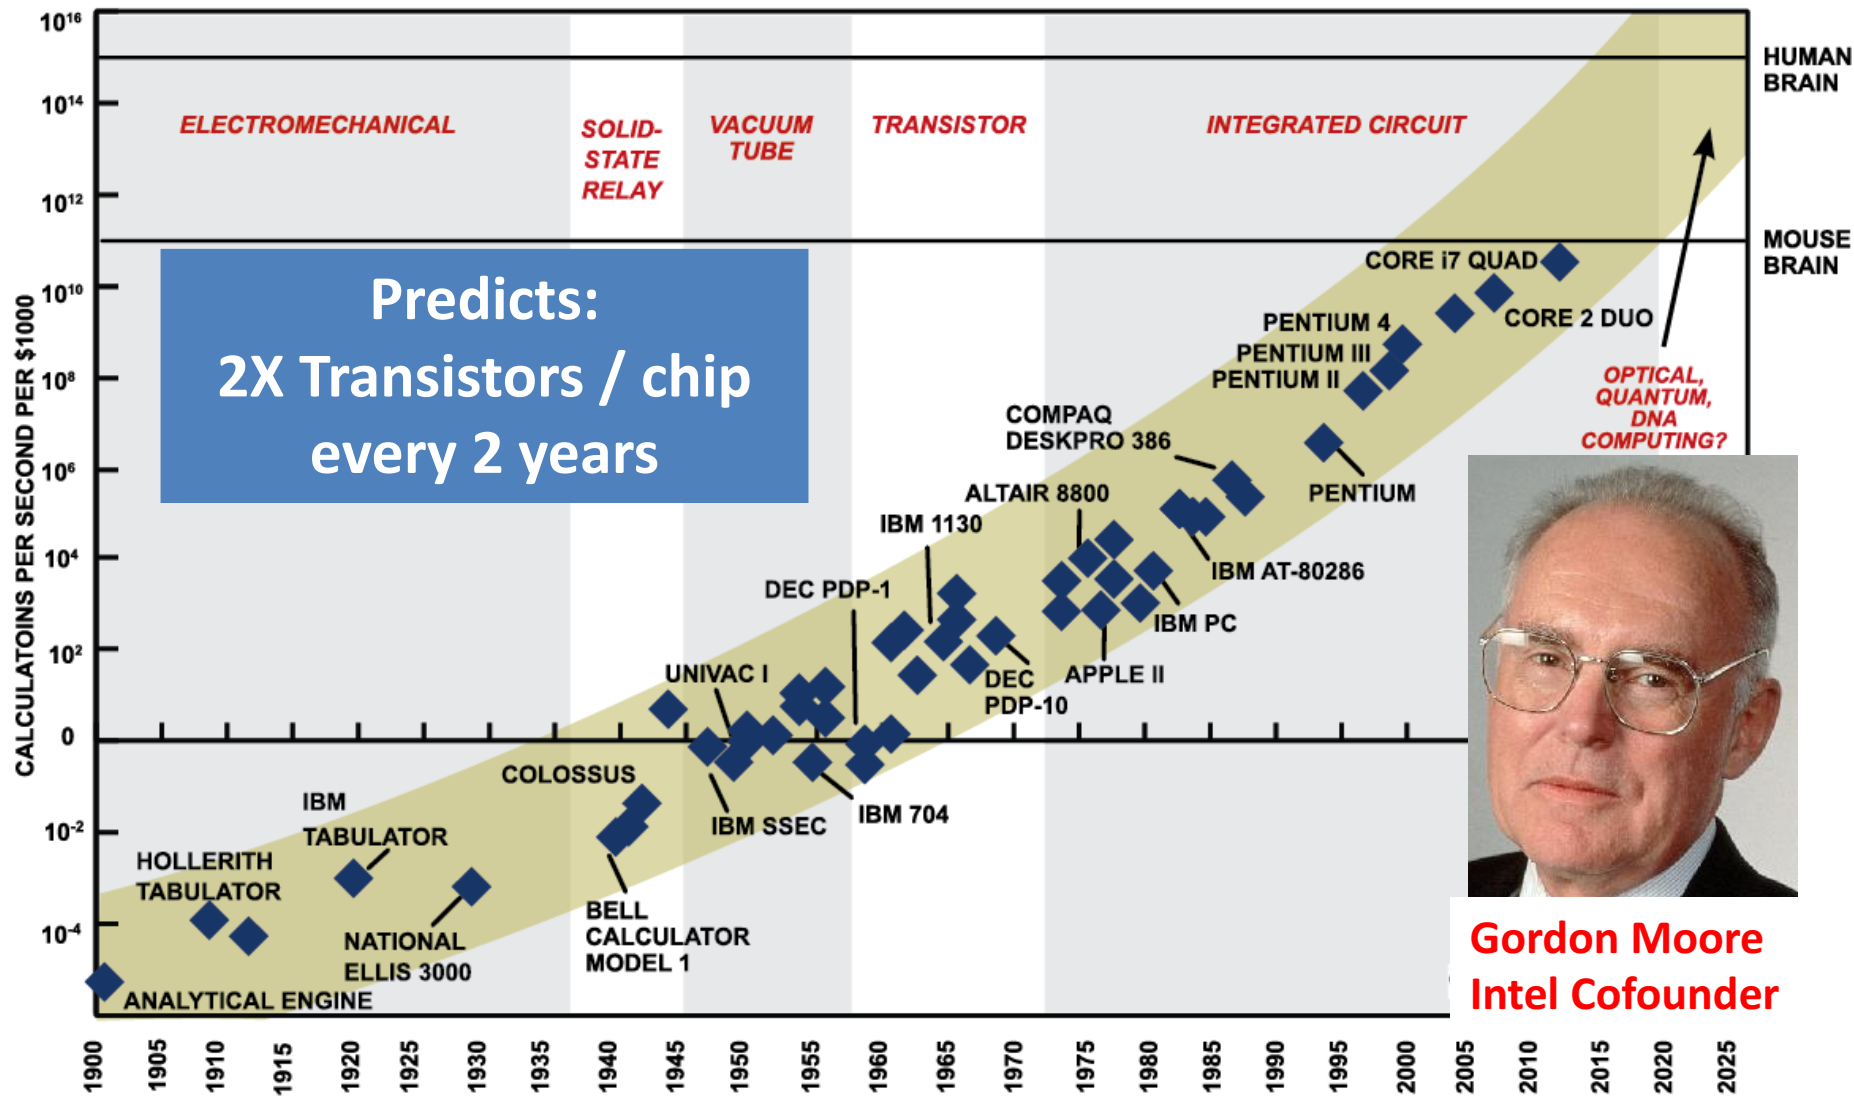
- **Parallel Requests**  
Assigned to computer  
e.g., Search “cats”
- **Parallel Threads**  
Assigned to core  
e.g., Lookup, Ads
- **Parallel Instructions**  
>1 instruction @ one time  
e.g., 5 pipelined instructions
- **Parallel Data**  
>1 data item @ one time  
e.g., Add of 4 pairs of words
- **Hardware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



# 6 Great Ideas in Computer Architecture

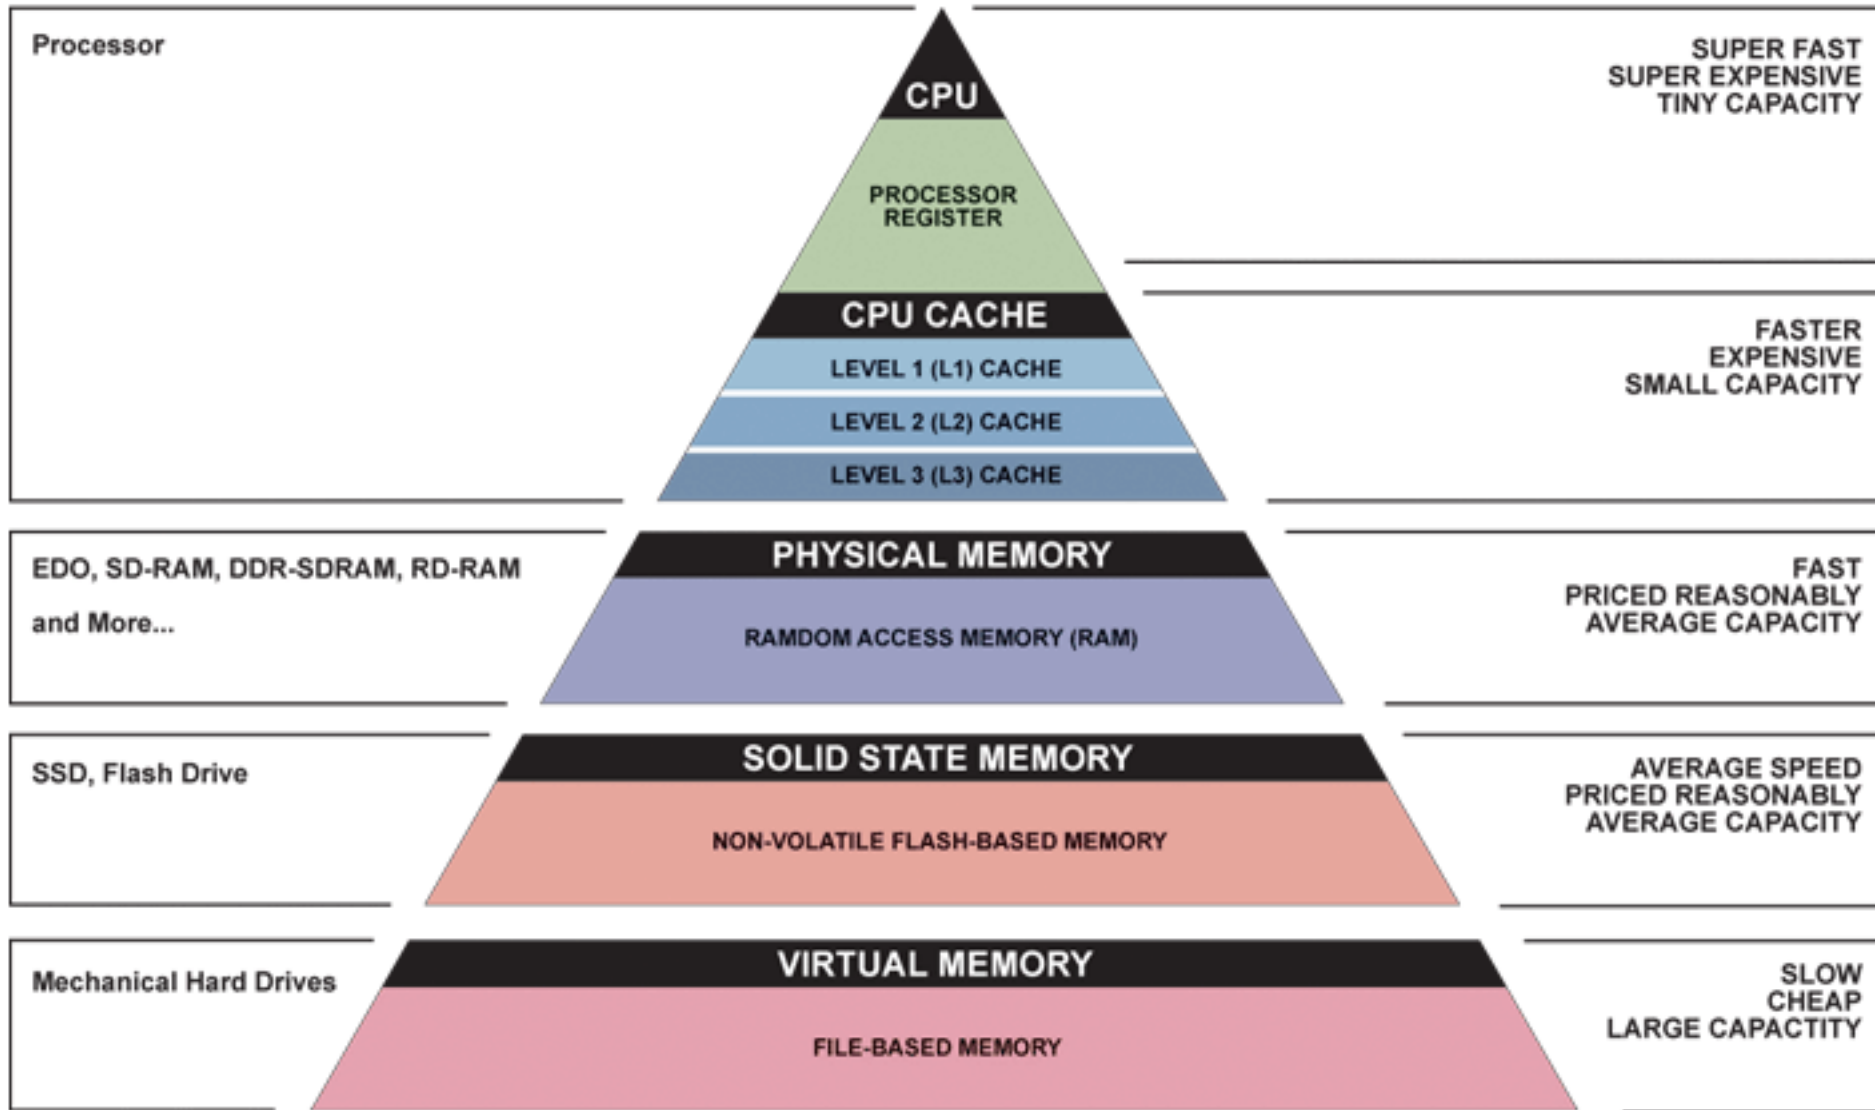
1. Abstraction  
(Layers of Representation/Interpretation)
2. Moore's Law (Designing through trends)
3. Principle of Locality (Memory Hierarchy)
4. Parallelism
5. Performance Measurement & Improvement
6. Dependability via Redundancy

# #2: Moore's Law

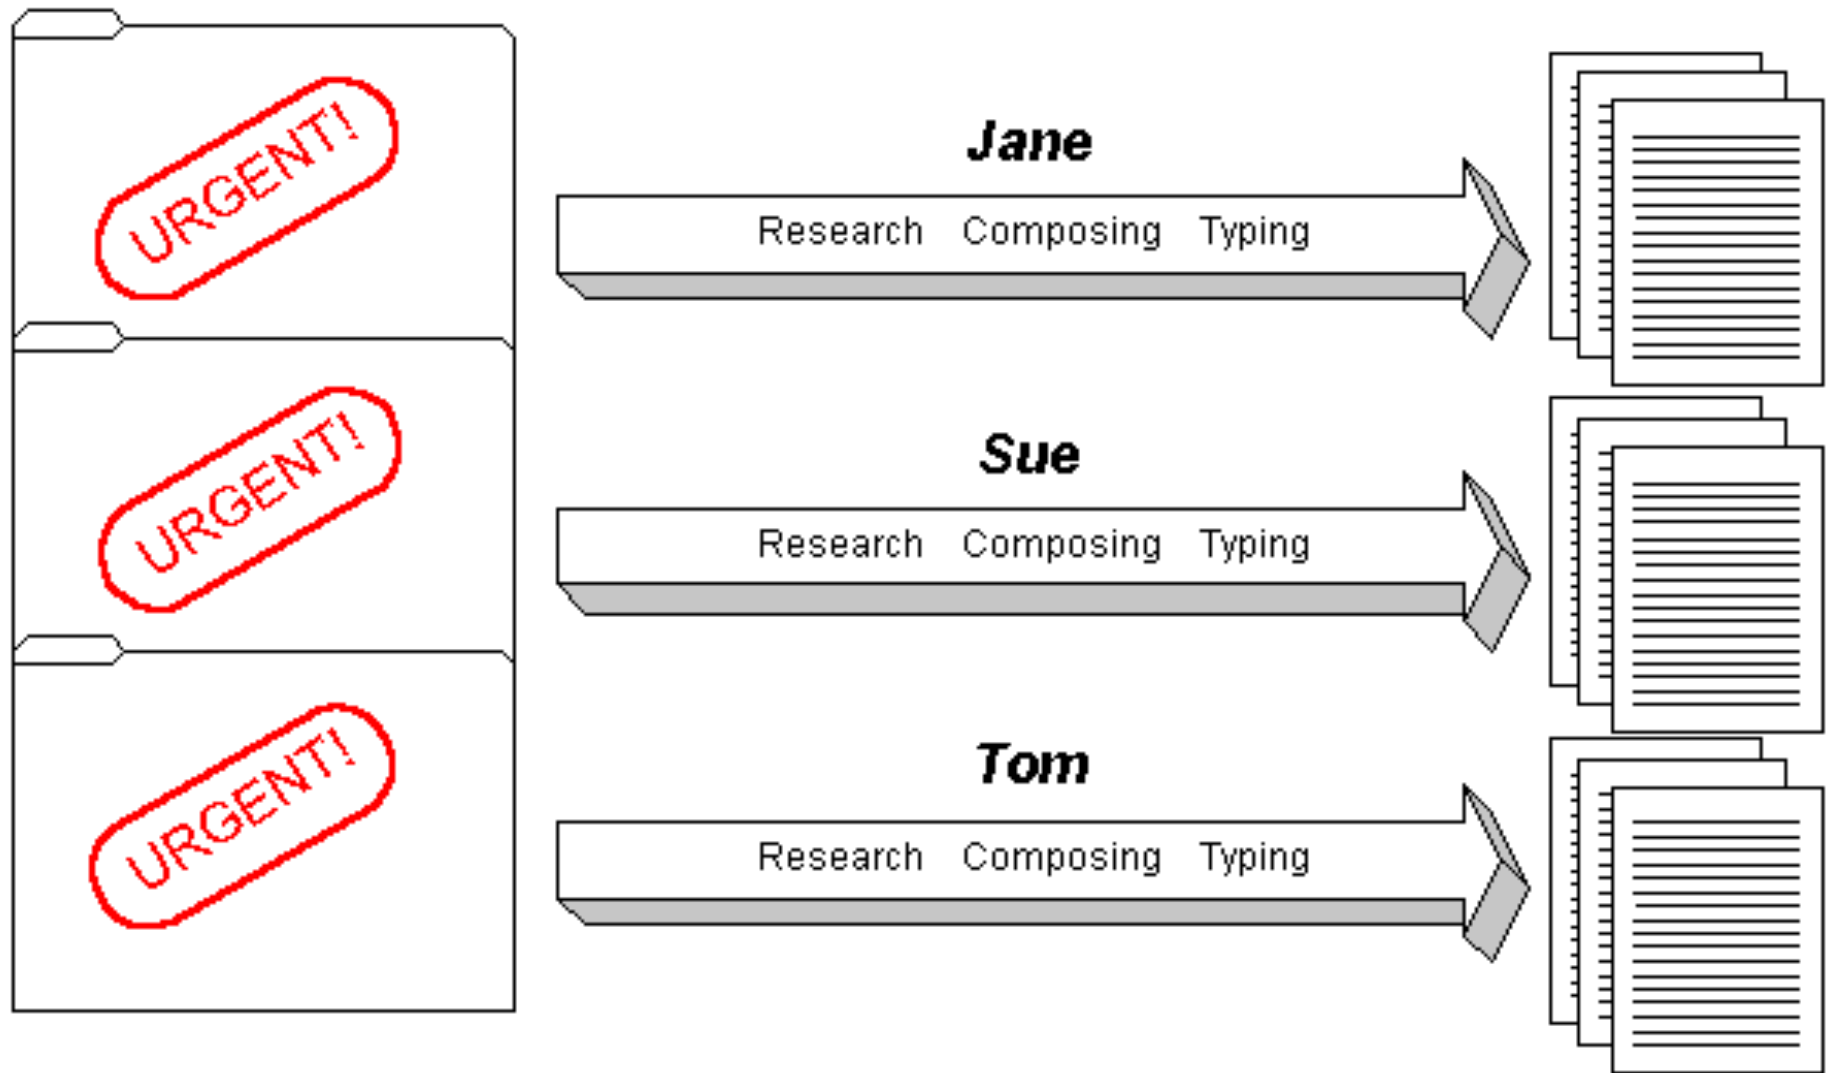


SOURCE: RAY KURZWEIL, "THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY", P.67, THE VIKING PRESS, 2006. DATAPOINTS BETWEEN 2000 AND 2012 REPRESENT BCA ESTIMATES.

# Great Idea #3: Principle of Locality/ Memory Hierarchy



# Great Idea #4: Parallelism



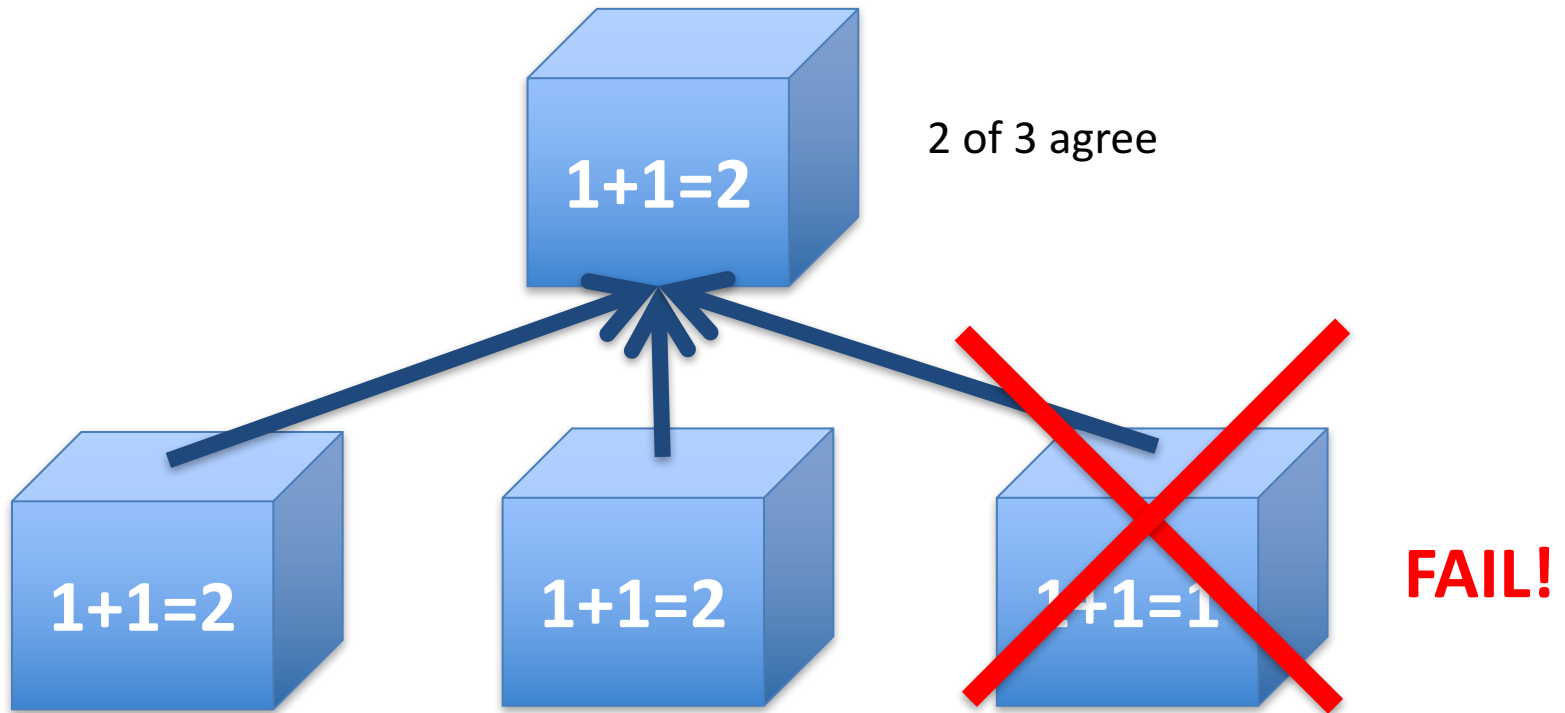
# Great Idea #5: Performance Measurement and Improvement

- Tuning 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*

# 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

# Key Concepts

- Inside computers, everything is a number
- But numbers usually stored with a fixed size
  - 8-bit bytes, 16-bit half words, 32-bit words, 64-bit double words, ...
- Integer and floating-point operations can lead to results too big/small to store within their representations: *overflow/underflow*



# Number Representation

# Number Representation

- Value of i-th digit is  $d \times Base^i$  where i starts at 0 and increases from right to left:
- $123_{10} = 1_{10} \times 10_{10}^2 + 2_{10} \times 10_{10}^1 + 3_{10} \times 10_{10}^0$   
 $= 1 \times 100_{10} + 2 \times 10_{10} + 3 \times 1_{10}$   
 $= 100_{10} + 20_{10} + 3_{10}$   
 $= 123_{10}$
- Binary (Base 2), Hexadecimal (Base 16), Decimal (Base 10) different ways to represent an integer
  - We use  $1_{\text{two}}, 5_{\text{ten}}, 10_{\text{hex}}$  to be clearer  
(vs.  $1_2, 4_8, 5_{10}, 10_{16}$ )

# Number Representation

- Hexadecimal digits:  
0,1,2,3,4,5,6,7,8,9,A,B,C,D,E,F
- $$\begin{aligned} \text{FFF}_{\text{hex}} &= 15_{\text{ten}} \times 16_{\text{ten}}^2 + 15_{\text{ten}} \times 16_{\text{ten}}^1 + 15_{\text{ten}} \times 16_{\text{ten}}^0 \\ &= 3840_{\text{ten}} + 240_{\text{ten}} + 15_{\text{ten}} \\ &= 4095_{\text{ten}} \end{aligned}$$
- $1111\ 1111\ 1111_{\text{two}} = \text{FFF}_{\text{hex}} = 4095_{\text{ten}}$
- May put blanks every group of binary, octal, or hexadecimal digits to make it easier to parse, like commas in decimal

# Signed Integers and Two's-Complement Representation

- Signed integers in C; want ½ numbers  $<0$ , want ½ numbers  $>0$ , and want one 0
- *Two's complement* treats 0 as positive, so 32-bit word represents  $2^{32}$  integers from  $-2^{31}$  ( $-2,147,483,648$ ) to  $2^{31}-1$  ( $2,147,483,647$ )
  - Note: one negative number with no positive version
  - Book lists some other options, all of which are worse
  - Every computer uses two's complement today
- *Most-significant bit* (leftmost) is the *sign bit*, since 0 means positive (including 0), 1 means negative
  - Bit 31 is most significant, bit 0 is least significant

# Two's-Complement Integers

Sign Bit

0000 0000 0000 0000 0000 0000 0000 0000<sub>two</sub> = 0<sub>ten</sub>

0000 0000 0000 0000 0000 0000 0000 0001<sub>two</sub> = 1<sub>ten</sub>

0000 0000 0000 0000 0000 0000 0000 0010<sub>two</sub> = 2<sub>ten</sub>

...

...

0111 1111 1111 1111 1111 1111 1111 1101<sub>two</sub> = 2,147,483,645<sub>ten</sub>

0111 1111 1111 1111 1111 1111 1111 1110<sub>two</sub> = 2,147,483,646<sub>ten</sub>

0111 1111 1111 1111 1111 1111 1111 1111<sub>two</sub> = 2,147,483,647<sub>ten</sub>

---

1000 0000 0000 0000 0000 0000 0000 0000<sub>two</sub> = -2,147,483,648<sub>ten</sub>

1000 0000 0000 0000 0000 0000 0000 0001<sub>two</sub> = -2,147,483,647<sub>ten</sub>

1000 0000 0000 0000 0000 0000 0000 0010<sub>two</sub> = -2,147,483,646<sub>ten</sub>

...

...

1111 1111 1111 1111 1111 1111 1111 1101<sub>two</sub> = -3<sub>ten</sub>

1111 1111 1111 1111 1111 1111 1111 1110<sub>two</sub> = -2<sub>ten</sub>

1111 1111 1111 1111 1111 1111 1111 1111<sub>two</sub> = -1<sub>ten</sub>

# Ways to Make Two's Complement

- For N-bit word, complement to  $2_{\text{ten}}^N$ 
  - For 4 bit number  $3_{\text{ten}} = 0011_{\text{two}}$ , two's complement (i.e.  $-3_{\text{ten}}$ ) would be

$$16_{\text{ten}} - 3_{\text{ten}} = 13_{\text{ten}} \text{ or } 10000_{\text{two}} - 0011_{\text{two}} = 1101_{\text{two}}$$

- Here is an easier way:

- Invert all bits and add 1

$$\begin{array}{r} 3_{\text{ten}} \quad 0011_{\text{two}} \\ \text{Bitwise complement} \quad 1100_{\text{two}} \\ + \quad 1_{\text{two}} \\ \hline -3_{\text{ten}} \quad 1101_{\text{two}} \end{array}$$

- Computers actually do it like this, too

# Two's-Complement Examples

- Assume for simplicity 4 bit width, -8 to +7 represented

$$\begin{array}{r} 3 \quad 0011 \\ +2 \quad 0010 \\ \hline 5 \quad 0101 \end{array}$$

$$\begin{array}{r} 3 \quad 0011 \\ + (-2) \quad 1110 \\ \hline 1 \quad 1 \quad 0001 \end{array}$$

$$\begin{array}{r} -3 \quad 1101 \\ + (-2) \quad 1110 \\ \hline -5 \quad 1 \quad 1011 \end{array}$$

*Overflow when  
magnitude of result  
too big small to fit  
into result  
representation*

$$\begin{array}{r} 7 \quad 0111 \\ +1 \quad 0001 \\ \hline -8 \quad 1000 \end{array}$$

$$\begin{array}{r} -8 \quad 1000 \\ + (-1) \quad 1111 \\ \hline +7 \quad 1 \quad 0111 \end{array}$$

Carry into MSB =  
Carry Out MSB

*Overflow!*

*Overflow!*

Carry into MSB  $\neq$   
Carry Out MSB

*Carry in = carry from less significant bits  
Carry out = carry to more significant bits*

Suppose we had a 5-bit word. What integers can be represented in two's complement?

☐ -32 to +31

☐ 0 to +31

☐ -16 to +15

☐ -15 to +16



Suppose we had a 5-bit word. What integers can be represented in two's complement?

☐ -32 to +31

☐ 0 to +31

☒ -16 to +15

☐ -15 to +16

# C Programming

# Quiz: Pointers

```
void foo(int *x, int *y)
{   int t;
    if ( *x > *y ) { t = *y; *y = *x; *x = t; }
}
int a=3, b=2, c=1;
foo(&a, &b);
foo(&b, &c);
foo(&a, &b);
printf("a=%d b=%d c=%d\n", a, b, c);
```

A: a=3 b=2 c=1

B: a=1 b=2 c=3

Result is: C: a=1 b=3 c=2

D: a=3 b=3 c=3

E: a=1 b=1 c=1

# Arrays and Pointers

```
int
foo(int array[],
    unsigned int size)
{
    ...
    printf("%d\n", sizeof(array));
}

int
main(void)
{
    int a[10], b[5];
    int c[] = {1, 3, 2, 5, 6};
    ... foo(a, 10)... foo(c, 5) ...
    printf("%d\n", sizeof(c));
}
```

What does this print (64bit) **8**

... because **array** is really a pointer (and a pointer is architecture dependent, but likely to be 8 on modern machines!)

What does this print? **40**

## Quiz:

```
int x[] = { 2, 4, 6, 8, 10 };  
int *p = x;  
int **pp = &p;  
(*pp)++;  
(*(*pp))++;  
printf("%d\n", *p);
```

Result is:

A: 2

B: 3

C: 4

D: 5

E: None of the above

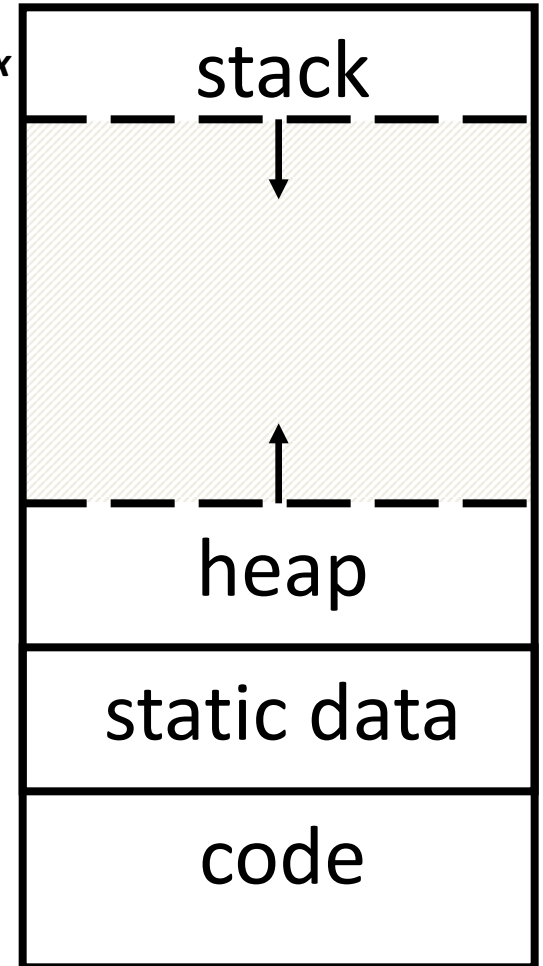
# C Memory Management

Memory Address  
(32 bits assumed here)

$\sim FFFF\ FFFF_{hex}$

- Program's *address space* contains 4 regions:
  - **stack**: local variables inside functions, grows downward
  - **heap**: space requested for dynamic data via `malloc()`; resizes dynamically, grows upward
  - **static data**: variables declared outside functions, does not grow or shrink. Loaded when program starts, can be modified.
  - **code**: loaded when program starts, does not change

$\sim 0000\ 0000_{hex}$

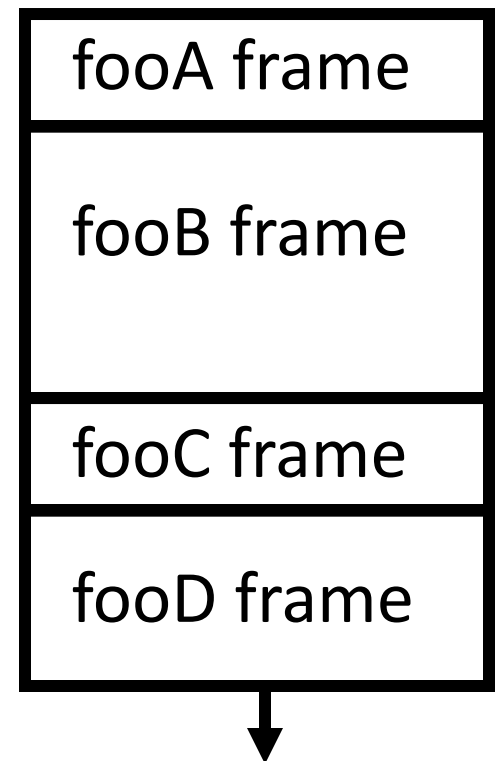


# The Stack

- Every time a function is called, a new frame is allocated on the stack
- Stack frame includes:
  - Return address (who called me?)
  - Arguments
  - Space for local variables
- Stack frames contiguous blocks of memory; stack pointer indicates start of stack frame
- When function ends, stack frame is tossed off the stack; frees memory for future stack frames
- We'll cover details later for MIPS processor

```
fooA() { fooB(); }  
fooB() { fooC(); }  
fooC() { fooD(); }
```

**Stack Pointer** →



# Faulty Heap Management

- What is wrong with this code?
- Memory leak!

```
int foo() {  
    int *value = malloc(sizeof(int));  
    *value = 42;  
    return *value;  
}
```



# Using Memory You Don't Own

- What is wrong with this code?

```
int* init_array(int *ptr, int new_size) {  
    ptr = realloc(ptr, new_size*sizeof(int));  
    memset(ptr, 0, new_size*sizeof(int));  
    return ptr;  
}
```

```
int* fill_fibonacci(int *fib, int size) {  
    int i;  
    init_array(fib, size);  
    /* fib[0] = 0; */ fib[1] = 1;  
    for (i=2; i<size; i++)  
        fib[i] = fib[i-1] + fib[i-2];  
    return fib;  
}
```

# Using Memory You Don't Own

- Improper matched usage of mem handles

```
int* init_array(int *ptr, int new_size) {  
    ptr = realloc(ptr, new_size*sizeof(int));  
    memset(ptr, 0, new_size*sizeof(int));  
    return ptr;  
}
```

Remember: `realloc` may move entire block

```
int* fill_fibonacci(int *fib, int size) {  
    int i;  
    /* oops, forgot: fib = */ init_array(fib, size);  
    /* fib[0] = 0; */ fib[1] = 1;  
    for (i=2; i<size; i++)  
        fib[i] = fib[i-1] + fib[i-2];  
    return fib;  
}
```

What if array is moved to new location?

# And In Conclusion, ...

- Pointers are an abstraction of machine memory addresses
- Pointer variables are held in memory, and pointer values are just numbers that can be manipulated by software
- In C, close relationship between array names and pointers
- Pointers know the type of the object they point to (except void \*)
- Pointers are powerful but potentially dangerous

# And In Conclusion, ...

- C has three main memory segments in which to allocate data:
  - Static Data: Variables outside functions
  - Stack: Variables local to function
  - Heap: Objects explicitly malloc-ed/free-d.
- Heap data is biggest source of bugs in C code

# MIPS

# Addition and Subtraction of Integers

## Example 1

- How to do the following C statement?  
 $a = b + c + d - e;$        $a = (b + c) + d - e;$   
 $b \rightarrow \$s1; c \rightarrow \$s2; d \rightarrow \$s3; e \rightarrow \$s4; a \rightarrow \$s0$
- Break into multiple instructions  

```
add $t0, $s1, $s2 # temp = b + c
add $t0, $t0, $s3 # temp = temp + d
sub $s0, $t0, $s4 # a = temp - e
```
- A single line of C may break up into several lines of MIPS.
- Notice the use of temporary registers – don't want to modify the variable registers \$s
- Everything after the hash mark on each line is ignored (comments)

# Overflow handling in MIPS

- Some languages detect overflow (Ada), some don't (most C implementations)
- MIPS solution is 2 kinds of arithmetic instructions:
  - These cause overflow to be detected
    - add (**add**)
    - add immediate (**addi**)
    - subtract (**sub**)
  - These do not cause overflow detection
    - add unsigned (**addu**)
    - add immediate unsigned (**addiu**)
    - subtract unsigned (**subu**)
- Compiler selects appropriate arithmetic
  - MIPS C compilers produce **addu, addiu, subu**

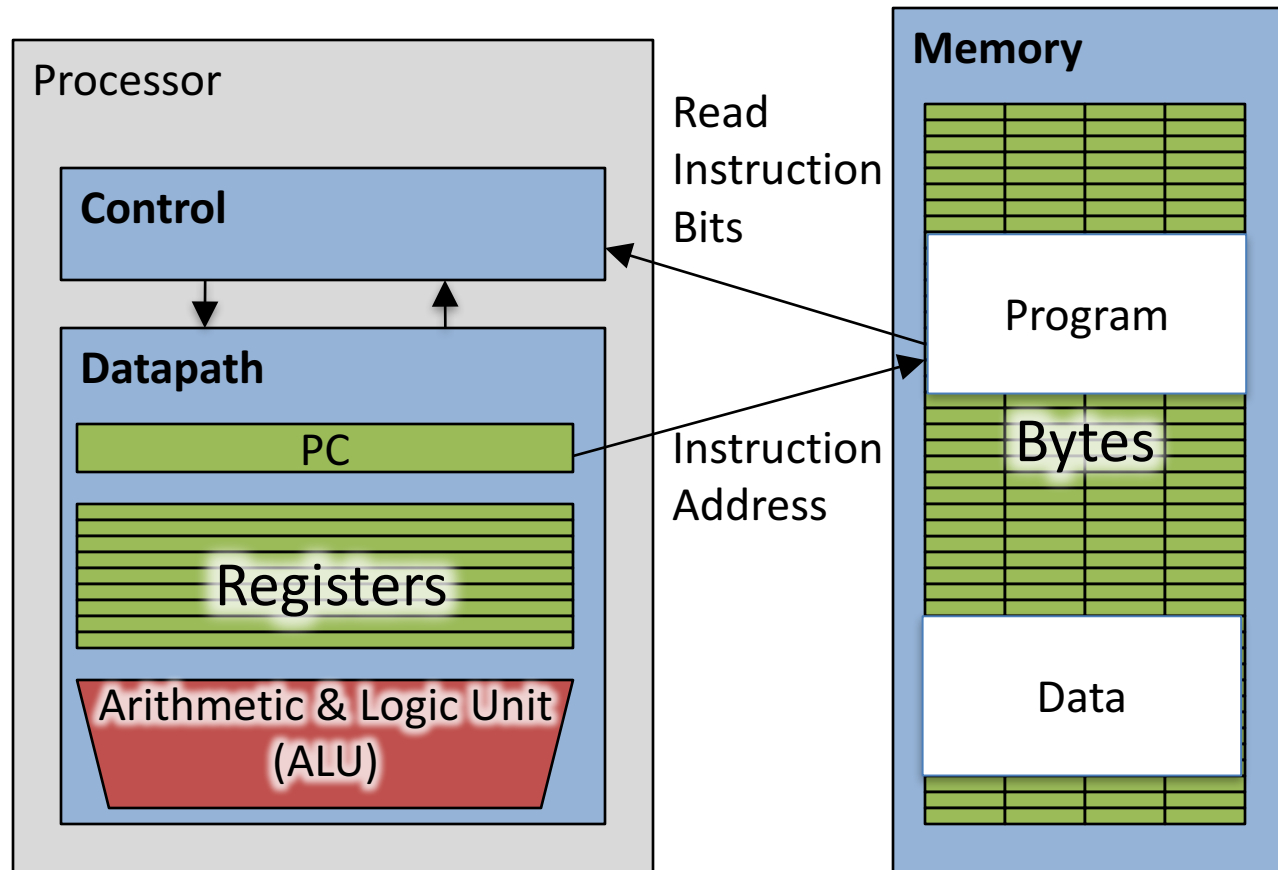
# Question:

We want to translate  $*x = *y + 1$  into MIPS  
( $x, y$  int pointers stored in:  $\$s0$   $\$s1$ )

- A:     addi    $\$s0, \$s1, 1$
- B:     lw        $\$s0, 1(\$s1)$   
       sw        $\$s1, 0(\$s0)$
- C:     lw        $\$t0, 0(\$s1)$   
       addi      $\$t0, \$t0, 1$   
       sw        $\$t0, 0(\$s0)$
- D:     sw        $\$t0, 0(\$s1)$   
       addi      $\$t0, \$t0, 1$   
       lw        $\$t0, 0(\$s0)$
- E:     lw        $\$s0, 1(\$t0)$   
       sw        $\$s1, 0(\$t0)$



# Executing a Program



- The PC (program counter) is internal register inside processor holding byte address of next instruction to be executed.
- Instruction is fetched from memory, then control unit executes instruction using datapath and memory system, and updates program counter (default is add +4 bytes to PC, to move to next sequential instruction)

# Question!

```
Start:  addi $s0,$zero,0
        slt  $t0,$s0,$s1
        beq  $t0,$zero,Exit
        sll  $t1,$s0,2
        addu $t1,$t1,$s5
        lw   $t1,0($t1)
        add  $s4,$s4,$t1
        addi $s0,$s0,1
        j    Start
Exit:
```

What is the code above?

- A: while loop
- B: do ... while loop
- C: for loop
- D: A or C
- E: Not a loop

# MIPS Function Call Conventions

- Registers faster than memory, so use them
- `$a0–$a3`: four *argument* registers to pass parameters (`$4 - $7`)
- `$v0, $v1`: two *value* registers to return values (`$2,$3`)
- `$ra`: one *return address* register to return to the point of origin (`$31`)

# Instruction Support for Functions (1/4)

```
... sum(a,b);... /* a,b:$s0,$s1 */  
}
```

**C**

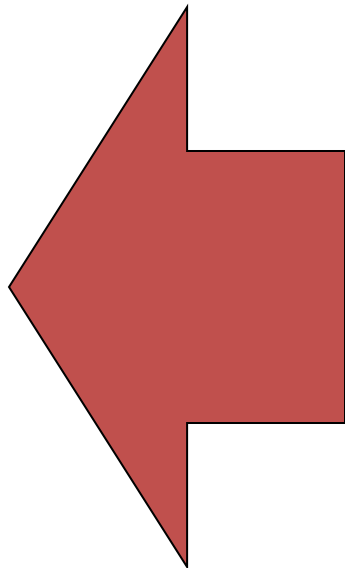
```
int sum(int x, int y) {  
    return x+y;  
}
```

---

address (shown in decimal)

**M**  
**I**  
**P**  
**S**

1000  
1004  
1008  
1012  
1016  
...  
2000  
2004



In MIPS, all instructions are 4 bytes, and stored in memory just like data. So here we show the addresses of where the programs are stored.

# Instruction Support for Functions (2/4)

```
... sum(a,b);... /* a,b:$s0,$s1 */  
}
```

C

```
int sum(int x, int y) {  
    return x+y;  
}
```

---

address (shown in decimal)

M  
I  
P  
S

```
1000 add    $a0,$s0,$zero    # x = a  
1004 add    $a1,$s1,$zero    # y = b  
1008 addi   $ra,$zero,1016    # $ra=1016  
1012 j      sum              # jump to sum  
1016 ...                    # next instruction  
...  
2000 sum:   add    $v0,$a0,$a1  
2004 jr     $ra    # new instr. "jump register"
```

# Instruction Support for Functions (3/4)

```
... sum(a,b);... /* a,b:$s0,$s1 */  
}
```


C

```
int sum(int x, int y) {  
    return x+y;  
}
```

---

- Question: Why use **jr** here? Why not use **j**?

- M  
I  
P  
S
- Answer: **sum** might be called by many places, so we can't return to a fixed place. The calling proc to **sum** must be able to say "return here" somehow.



```
2000 sum: add $v0,$a0,$a1  
2004 jr $ra # new instr. "jump register"
```

# Instruction Support for Functions (4/4)

- Single instruction to jump and save return address:  
jump and link (**j~~a~~l**)
- Before:  

```
1008 addi $ra,$zero,1016    # $ra=1016
1012 j  sum                # goto sum
```
- After:  

```
1008 jal sum    # $ra=1012, goto sum
```
- Why have a **j~~a~~l**?
  - Make the common case fast: function calls very common.
  - Don't have to know where code is in memory with **j~~a~~l**!

# Question

- Which statement is FALSE?
  - A: MIPS uses `jal` to invoke a function and `jr` to return from a function
  - B: `jal` saves `PC+1` in `$ra`
  - C: The callee can use temporary registers (`$ti`) without saving and restoring them
  - D: The caller can rely on save registers (`$si`) without fear of callee changing them



# Stack Before, During, After Call

High address

\$fp →

\$sp →

Low address

a.

\$fp →

Saved argument  
registers (if any)

Saved return address

Saved saved  
registers (if any)

Local arrays and  
structures (if any)

\$sp →

b.

\$fp →

\$sp →

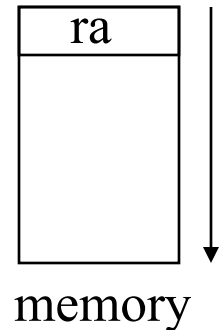
c.

# Basic Structure of a Function

## *Prologue*

```
entry_label:  
addi $sp,$sp, -framesize  
sw $ra, framesize-4($sp)  # save $ra  
save other regs if need be
```

*Body ... (call other functions...)*



## *Epilogue*

```
restore other regs if need be  
lw $ra, framesize-4($sp)  # restore $ra  
addi $sp,$sp, framesize  
jr $ra
```

# Instruction Formats

- **I-format**: used for instructions with immediates, **lw** and **sw** (since offset counts as an immediate), and branches (**beq** and **bne**)
  - (but not the shift instructions; later)
- **J-format**: used for **j** and **jal**
- **R-format**: used for all other instructions
- It will soon become clear why the instructions have been partitioned in this way

# R-Format Instructions (1/5)

- Define “**fields**” of the following number of bits each:  $6 + 5 + 5 + 5 + 5 + 6 = 32$

6	5	5	5	5	6
---	---	---	---	---	---

- For simplicity, each field has a name:

<b>opcode</b>	<b>rs</b>	<b>rt</b>	<b>rd</b>	<b>shamt</b>	<b>funct</b>
---------------	-----------	-----------	-----------	--------------	--------------

- Important:** On these slides and in book, each field is viewed as a 5- or 6-bit unsigned integer, not as part of a 32-bit integer
  - Consequence: 5-bit fields can represent any number 0-31, while 6-bit fields can represent any number 0-63

# I-Format Instructions (2/4)

- Define “fields” of the following number of bits each:  
 $6 + 5 + 5 + 16 = 32$  bits

6	5	5	16
---	---	---	----

- Again, each field has a name:

<b>opcode</b>	<b>rs</b>	<b>rt</b>	<b>immediate</b>
---------------	-----------	-----------	------------------

- **Key Concept:** Only one field is inconsistent with R-format. Most importantly, **opcode** is still in same location.

## I-Format Example (2/2)

- MIPS Instruction:

`addi $21, $22, -50`

**Decimal/field representation:**

8	22	21	-50
---	----	----	-----

**Binary/field representation:**

001000	10110	10101	1111111111001110
--------	-------	-------	------------------

**hexadecimal representation: 22D5 FFCE<sub>hex</sub>**

# Branch Example (1/2)

- MIPS Code:

```
Loop: beq    $9, $0, End
      addu   $8, $8, $10
      addiu  $9, $9, -1
      j      Loop
End:
```



- I-Format fields:

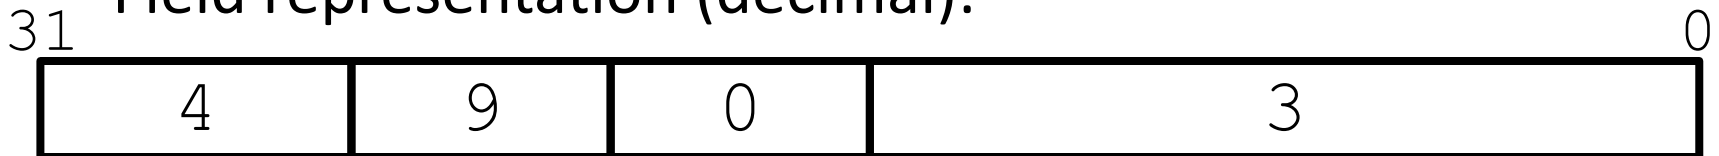
opcode = 4	(look up on Green Sheet)
rs = 9	(first operand)
rt = 0	(second operand)
immediate = 3	

# Branch Example (2/2)

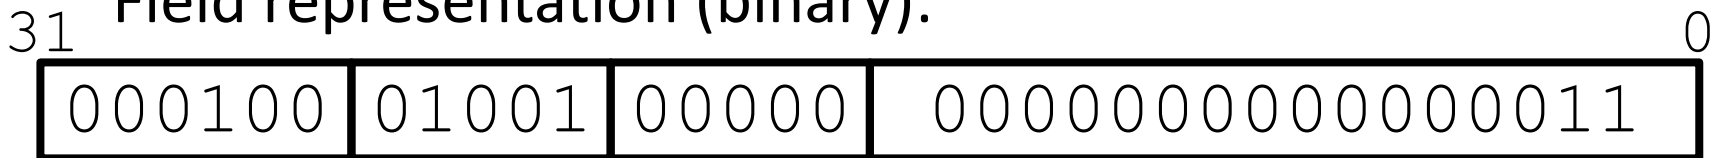
- MIPS Code:

```
Loop: beq    $9, $0, End
      addu   $8, $8, $10
      addiu  $9, $9, -1
      j      Loop
End:
```

Field representation (decimal):



Field representation (binary):





# J-Format Instructions (2/4)

- Define two “fields” of these bit widths:



- As usual, each field has a name:



- **Key Concepts:**

- Keep `opcode` field identical to R-Format and I-Format for consistency
- Collapse all other fields to make room for large target address

# Summary

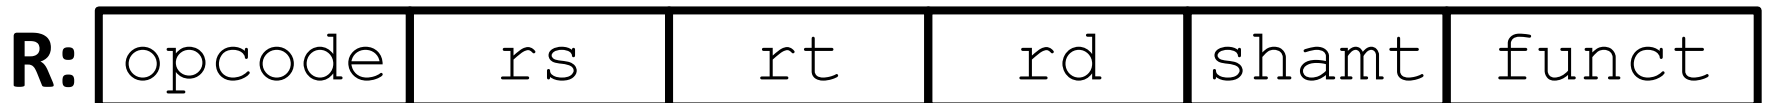
- **I-Format:** instructions with immediates, `lw/sw` (offset is immediate), and `beq/bne`
  - But not the shift instructions
  - Branches use PC-relative addressing



- **J-Format:** `j` and `jal` (but not `jr`)
  - Jumps use absolute addressing



- **R-Format:** all other instructions



# Assembler Pseudo-Instructions

- Certain C statements are implemented unintuitively in MIPS
  - e.g. assignment ( $a=b$ ) via `add $zero`
- MIPS has a set of “pseudo-instructions” to make programming easier
  - More intuitive to read, but get translated into actual instructions later

- Example:

```
move dst, src
```

translated into

```
addi dst, src, 0
```

# Multiply and Divide

- Example pseudo-instruction:

```
mul $rd,$rs,$rt
```

- Consists of **mult** which stores the output in special **hi** and **lo** registers, and a move from these registers to **\$rd**

```
mult $rs,$rt
```

```
mflo $rd
```

- **mult** and **div** have nothing important in the **rd** field since the destination registers are **hi** and **lo**
- **mfhi** and **mflo** have nothing important in the **rs** and **rt** fields since the source is determined by the instruction (see COD)

# Question

Which of the following place the address of LOOP in \$v0?

1) `la $t1, LOOP`  
`lw $v0, 0($t1)`

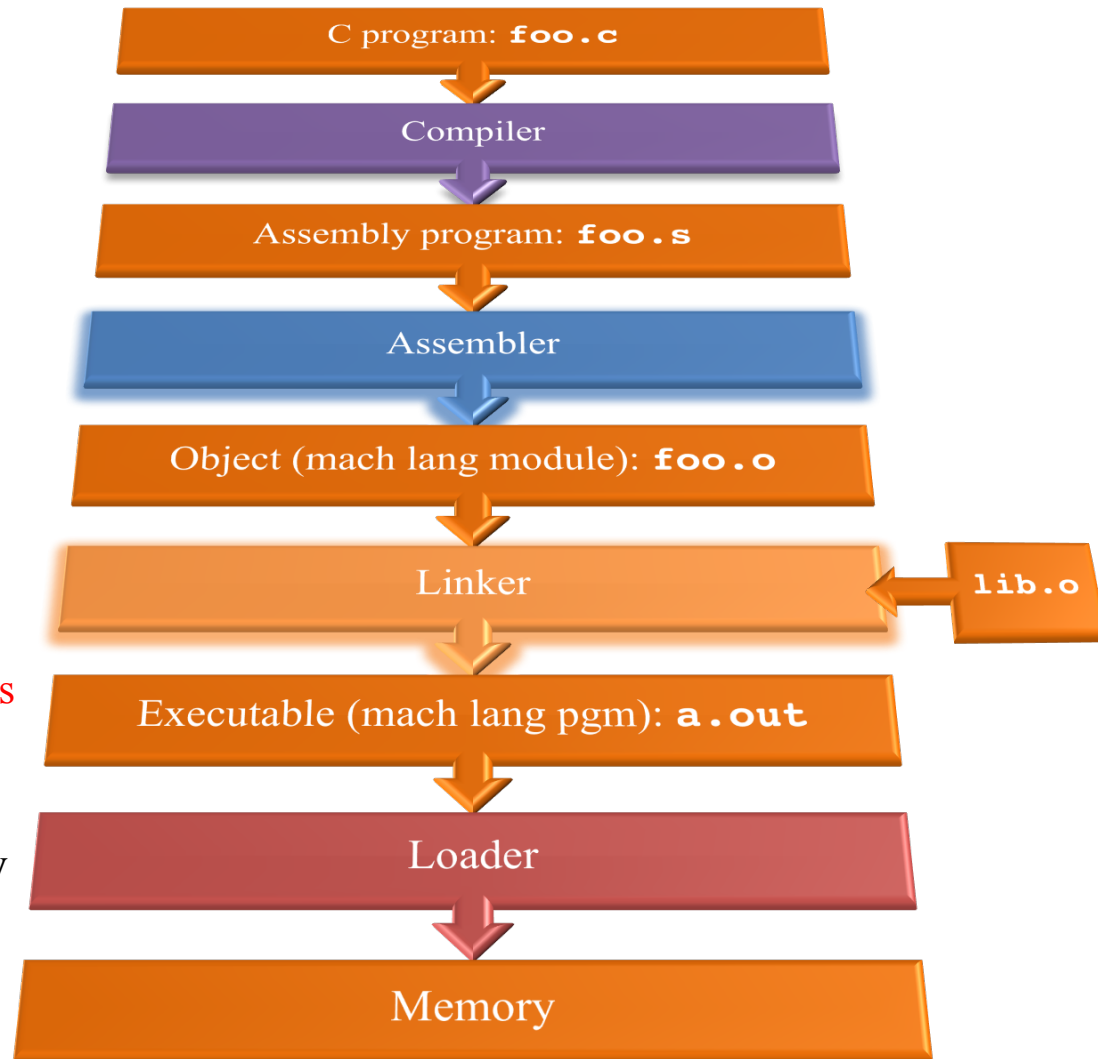
2) `jal LOOP`  
`LOOP: addu $v0, $ra, $zero`

3) `la $v0, LOOP`

	1	2	3
A)	T	T	T
B)	T	T	F
C)	F	T	T
D)	F	T	F
E)	F	F	T

# Steps in compiling a C program

- Compiler converts a single HLL file into a single assembly language file.
- Assembler removes pseudo-instructions, converts what it can to machine language, and creates a checklist for the linker (relocation table). A `.s` file becomes a `.o` file.
  - Does 2 passes to resolve addresses, handling internal forward references
- Linker combines several `.o` files and resolves absolute addresses.
  - Enables separate compilation, libraries that need not be compiled, and resolves remaining addresses
- Loader loads executable into memory and begins execution.



# Pseudo-instruction Replacement

- Assembler treats convenient variations of machine language instructions as if real instructions

Pseudo:

```
subu $sp,$sp,32
```

```
sd $a0, 32($sp)
```

```
mul $t7,$t6,$t5
```

```
addu $t0,$t6,1
```

```
ble $t0,100,loop
```

```
la $a0, str
```

Real:

```
addiu $sp,$sp,-32
```

```
sw $a0, 32($sp)
```

```
sw $a1, 36($sp)
```

```
mult $t6,$t5
```

```
mflo $t7
```

```
addiu $t0,$t6,1
```

```
slti $at,$t0,101
```

```
bne $at,$0,loop
```

```
lui $at,left(str)
```

```
ori $a0,$at,right(str)
```

# Question

At what point in process are all the machine code bits generated for the following assembly instructions:

1) `addu $6, $7, $8`

2) `jal fprintf`

A: 1) & 2) After compilation

B: 1) After compilation, 2) After assembly

C: 1) After assembly, 2) After linking

D: 1) After assembly, 2) After loading

E: 1) After compilation, 2) After linking



# INTRO TO CACHES

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

*Software*

*Hardware*

Warehouse  
Scale  
Computer

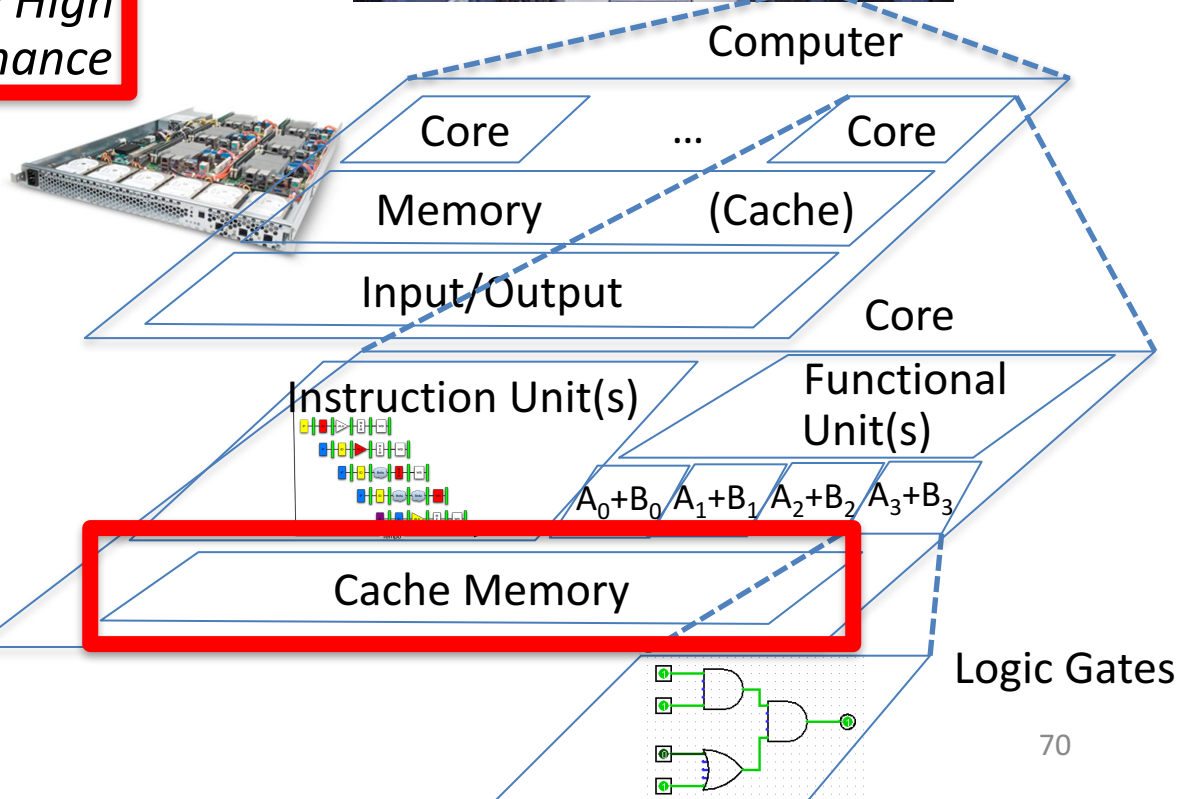


Smart  
Phone



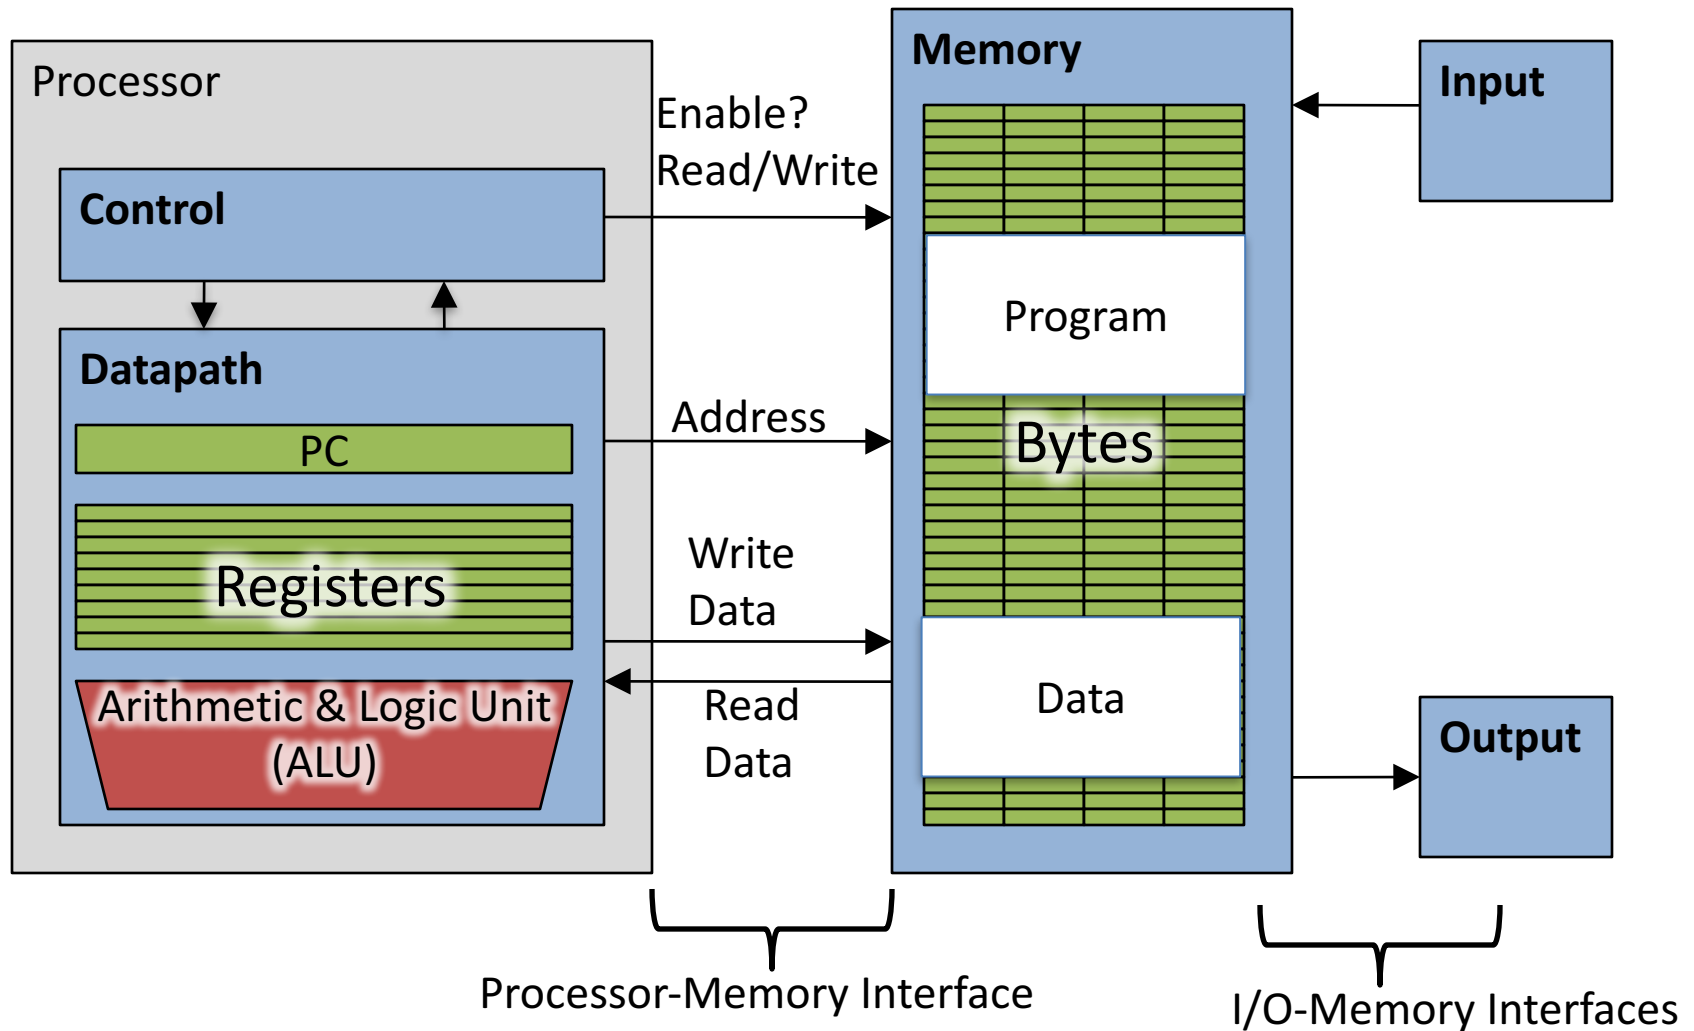
*Harness  
Parallelism &  
Achieve High  
Performance*

How do  
we know?



- 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 @ one time
- Programming Languages

# Components of a Computer



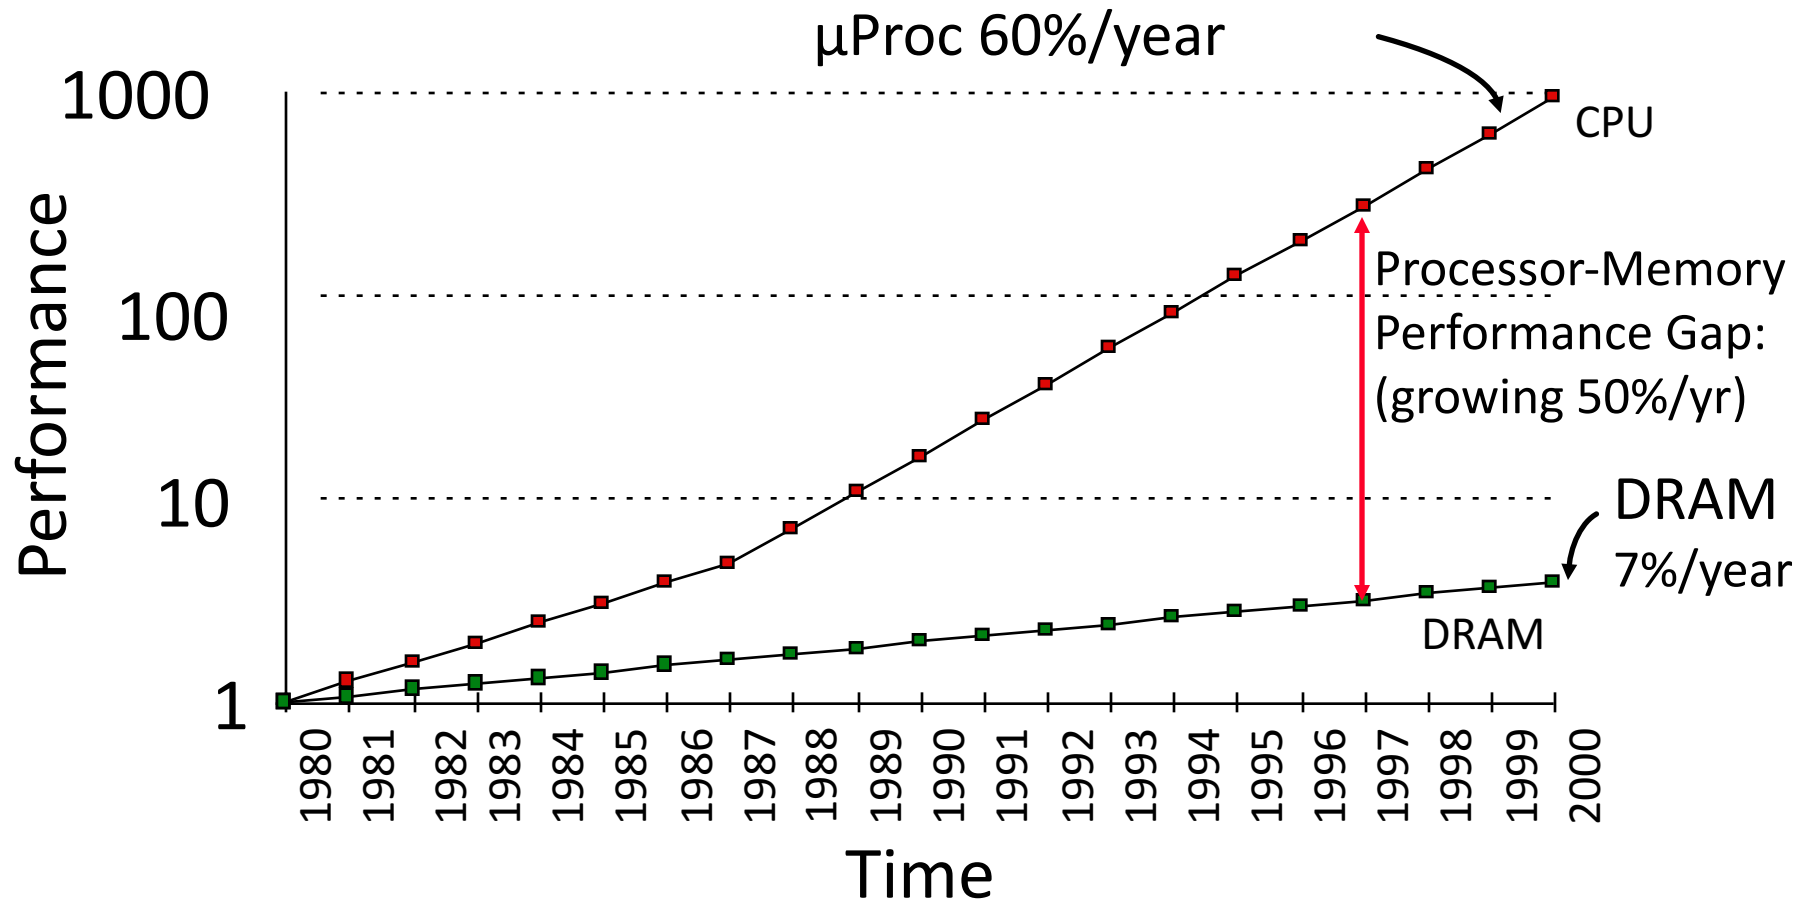
# Problem: Large memories slow?

## Library Analogy

- Finding a book in a large library takes time
  - Takes time to search a large card catalog – (mapping title/author to index number)
  - Round-trip time to walk to the stacks and retrieve the desired book.
- Larger libraries makes both delays worse
- Electronic memories have the same issue, *plus* the technologies that we use to store an individual bit get slower as we increase density (SRAM versus DRAM versus Magnetic Disk)

***However what we want is a large yet fast memory!***

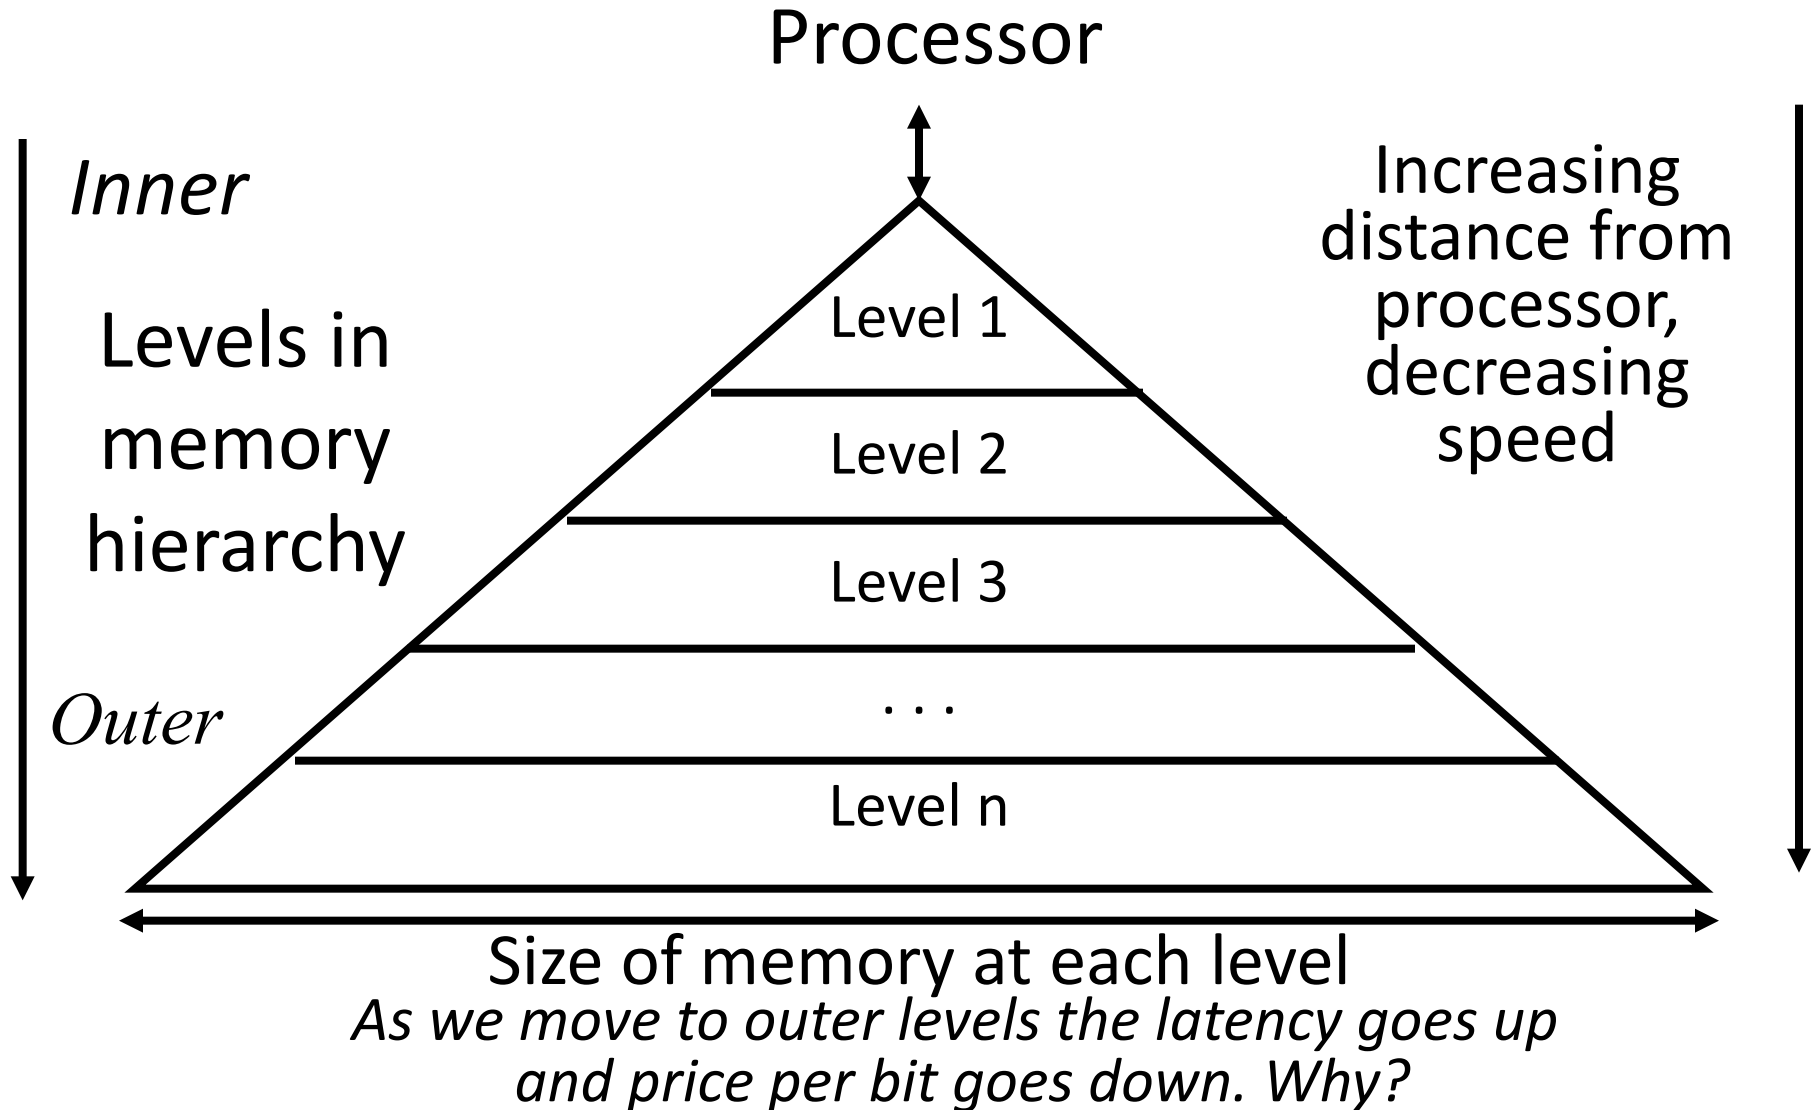
# Processor-DRAM Gap (latency)



1980 microprocessor executes ~one instruction in same time as DRAM access  
2015 microprocessor executes ~1000 instructions in same time as DRAM access

***Slow DRAM access could have disastrous impact on CPU performance!***

# Big Idea: Memory Hierarchy

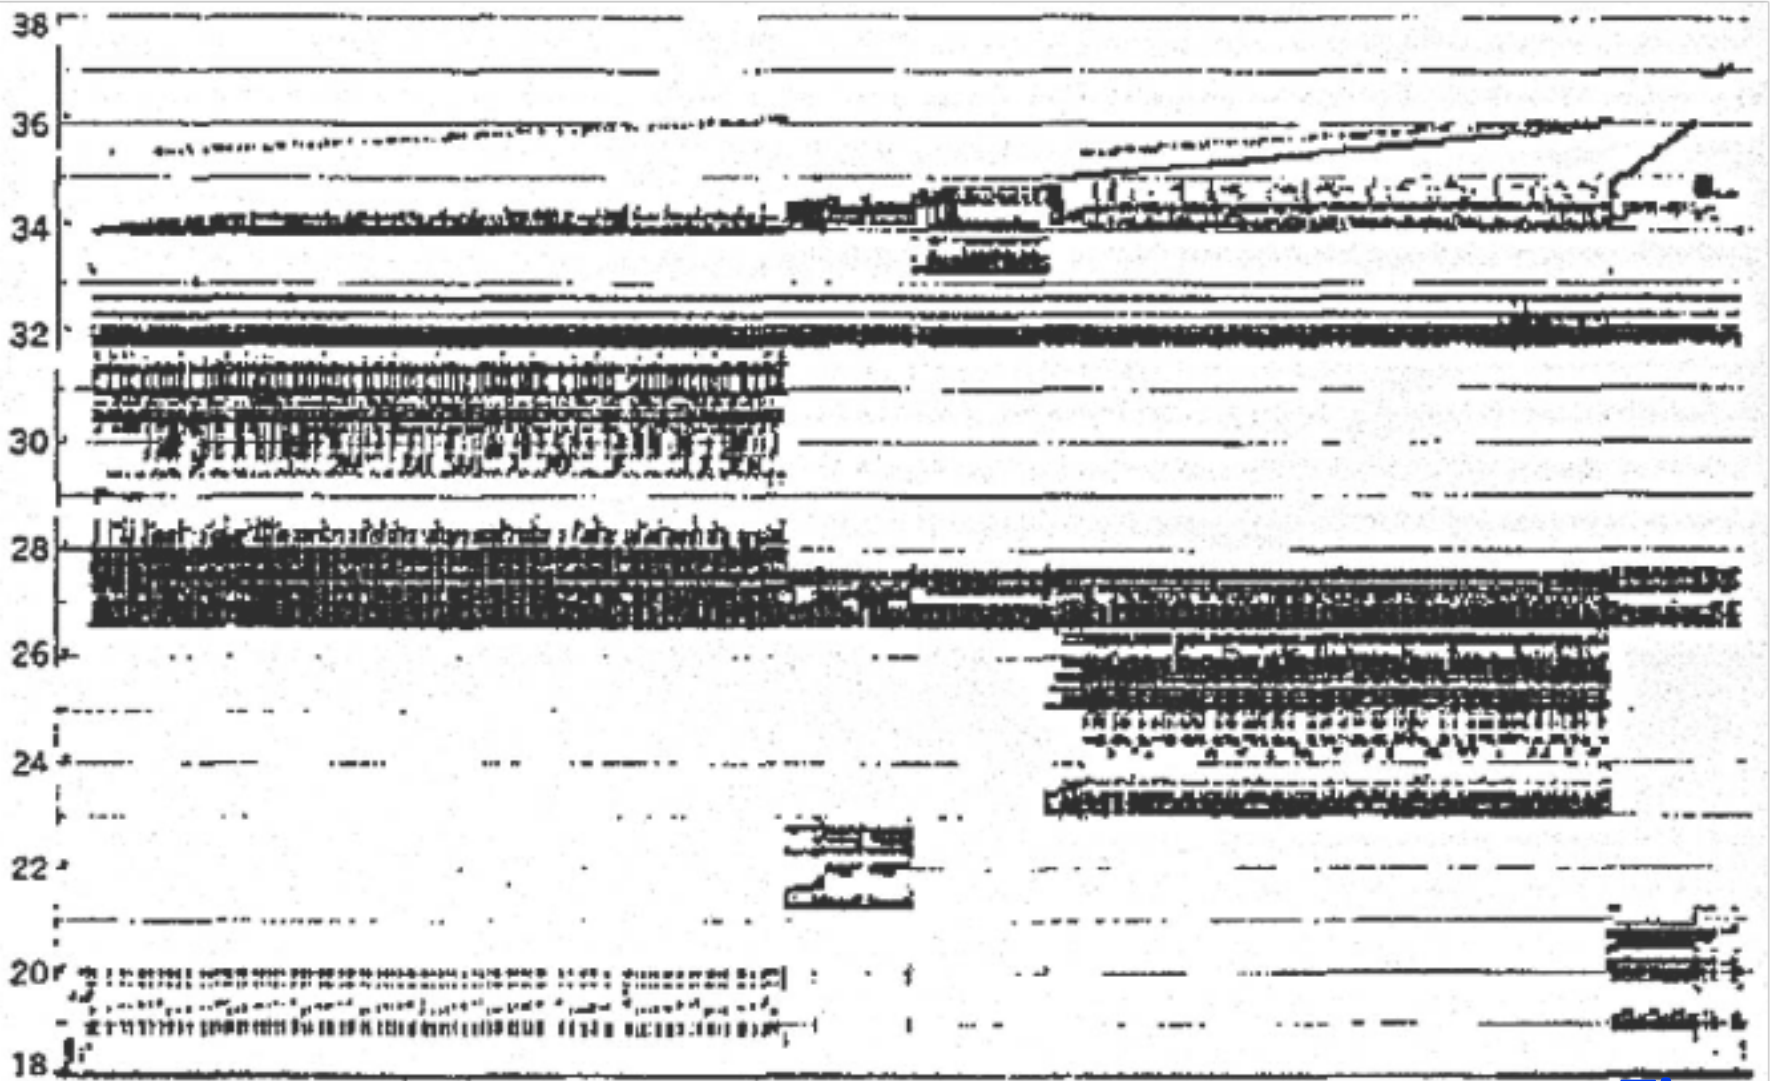


# What to do: Library Analogy

- Want to write a report using library books
- Go to library, look up relevant books, fetch from stacks, and place on desk in library
- If need more, check them out and keep on desk
  - But don't return earlier books since might need them
- You hope this collection of ~10 books on desk enough to write report, despite 10 being only a tiny fraction of books available

# Real Memory Reference Patterns

Memory Address (one dot per access)



Time

Donald J. Hatfield, Jeanette Gerald: Program Restructuring for Virtual Memory. IBM Systems Journal 10(3): 168-192 (1971)

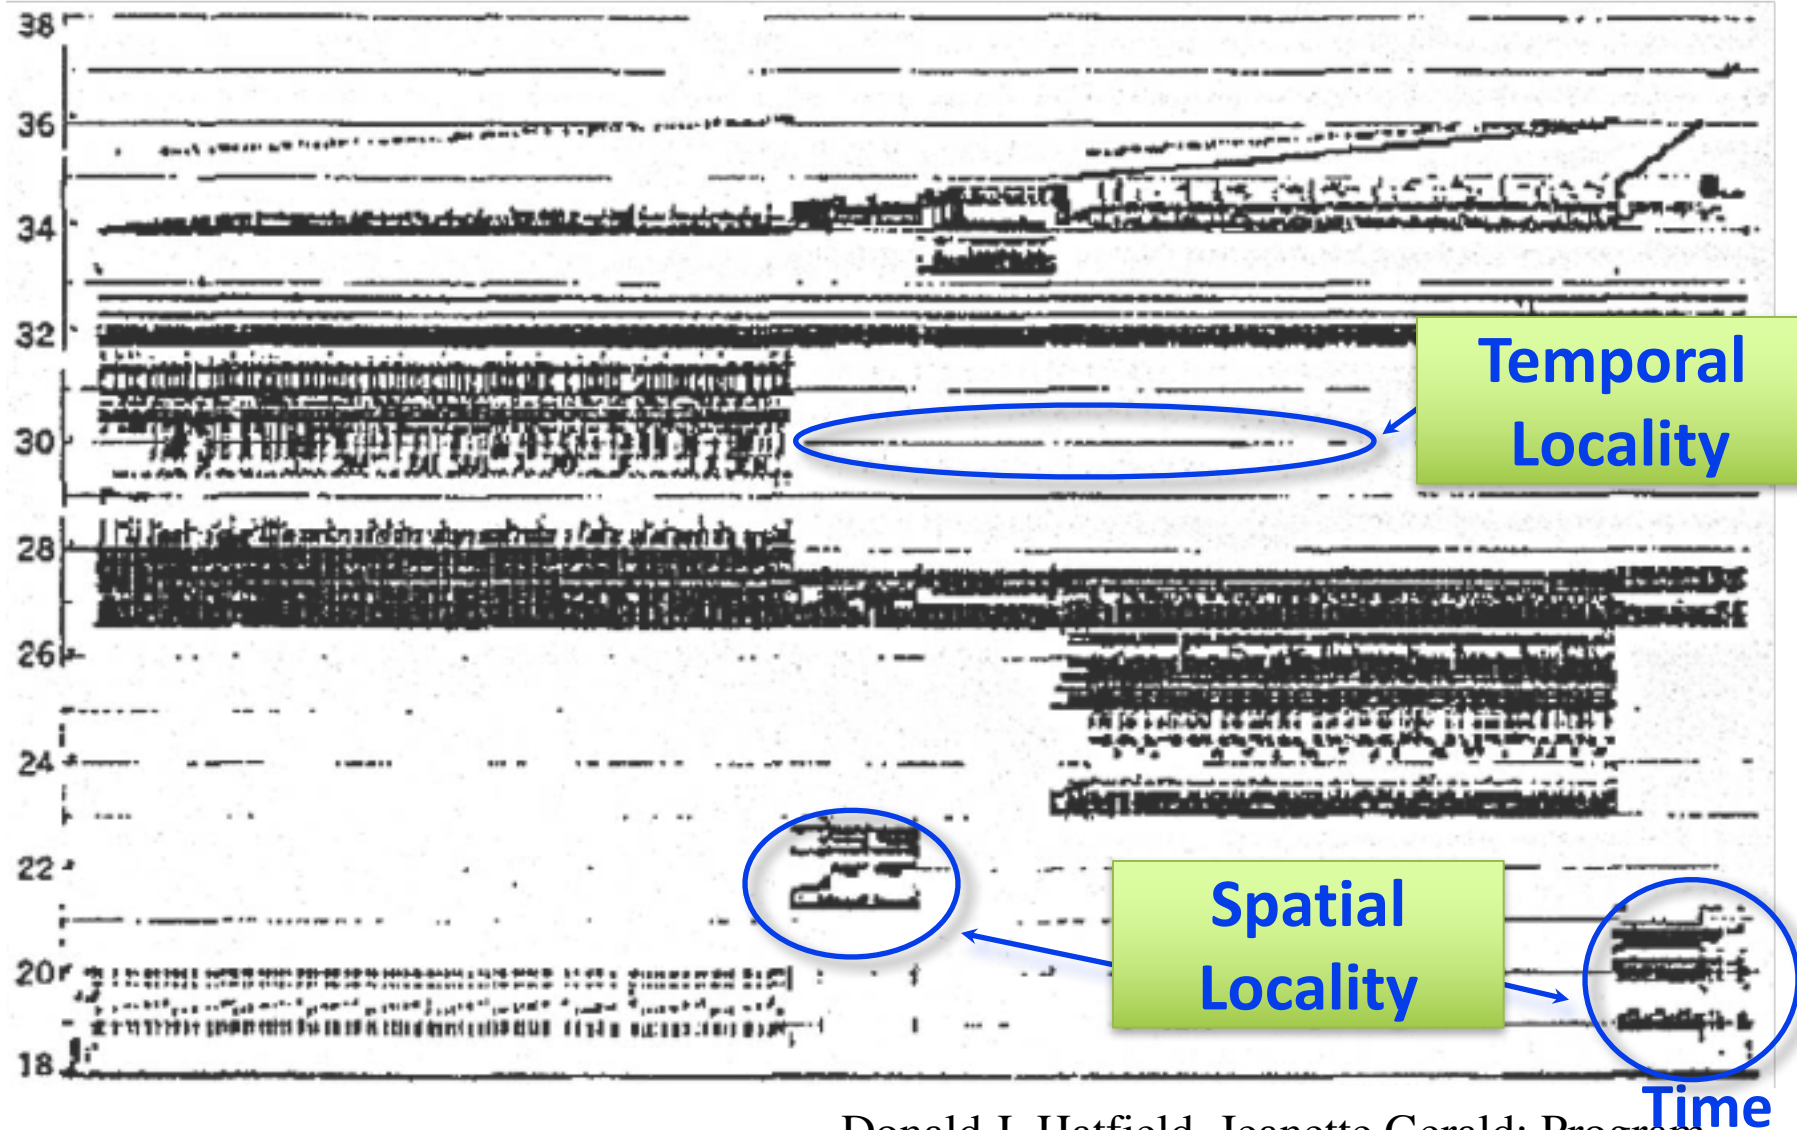


# Big Idea: Locality

- *Temporal Locality* (locality in time)
  - Go back to same book on desktop multiple times
  - If a memory location is referenced, then it will tend to be referenced again soon
- *Spatial Locality* (locality in space)
  - When go to book shelf, pick up multiple books on J.D. Salinger since library stores related books together
  - If a memory location is referenced, the locations with nearby addresses will tend to be referenced soon

# Memory Reference Patterns

Memory Address (one dot per access)

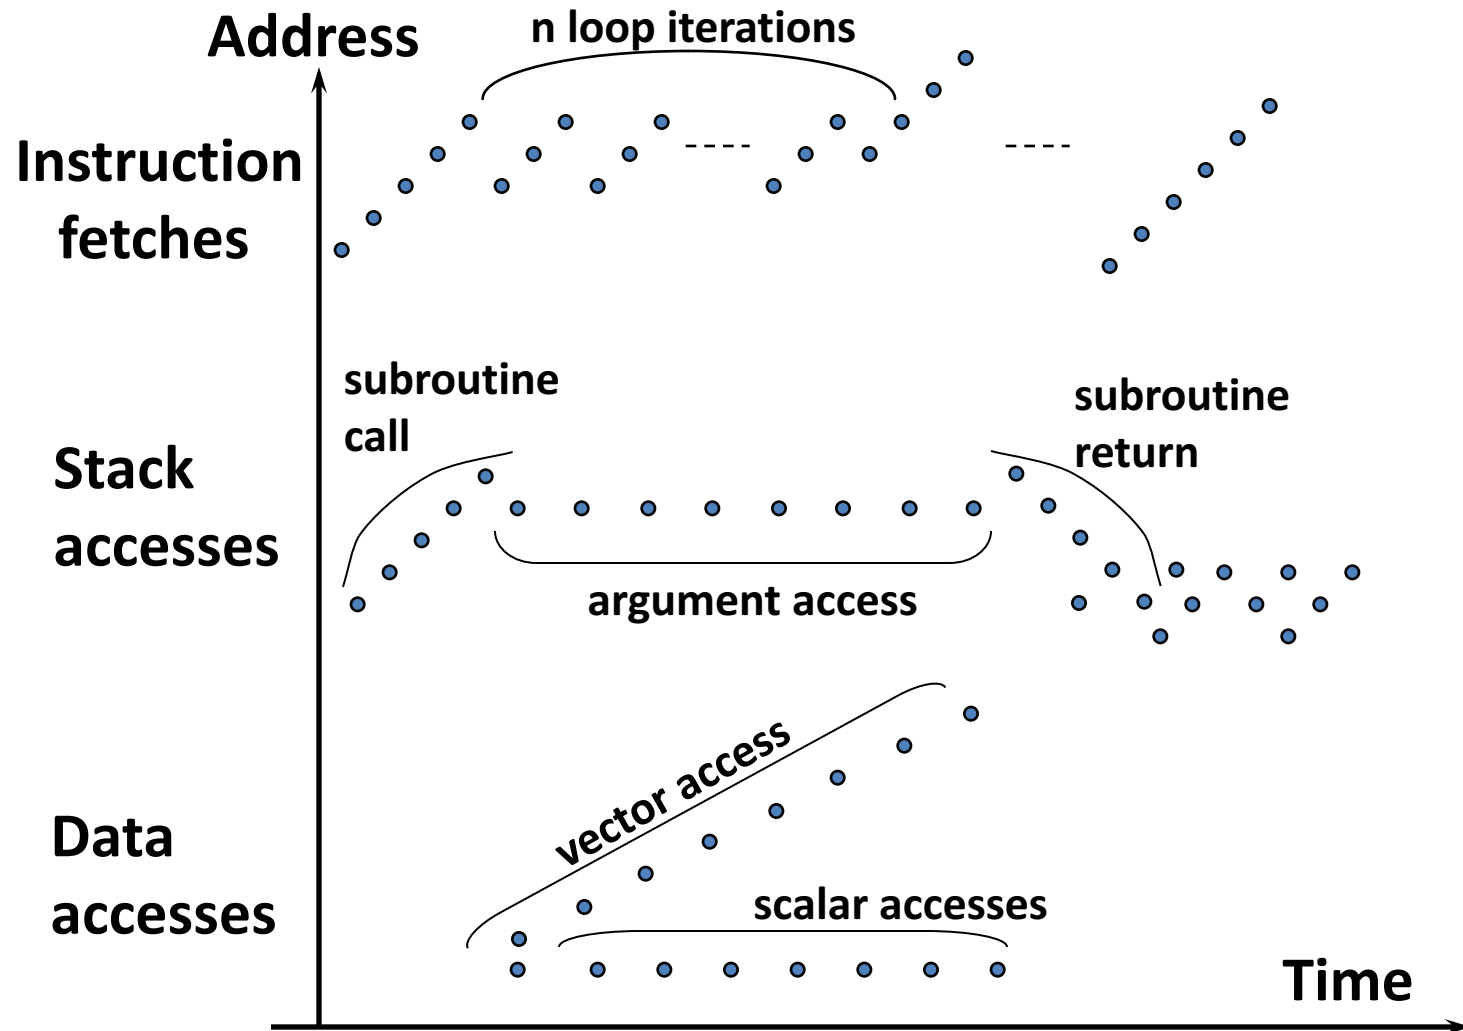


Donald J. Hatfield, Jeanette Gerald: Program Restructuring for Virtual Memory. IBM Systems Journal 10(3): 168-192 (1971)

# Principle of Locality

- *Principle of Locality*: Programs access small portion of address space at any instant of time (spatial locality) and repeatedly access that portion (temporal locality)
- What program structures lead to **temporal** and **spatial locality** in **instruction** accesses?
- In **data** accesses?

# Memory Reference Patterns



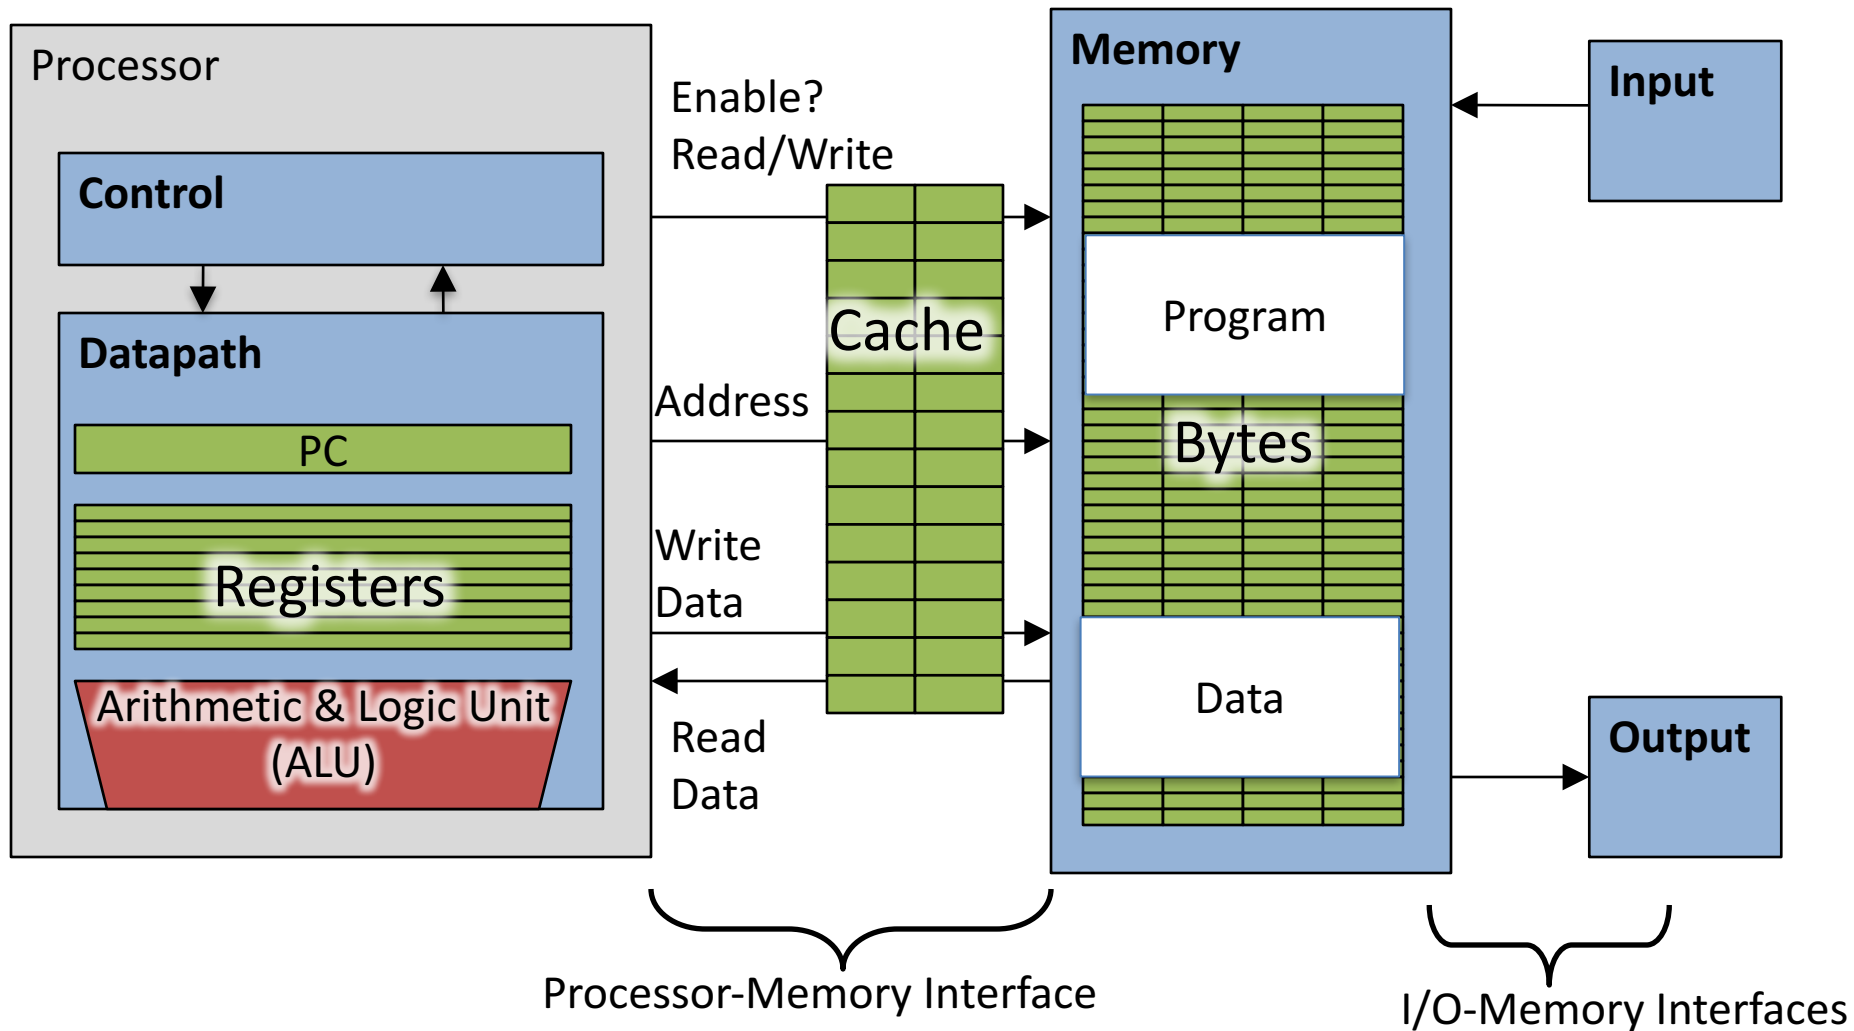
# Cache Philosophy

- Programmer-invisible hardware mechanism to give illusion of speed of fastest memory with size of largest memory
  - Works fine even if programmer has no idea what a cache is
  - However, performance-oriented programmers today sometimes “reverse engineer” cache design to design data structures to match cache

# Memory Access without Cache

- Load word instruction: `lw $t0, 0($t1)`
- `$t1` contains  $1022_{\text{ten}}$ , `Memory[1022] = 99`
  1. Processor issues address  $1022_{\text{ten}}$  to Memory
  2. Memory reads word at address  $1022_{\text{ten}}$  (99)
  3. Memory sends 99 to Processor
  4. Processor loads 99 into register `$t0`

# Adding Cache to Computer



# Memory Access with Cache

- Load word instruction: `lw $t0, 0($t1)`
- `$t1` contains  $1022_{ten}$ , `Memory[1022] = 99`
- With cache: Processor issues address  $1022_{ten}$  to Cache
  1. Cache checks to see if has copy of data at address  $1022_{ten}$ 
    - 2a. If finds a match (Hit): cache reads 99, sends to processor
    - 2b. No match (Miss): cache sends address 1022 to Memory
      - I. Memory reads 99 at address  $1022_{ten}$
      - II. Memory sends 99 to Cache
      - III. Cache replaces word with new 99
      - IV. Cache sends 99 to processor
  2. Processor loads 99 into register `$t0`



# Cache “Tags”

- Need way to tell if have copy of location in memory so that can decide on hit or miss
- On cache miss, put memory address of block in “tag address” of cache block

1022 placed in tag next to data from memory (99)

Tag	Data
252	12
1022	99
131	7
2041	20

From earlier instructions

# Anatomy of a 16 Byte Cache, 4 Byte Block

- Operations:
  1. Cache Hit
  2. Cache Miss
  3. Refill cache from memory
- Cache needs Address Tags to decide if Processor Address is a Cache Hit or Cache Miss
  - Compares all 4 tags

