



CS4803DGC Design Game Console
Spring 2009
Prof. Hyesoon Kim

Georgia Tech College of Computing

The slide features a grid of images including a building, a blue square, a green circuit board, a purple circuit board, a green square, and a brown square.

Sources

- LC3B1, LC3B2, LC3B3

Georgia Tech College of Computing

The slide has a title bar with four colored icons: orange, blue, purple, and green.

What is Architecture?

Problem
Algorithm
ISA
u-architecture
Circuits
Electrons

ISA: Interface between s/w & h/w

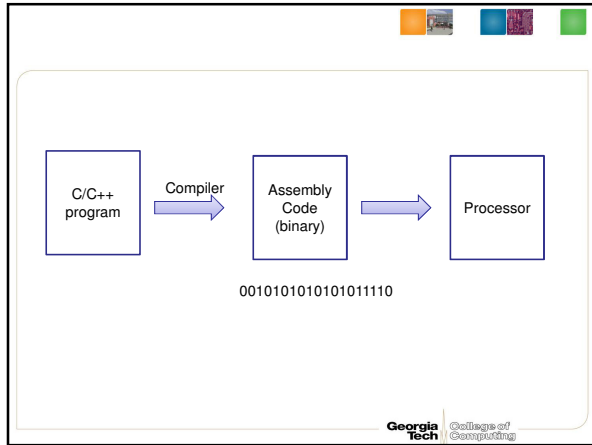
Georgia Tech College of Computing

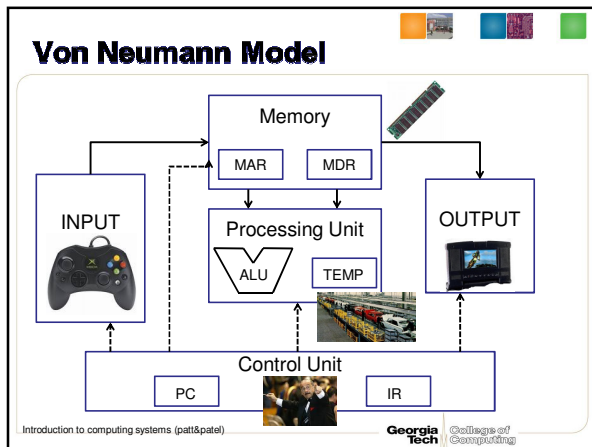
The slide has a title bar with four colored icons: orange, blue, purple, and green.

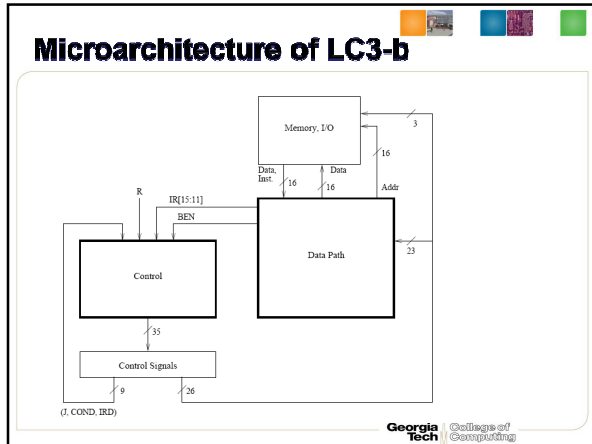
ISA (Instruction Set Architecture)

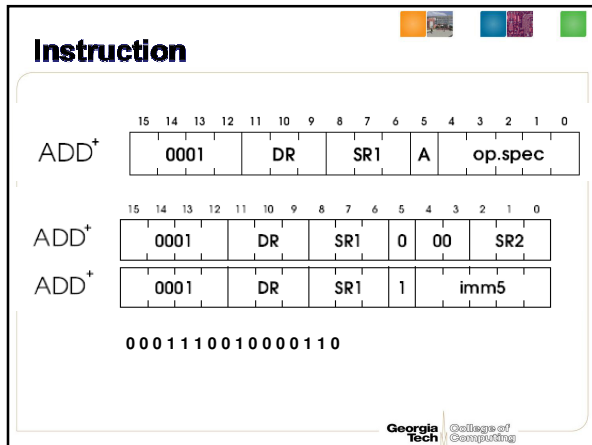
- Opcode
- Data type
- Addressing mode: register, immediate, displacement
 - E.g.) x86 SIB
- Memory addressing: unaligned, aligned

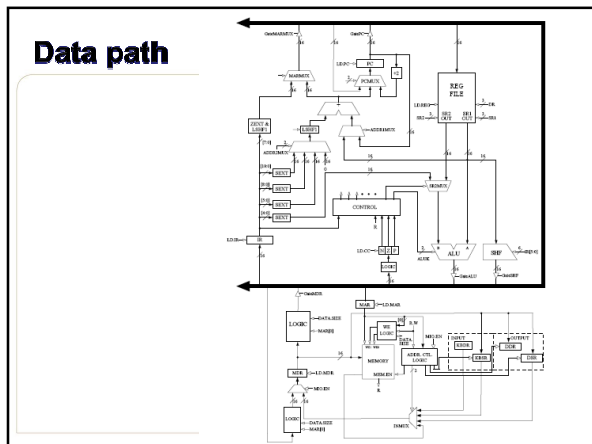
Georgia Tech College of Computing

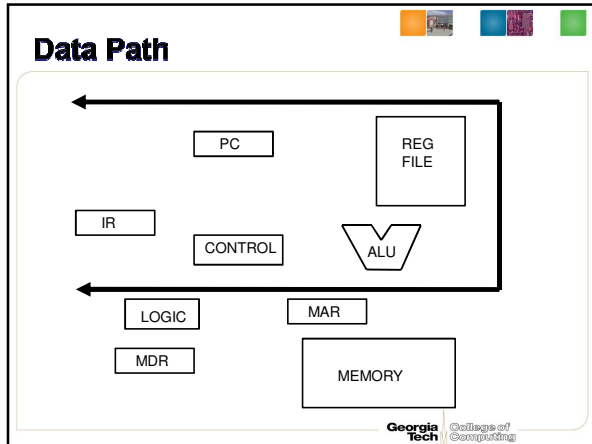


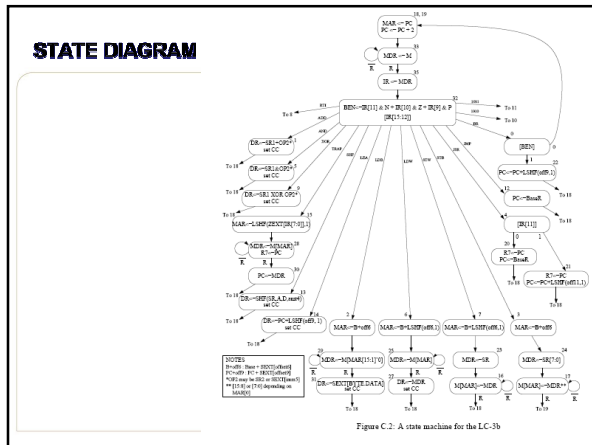


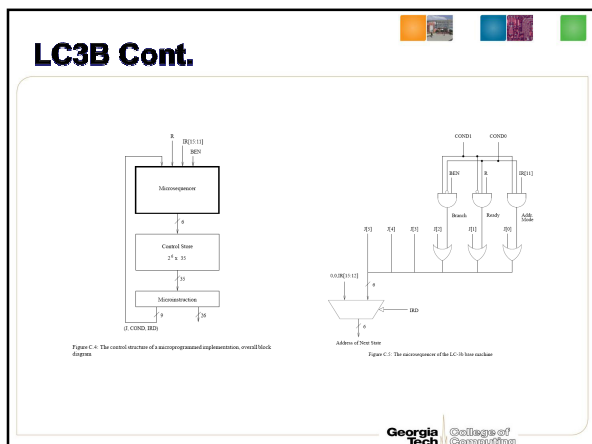












Homework #1

- Please read LC3b1, LC3b2, LC3b3 first.
- Be careful with Conditional Code

Georgia Tech College of Computing

PIPELINE DESIGN

Georgia Tech College of Computing

Overview of a Processor

0x0001	LD R1, MEM[R0]
0x0002	ADD R2, R2, #1
0x0003	BRZERO 0x0001

Non-pipelined: 15 cycles

Pipelined: 7 cycles

Georgia Tech College of Computing

Simple 5-stage Pipeline

- IF: Instruction fetch cycle
- ID: Instruction decode/register fetch cycle
- EX: Execution/effective address cycle
- MEM: Memory access cycle
- WB: Write-back cycle

Georgia Tech College of Computing

Basic Pipeline

Georgia Tech College of Computing

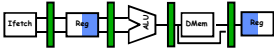
ILP is Bounded

- For any sequence of instructions, the available parallelism is limited
- Hazards/Dependencies are what limit the ILP
 - Data dependencies
 - Control dependencies
 - Memory dependencies

Georgia Tech College of Computing

Dependences/Dependencies

- Data Dependencies
 - RAW: Read-After-Write (True Dependence)
 - WAR: Anti-Dependence
 - WAW: Output Dependence
- Control Dependence
 - When following instructions depend on the outcome of a previous branch/jump



Georgia Tech College of Computing

Data Dependencies

- Register dependencies
 - RAW, WAR, WAW, based on register *number*
- Memory dependencies
 - Based on memory *address*
 - This is harder
 - Register names known at decode
 - Memory addresses not known until execute
 - Also called memory disambiguation

Georgia Tech College of Computing

Types of Data Dependencies

(Assume A comes before B in program order)

- RAW (Read-After-Write)
 - A writes to a location, B reads from the location, therefore B has a RAW dependency on A
 - Also called a “true dependency”

Georgia Tech College of Computing

Data Dep's (cont'd)

- WAR (Write-After-Read)
 - A reads from a location, B writes to the location, therefore B has a WAR dependency on A
 - If B executes before A has read its operand, then the operand will be lost
 - Also called an anti-dependence

Georgia Tech College of Computing

Data Dep's (cont'd)

- Write-After-Write
 - A writes to a location, B writes to the same location
 - If B writes first, then A writes, the location will end up with the wrong value
 - Also called an output-dependence

Georgia Tech College of Computing

Control Dependencies

- If we have a conditional branch, until we actually know the outcome, *all* later instructions must wait
 - That is, all instructions are control dependent on all earlier branches
 - This is true for unconditional branches as well (e.g., can't return from a function until we've loaded the return address)

Georgia Tech College of Computing

Memory Dependencies

- Basically similar to regular (register) data dependencies: RAW, WAR, WAW
- However, the exact location is not known until run-time :
 - A: STORE R1, 0[R2]
 - B: LOAD R5, 24[R8]
 - C: STORE R3, -8[R9]
- RAW exists if $(R2+0) == (R8+24)$
- WAR exists if $(R8+24) == (R9 - 8)$
- WAW exists if $(R2+0) == (R9 - 8)$

Georgia Tech | College of Computing

Identifying Dependences

I1: R2 = 17
I2: R1 = 49
I3: R3 = -8
I4: R5 = LOAD 0[R3]
I5: R4 = R1 + R2
I6: R7 = R4 - R3
I7: R6 = R4 * R5
I8: R5 = R6 + 10

List RAW, WAR, WAW

Georgia Tech | College of Computing

Hazards

- When two instructions that have one or more dependences between them occur close enough that changing the instruction order will change the outcome of the program
- Not all dependencies lead to hazards!

Georgia Tech | College of Computing

Control Hazard

Control hazard

```

In s t r.
O r d e r
add r1, r2, r3
uncond. TARG
and r6, r1, r7
TARG
or r1, r2, r9
br (r1) TARG1
    
```

Control hazard and data hazard

Georgia Tech College of Computing

Data Forwarding

Time (clock cycles)

```

In s t r.
O r d e r
add r1, r2, r3
sub r4, r1, r3
and r6, r1, r7
or r8, r1, r9
xor r10, r1, r4
    
```

- Destination register is a name for instr's result
- Source registers are names for sources
- Forwarding logic builds data dependence (flow) graph for instructions in the execution window

Georgia Tech College of Computing

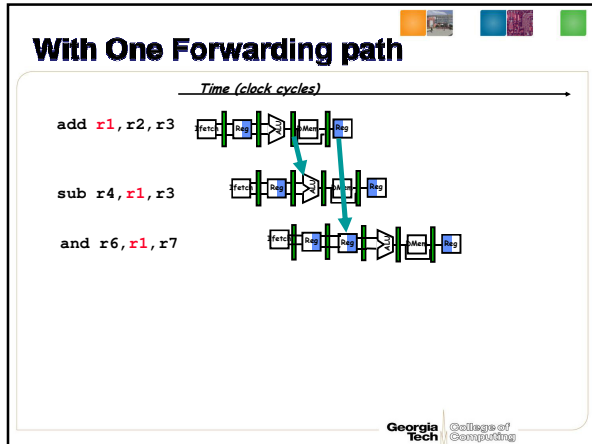
W/O Forwarding

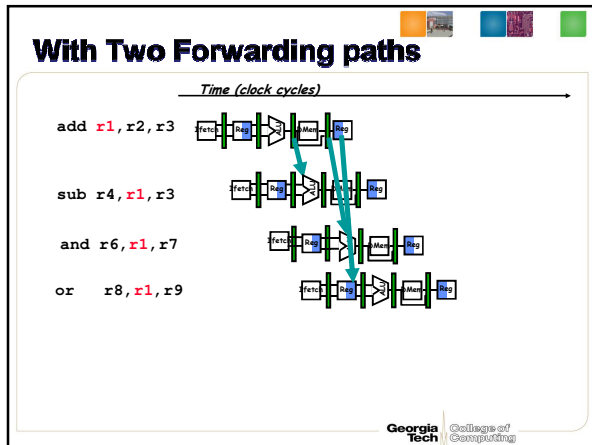
Time (clock cycles)

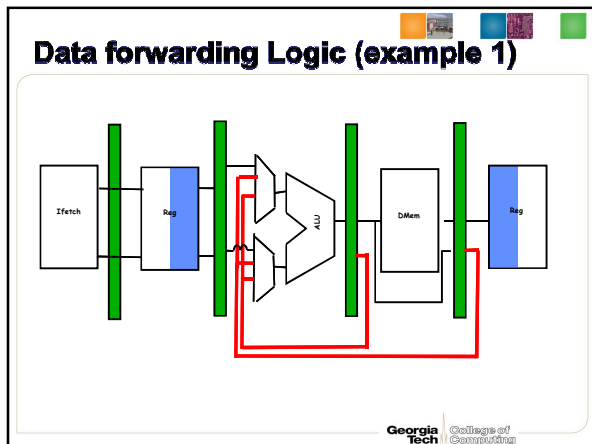
```

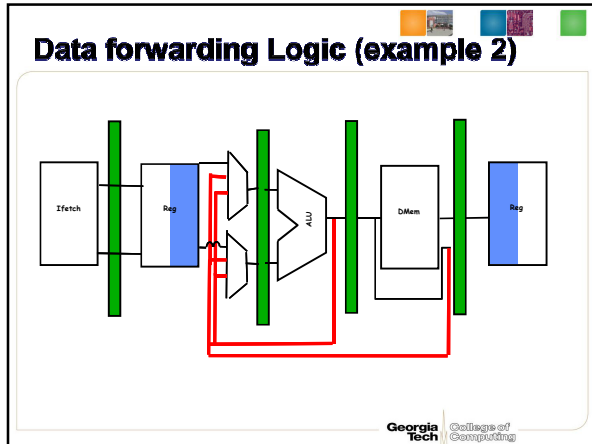
add r1, r2, r3
sub r4, r1, r3
and r6, r1, r7
    
```

Georgia Tech College of Computing

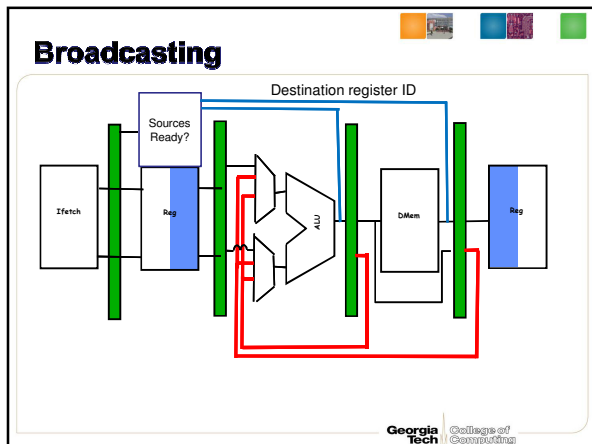


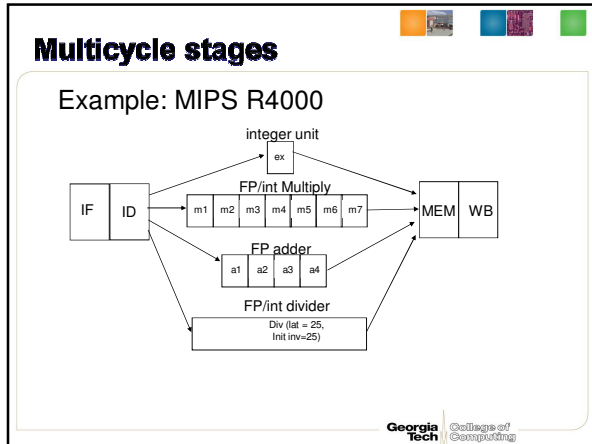


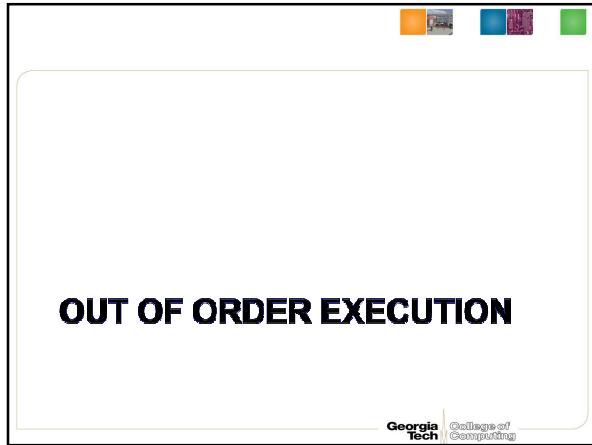


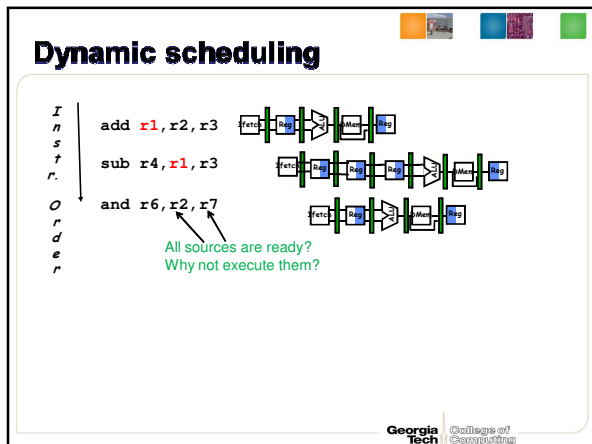


- Is that all?
 - How Do we know whether Add is writing a data to R1?
 - Broadcasting Register id
- Georgia Tech College of Computing









Example

```

mul r1, r2, r3
add r4, r1, r3
add r6, r2, r5
mul r6, r3, r4
    
```

FU	Busy	SRC1		SRC2		DR			
		ID	R	FUNC	ID	R	FUNC	ID	R
ADD									
ADD									
MUL									
MUL									

Instruction	Issue	Read operation	Execution complete	Write results

Georgia Tech College of Computing

Example

Cycle = 3

```

mul r1, r2, r3
add r4, r1, r3
add r6, r2, r5
mul r6, r3, r4
    
```

Execution latency: mul 4 cycle
add 1 cycle

FU	Busy	SRC1		SRC2		DR			
		ID	R	FUNC	ID	R	FUNC	ID	R
ADD1		r1		mul1	r3	V		r4	
ADD2									
MUL1	V	r2	V		r3	V		r1	
MUL2									

Instruction	Issue	Read operation	Execution complete	Write results
Mul1 (I1)	V	V		
add (I2)	V			

Georgia Tech College of Computing

Example

Cycle = 4

```

mul r1, r2, r3
add r4, r1, r3
add r6, r2, r5
mul r6, r3, r4
    
```

FU	Busy	SRC1		SRC2		DR			
		ID	R	FUNC	ID	R	FUNC	ID	R
ADD1		r1		mul1	r3	V		r4	
ADD2	V	r2	V		r5	V		r6	
MUL1	V	r2	V		r3	V		r1	
MUL2									

Instruction	Issue	Read operation	Execution complete	Write results
Mul1 (I1)	V	V	Progress (3)	
add (I2)	V			
Add (I3)	V			

Georgia Tech College of Computing

Handling Branches: I

br target 0x800
add r1, r2, r3 0x804
target sub r2, r3, r4 0x900

Stage information not the latches

cycle	PC (latch)	FE	ID	EX	MEM	WB
1	0x800	br				
2	0x804	add	br			
3	0x804	add	br	br		
4	0x804	add	br	br	br	
5	0x804	add	br	br	br	br
6	0x900	sub	br	br	br	br

Handling Branches: II

br target 0x800
add r1, r2, r3 0x804
target sub r2, r3, r4 0x900

cycle	PC (latch)	FE	ID	EX	MEM	WB
1	0x800	br				
2	0x804	add	br			
3	0x804	add	br	br		
4	0x804	add	br	br	br	
5	0x900	sub	br	br	br	br
6	0x904	add	sub	br	br	br

Handling Branches: III

br target 0x800
add r1, r2, r3 0x804
target sub r2, r3, r4 0x900

cycle	PC (latch)	FE	ID	EX	MEM	WB
1	0x800	br				
2	0x804	add	br			
3	0x804	add	br	br		
4	0x900	sub	br	br	br	
5	0x904	add	sub	br	br	br
6	0x908	mul	add	sub	br	br

Handling Branches: IV

```

br target 0x800
add r1, r2, r3 0x804
target sub r2, r3, r4 0x900
    
```

cycle	PC (latch)	FE	ID	EX	MEM	WB
1	0x800	br				
2	0x804	add	br			
3	0x804	sub	br			
4	0x904	mul	sub	br		
5	0x908	add	mul	sub	br	
6	0x90b	div	add	mul	sub	br

Handling Branches: IV

Any problem?

What is the critical length?
 Case:IV Branch unit + mux + I-cache access time
 Case:III I-cache access time
 Why does it matter?
 cycle time

Handling Branches: V

```

br target 0x800
add r1, r2, r3 0x804
target sub r2, r3, r4 0x900
    
```

cycle	PC (latch)	FE	ID	EX	MEM	WB
1	0x800	br				
2	0x804	add	br			
3	0x804	add	br			
4	0x900	sub	br	br		
5	0x904	add	sub	br	br	
6	0x908	mul	add	sub	br	br

What if we


```

0x800    sub r1, r2,r3
0x804    add r4, r2,r3
0x808    br   target
0x80b
0x810
0x900 target mul r2, r3,r4
            
```

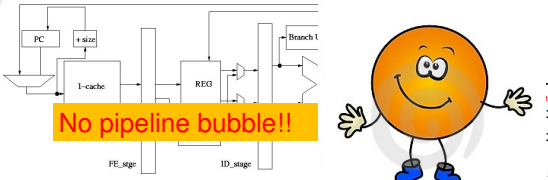
```

0x900 target mul r2, r3,r4
            
```


Change the rule!
Always execute the next two instructions after a branch



Rule: Always execute the next two instructions after a branch




cycle	Fetch addr	FE	ID	E		
1	0x800	br				
2	0x804	sub	br			
3	0x808	add	sub	br		
4	0x900	mul	add	sub	br	
5	0x904	div	mul	add	sub	br
6	0x908	add	div	mul	add	sub
7	0x90b	sub	add	div	mul	add



Delayed branch


- N-cycle delay slot
- The compiler fills out useful instructions inside the delay slot
- Different options:
 - Fill the slot from before the branch instruction
 - Restriction: branch must not depend on result of the filled instruction
 - Fill the slot from the target of the branch instruction
 - Restriction: should be OK to execute instruction even if not taken
 - Fill the slot from fall through of the branch
 - Restriction: should be OK to execute instruction even if taken

Still Cancel or nullifying instructions



Remark

- Many DSP architecture, older RISC, MIPS, PA-RISC, SPARC.
- Delayed branches are architecturally ~~invisible~~ ^{visible}
 - Advantage:
 - better performance
 - Disadvantage:
 - what if implementation changes?
 - Deeper pipeline-> more branch delays?
- Interrupt/exceptions?
 - Where to go back?
- Combining with a branch predictor?



Georgia Tech College of Computing
