Pipelined Design

CPE380, Spring 2022

Hank Dietz

http://aggregate.org/hankd/

University of Kentucky

Different Implementations

- **Multi-cycle** MIPS, multiple CPI:

  `http://aggregate.org/EE380/multiv.html`

- **Single-cycle** MIPS, 1 CPI, but slow clock:

  `http://aggregate.org/EE380/onebeq.html`

- **Pipelined** MIPS, multiple CPI, but fast clock and *throughput* up to 1 instruction/cycle

Basic Pipelining

- Single-cycle **control signals move through the pipe** along with the data
- Divide single-cycle design into **equal-delay stages**, adding **buffers between stages**
  - Ideally, ***n*** stages gives ***n***X throughput
  - Usually <***n***X, and ***n*** can't be arbitrarily large
- Throughput comes from having useful work in all stages all the time – avoiding **bubbles**

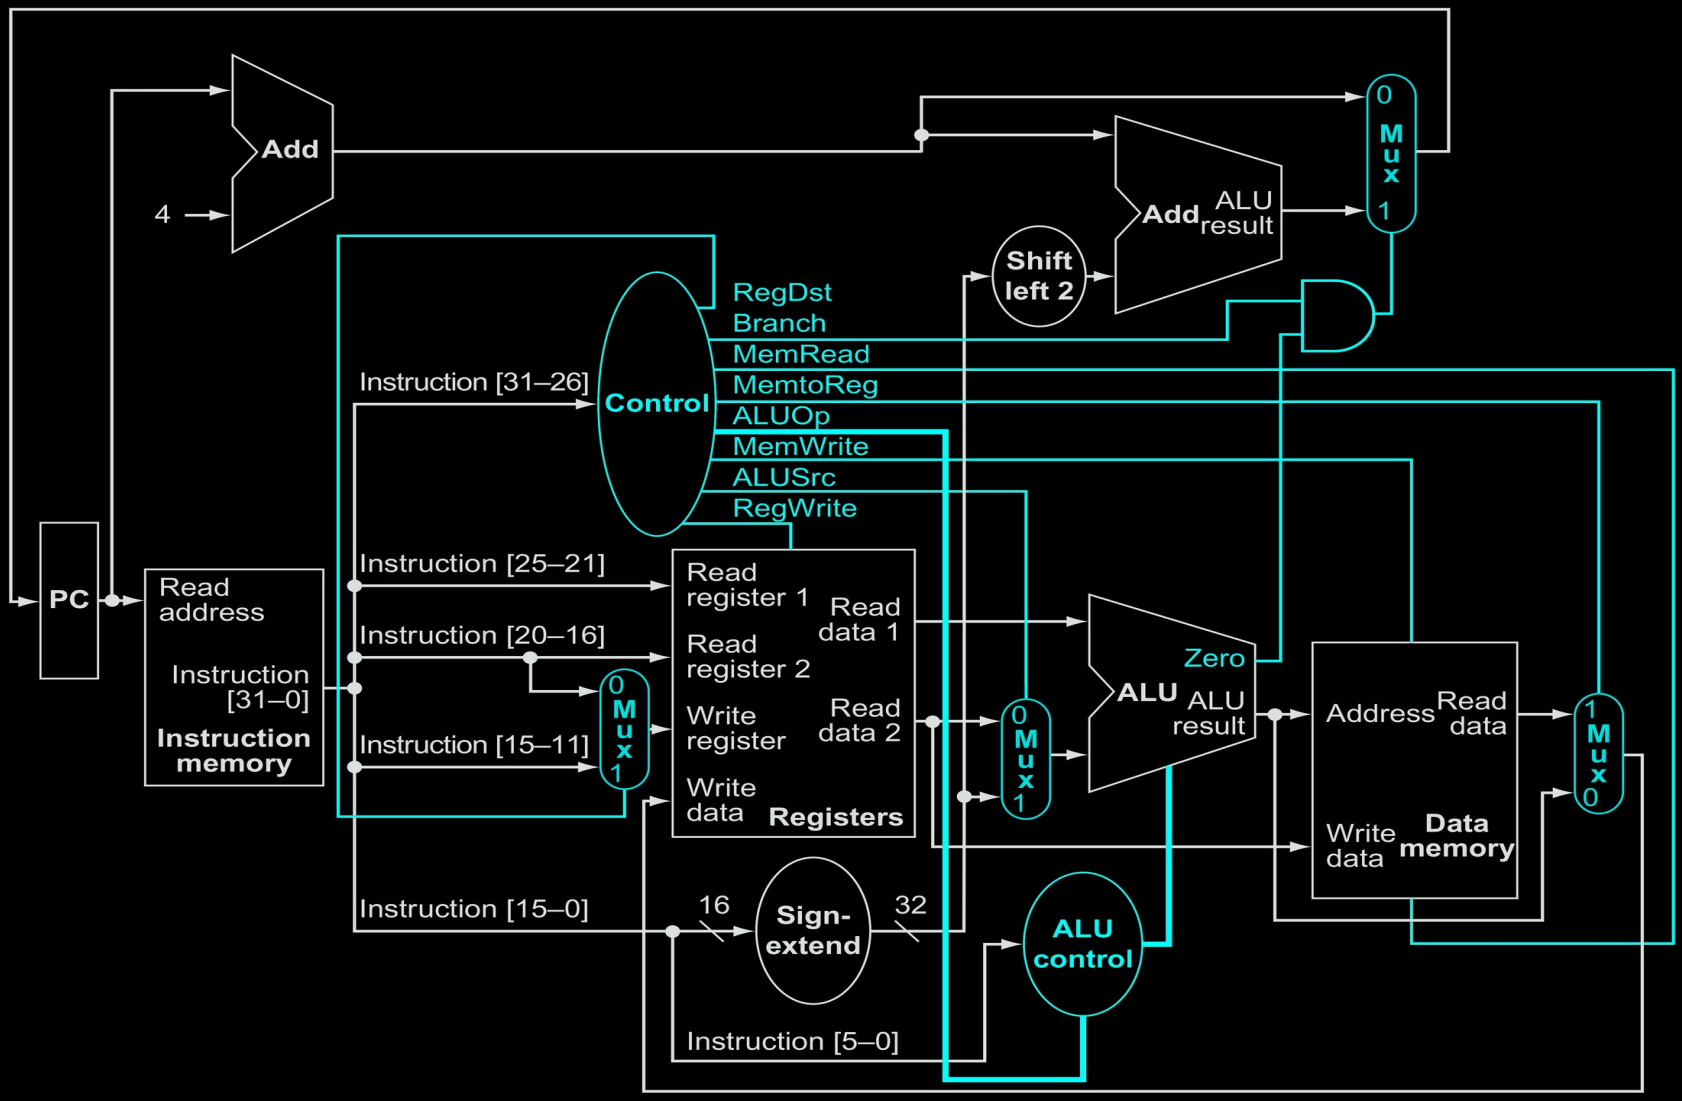

The single cycle design...

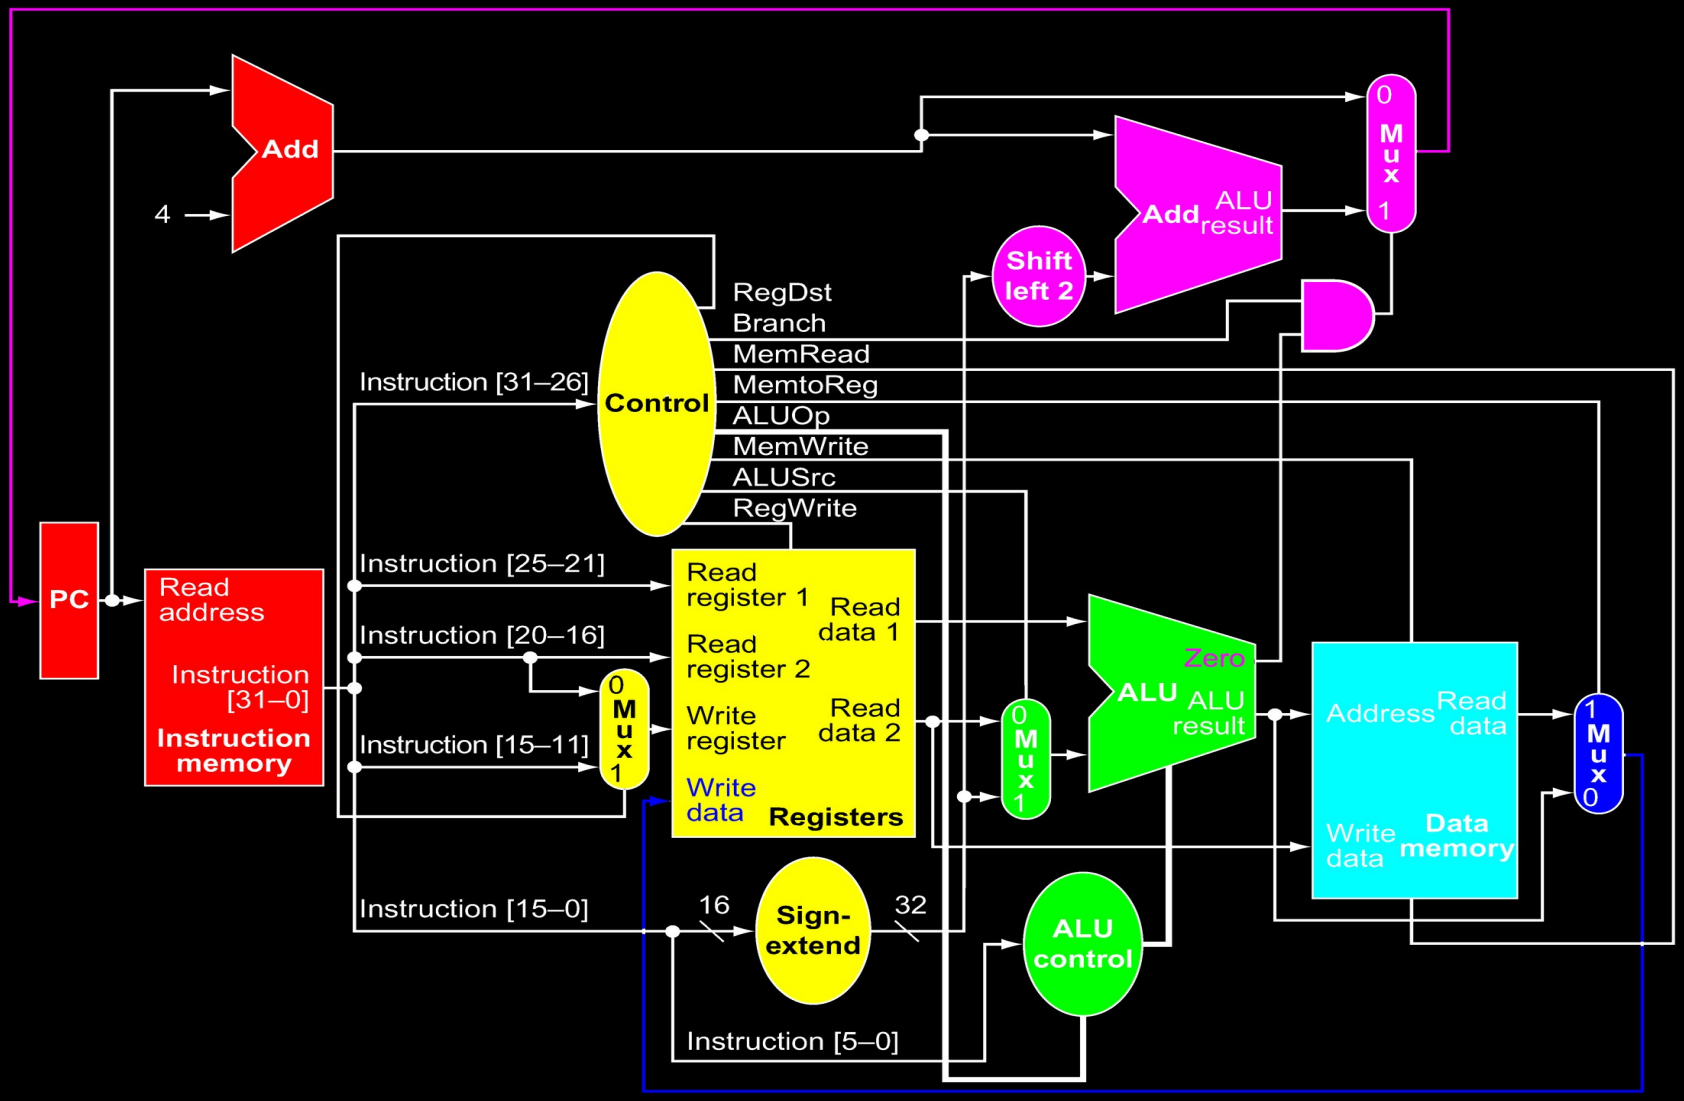

IF: Instruction Fetch EX: Execute WB: Write Back

ID: Instruction Decode MEM: Memory Access what is this?

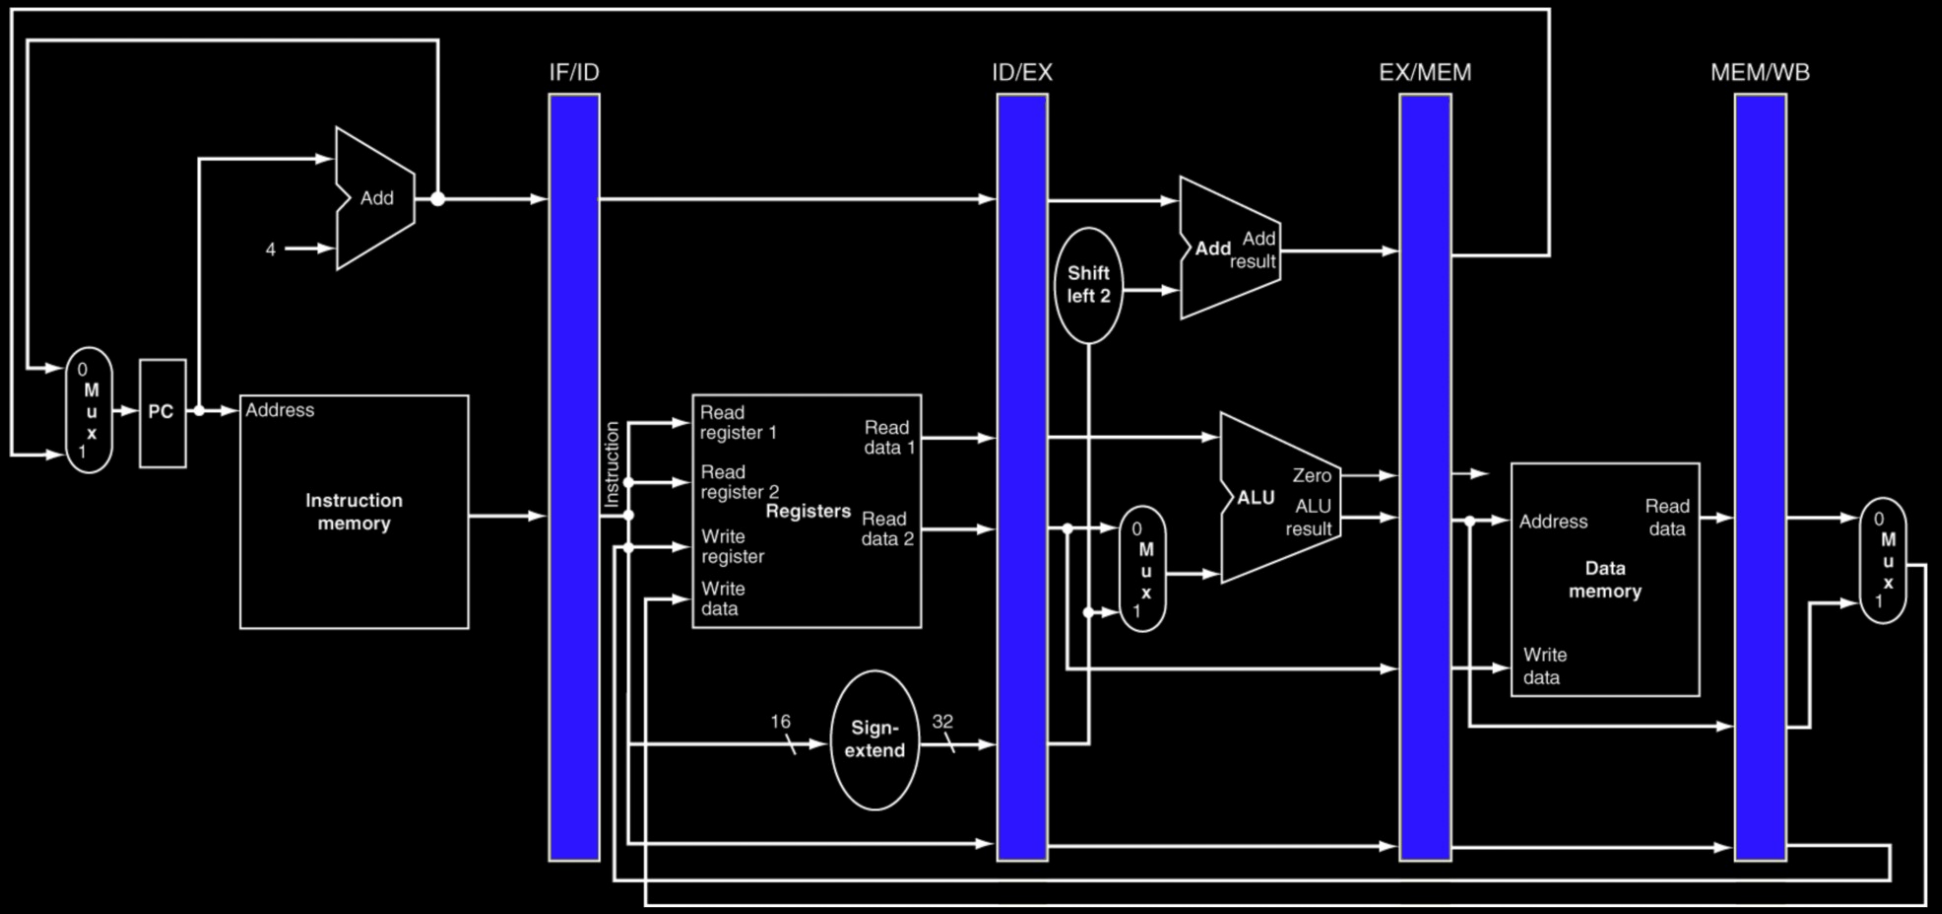

- Add buffers between pipe stages...

Pipeline Throughput

IF	ID	EX	MEM	WB
addu $t0,$t1,$t2				–
	addu $t0,$t1,$t2			–
		addu $t0,$t1,$t2		–
			addu $t0,$t1,$t2	–
and $t3,$t4,$t5				addu $t0,$t1,$t2
	and $t3,$t4,$t5			–
		and $t3,$t4,$t5		–
			and $t3,$t4,$t5	–
. . .				and $t3,$t4,$t5

IF	ID	EX	MEM	WB
addu $t0,$t1,$t2				–
and $t3,$t4,$t5	addu $t0,$t1,$t2			–
addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2		–
lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2	–
sw $t1,4($t9)	lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2
beq $t1,$t2,lab	sw $t1,4($t9)	lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5
…	beq $t1,$t2,lab	sw $t1,4($t9)	lw $t7,0($t8)	addiu $t6,$0,1
…	…	beq $t1,$t2,lab	sw $t1,4($t9)	lw $t7,0($t8)
…	…	…	beq $t1,$t2,lab	sw $t1,4($t9)
…	…	…	…	beq $t1,$t2,lab

Pipeline Bubble: nop

IF	ID	EX	MEM	WB
addu $t0,$t1,$t2				-
and $t3,$t4,$t5	addu $t0,$t1,$t2			-
addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2		-
lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2	-
sw $t1,4($t9)	lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2
beq $t1,$t2,lab	sw $t1,4($t9)	lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5
…	beq $t1,$t2,lab	sw $t1,4($t9)	lw $t7,0($t8)	addiu $t6,$0,1
…	…	beq $t1,$t2,lab	sw $t1,4($t9)	lw $t7,0($t8)
…	…	…	beq $t1,$t2,lab	sw $t1,4($t9)
…	…	…	…	beq $t1,$t2,lab

IF	ID	EX	MEM	WB
addu $t0,$t1,$t2				-
and $t3,$t4,$t5	addu $t0,$t1,$t2			-
addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2		-
lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2	-
sw $t1,4($t7)	lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5	addu $t0,$t1,$t2
beq $t1,$t2,lab	sw $t1,4($t7)	lw $t7,0($t8)	addiu $t6,$0,1	and $t3,$t4,$t5
beq $t1,$t2,lab	sw $t1,4($t7)	nop	lw $t7,0($t8)	addiu $t6,$0,1
beq $t1,$t2,lab	sw $t1,4($t7)	nop	nop	lw $t7,0($t8)
beq $t1,$t2,lab	sw $t1,4($t7)	nop	nop	nop
…	beq $t1,$t2,lab	sw $t1,4($t7)	nop	nop
…	…	beq $t1,$t2,lab	sw $t1,4($t7)	nop
…	…	…	beq $t1,$t2,lab	sw $t1,4($t7)
…	…	…	…	beq $t1,$t2,lab

Setting The Stages

Instruction	IF	ID	EX	MEM	WB	Circuit delay
addu	250	100	300	–	100	750ps
and	250	100	100	–	100	550ps
addiu	250	100	300	–	100	750ps
lw	250	100	300	250	100	1000ps
sw	250	100	300	100	100	850ps
beq	250	100	300	–	–	750ps

- Single-cycle limited by `lw` to 1GHz clock
- 5-stage pipeline limited by **EX**
  - 300ps latency might allow 3.33GHz clock
  - If buffer between stages is 100ps, 2.5GHz
- Multi-cycle might be 5 cycles @2.5GHz

Why A Bubble?

- Structural hazards: can't do 2 things on one unit of HW simultaneously – single-cycle cures this!
- Data dependence: need results from previous computations, which might not be done yet
- Control dependence
  - Computation of branch target address
  - Conditional branch/jump taken vs. not taken

Dependence Analysis

- Use or R: reads the value bound to a name
- Def or W: binds a new value to a name
- True dependence: carries a value, D→U, RAW

```
add $t0,$t1,$t2    or $t3,$t0,$t4
```

- Anti-dependence: kills a value, U←D, WAR

```
add $t0,$t1,$t2    or $t1,$t3,$t4
```

- Output dependence: kills a value, D→D, WAW

```
add $t0,$t1,$t2    or $t0,$t3,$t4
```

When Dependence Matters

- True dependence causes a delay

```
add $t0,$t1,$t2
or  $t3,$t0,$t4   //wait for $t0
```

- **Other types don't**

How To Handle Dependence

- Have programmer/assembler pad with NOPs
  - Violates ISA concept (how much padding?); too little padding gives wrong answers, too much wastes time
  - Makes program bigger

IF	ID	EX	MEM	WB
add $t0,$t1,$t2	…	…	…	…
nop	add $t0,$t1,$t2	…	…	…
nop	nop	add $t0,$t1,$t2	…	…
nop	nop	nop	add $t0,$t1,$t2	…
or $t3,$t0,$t4	nop	nop	nop	add $t0,$t1,$t2
…	or $t3,$t0,$t4	nop	nop	nop
…	…	or $t3,$t0,$t4	nop	nop
…	…	…	or $t3,$t0,$t4	nop
…	…	…	…	or $t3,$t0,$t4

How To Handle Dependence

- Hardware interlock
  - rs or rt in ID same as dest in EX, MEM, WB
  - Detects dependence & stalls until satisfied
  - **`nop`** is really **`or`** with side-effects disabled: MemRead, MemWrite, RegWrite, & Branch

IF	ID	EX	MEM	WB
`add $t0,$t1,$t2`	…	…	…	…
`or $t3,$t0,$t4`	`add $t0,$t1,$t2`	…	…	…
…	`or $t3,$t0,$t4`	`add $t0,$t1,$t2`	…	…
…	`or $t3,$t0,$t4`	`nop`	`add $t0,$t1,$t2`	…
…	`or $t3,$t0,$t4`	`nop`	`nop`	`add $t0,$t1,$t2`
…	…	`or $t3,$t0,$t4`	`nop`	`nop`
…	…	…	`or $t3,$t0,$t4`	`nop`
…	…	…	…	`or $t3,$t0,$t4`

How To Handle Dependence

- Value forwarding: use interlock circuitry to find value & forward it to ID stage output buffer
  - ALU result ready from EX, so **NO DELAY**!
  - **`lw`** result ready from MEM, so 1 cycle delay
  - Value still gets stored in register in WB
  - Works great, but adds lots of datapath…

IF	ID	EX	MEM	WB
`lw $t0,0($t1)`	…	…	…	…
`or $t3,$t0,$t4`	`lw $t0,0($t1)`	…	…	…
…	`or $t3,$t0,$t4`	`lw $t0,0($t1)`	…	…
…	`or $t3,$t0,$t4`	`nop`	`lw $t0,0($t1)`	…
…	…	`or $t3,$t0,$t4`	`nop`	`lw $t0,0($t1)`
…	…	…	`or $t3,$t0,$t4`	`nop`
…	…	…	…	`or $t3,$t0,$t4`

Using Dependence Information

- True dependence causes a delay

```
lw   $t0,0($t1)
or   $t3,$t0,$t4   //wait for $t0
xor  $t5,$t6,$t7   //takes a cycle
```

- Use **code motion** or **out-of-order execution** to **execute useful instructions while waiting!**

```
lw   $t0,0($t1)
xor  $t5,$t6,$t7   //this is free!
or   $t3,$t0,$t4   //wait for $t0
```

Reordering Instructions

- **Code motion** by compiler/assembler:
  - \+ Sees whole program, has time for analysis
  - – Limited number of registers, imperfect info (e.g., does not know function unit details)
- **Out-of-order execution** by hardware:
  - \+ Perfect info, can use register renaming (with more registers than an instruction can name)
  - – Limited to instructions buffered
- **Modern systems do both**

Computing "May Move"

- No dependence allows code to move
- Anti-dependence and output-dependence impose an ordering constraint… but renaming can remove that constraint:

```
add $t0,$t1,$t2        add $t0,$t1,$t2
or  $t1,$t3,$t4        or  $t0,$t3,$t4
      becomes                becomes
add $t0,$t1,$t2        add $t0,$t1,$t2
or  $t1',$t3,$t4       or  $t0',$t3,$t4
```

A Real Processor: AMD Athlon

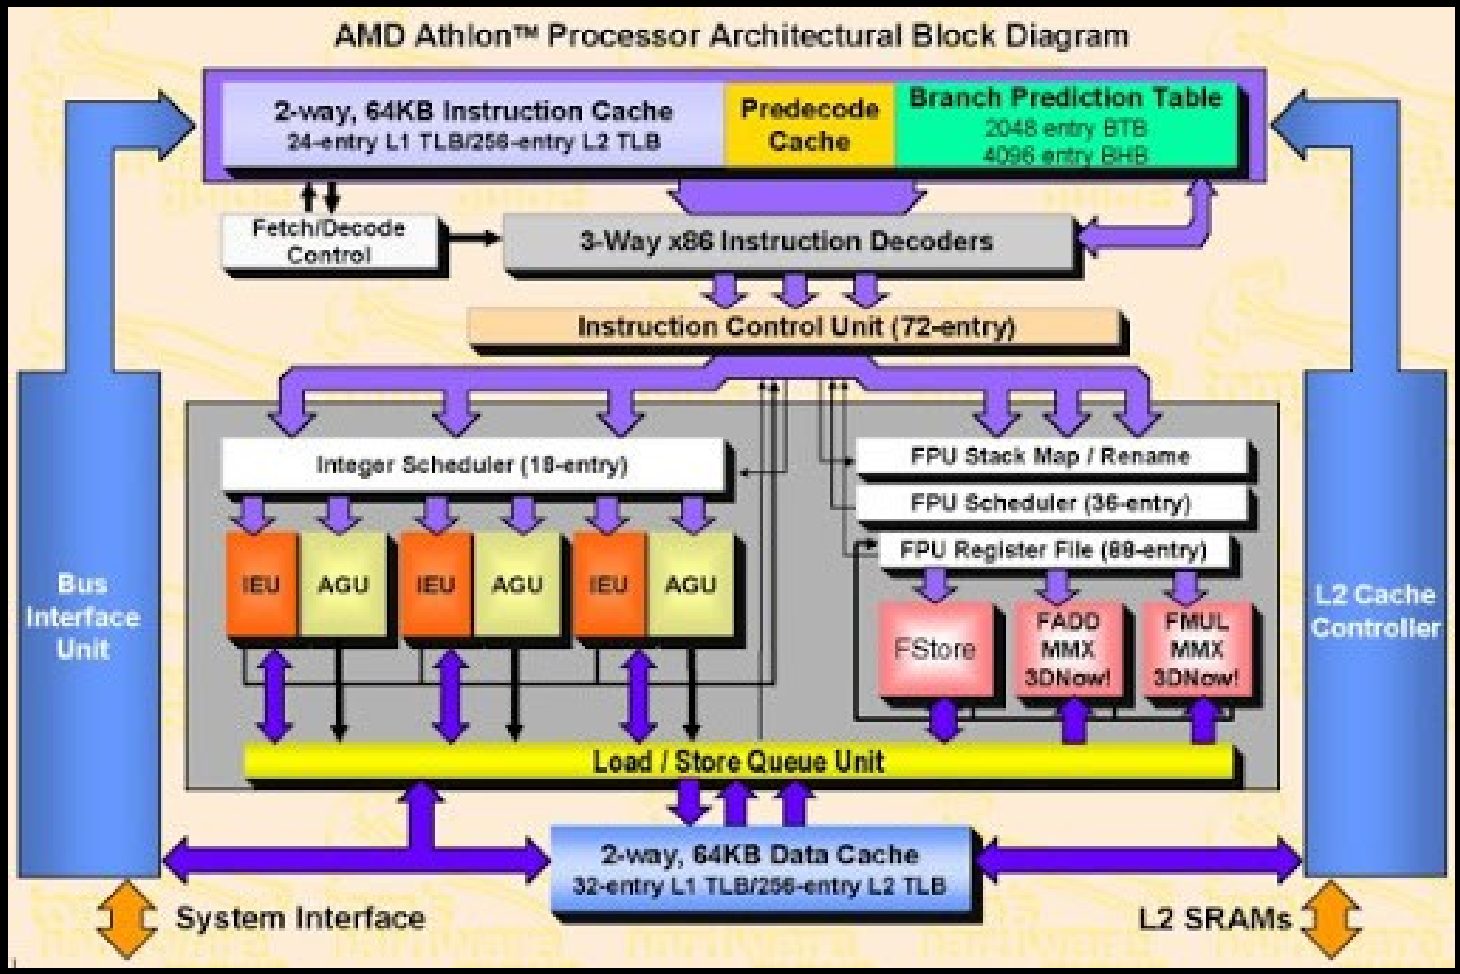

E.g., IA32 has 8 float registers, but this has 88!

Parallel Execution Orders

- IF stage can fork to feed multiple pipes
  - SIMD: GPU, SWAR, and Vector
  - VLIW: Very Long Instruction Word
  - EPIC: Explicitly Parallel Instruction Computer
  - Superscalar: instructions grouped at runtime

Athlon can sustain 3 IPC:
3 Integer Execution Units
3 Address Generation Units
Float store, add, & multiply

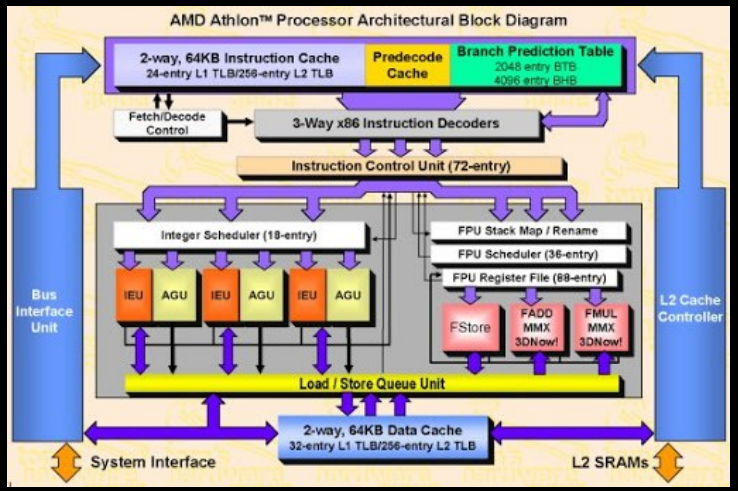

Computing The Branch Target

- Branch instructions encode offset, not address
  - Add PC + offset *typically* takes a cycle
  - What do we do while waiting for address?
- Hardware interlock & stall
- Delayed branch: branch *after* next instruction
- **BTB**: **Branch Target Buffer**
  - Caches branch {PC: target} pairs and looks 'em up at same time as fetching instruction
  - No help 1st time, but no delay on repeats

To Take, Or Not To Take?

- When do we know if conditional jump or branch is taken or not taken?

Determined only after EX stage?

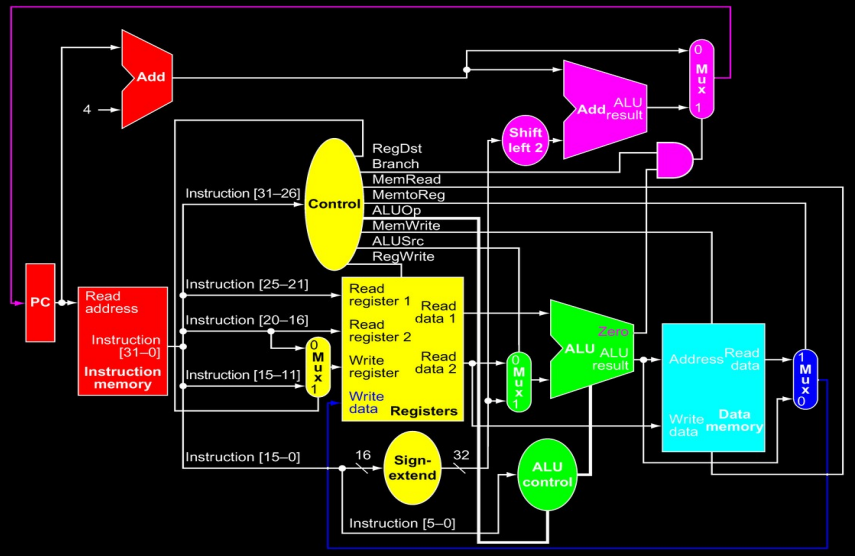

- NOP padding or HW interlock very inefficient!
  - Usually determined in a late pipe stage
  - Blocks ALL progress (including superscalar)

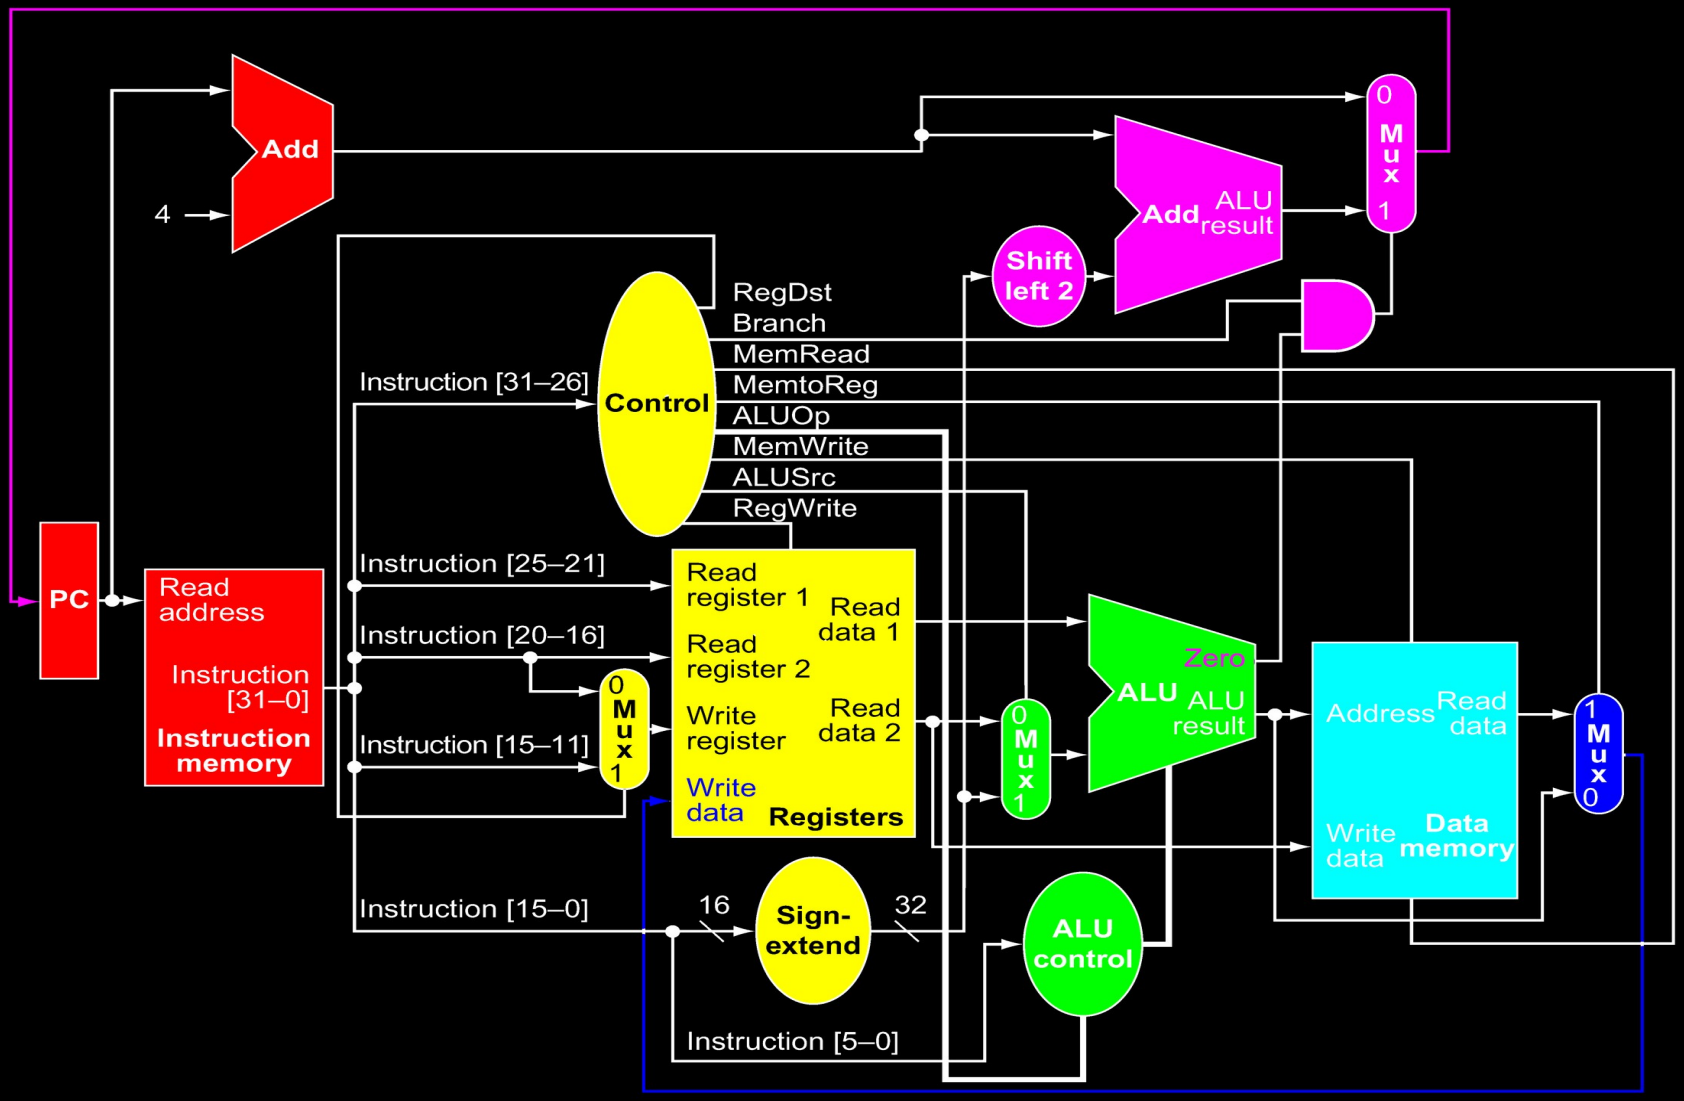

IF: Instruction Fetch EX: Execute WB: Write Back

ID: Instruction Decode MEM: Memory Access what is this?

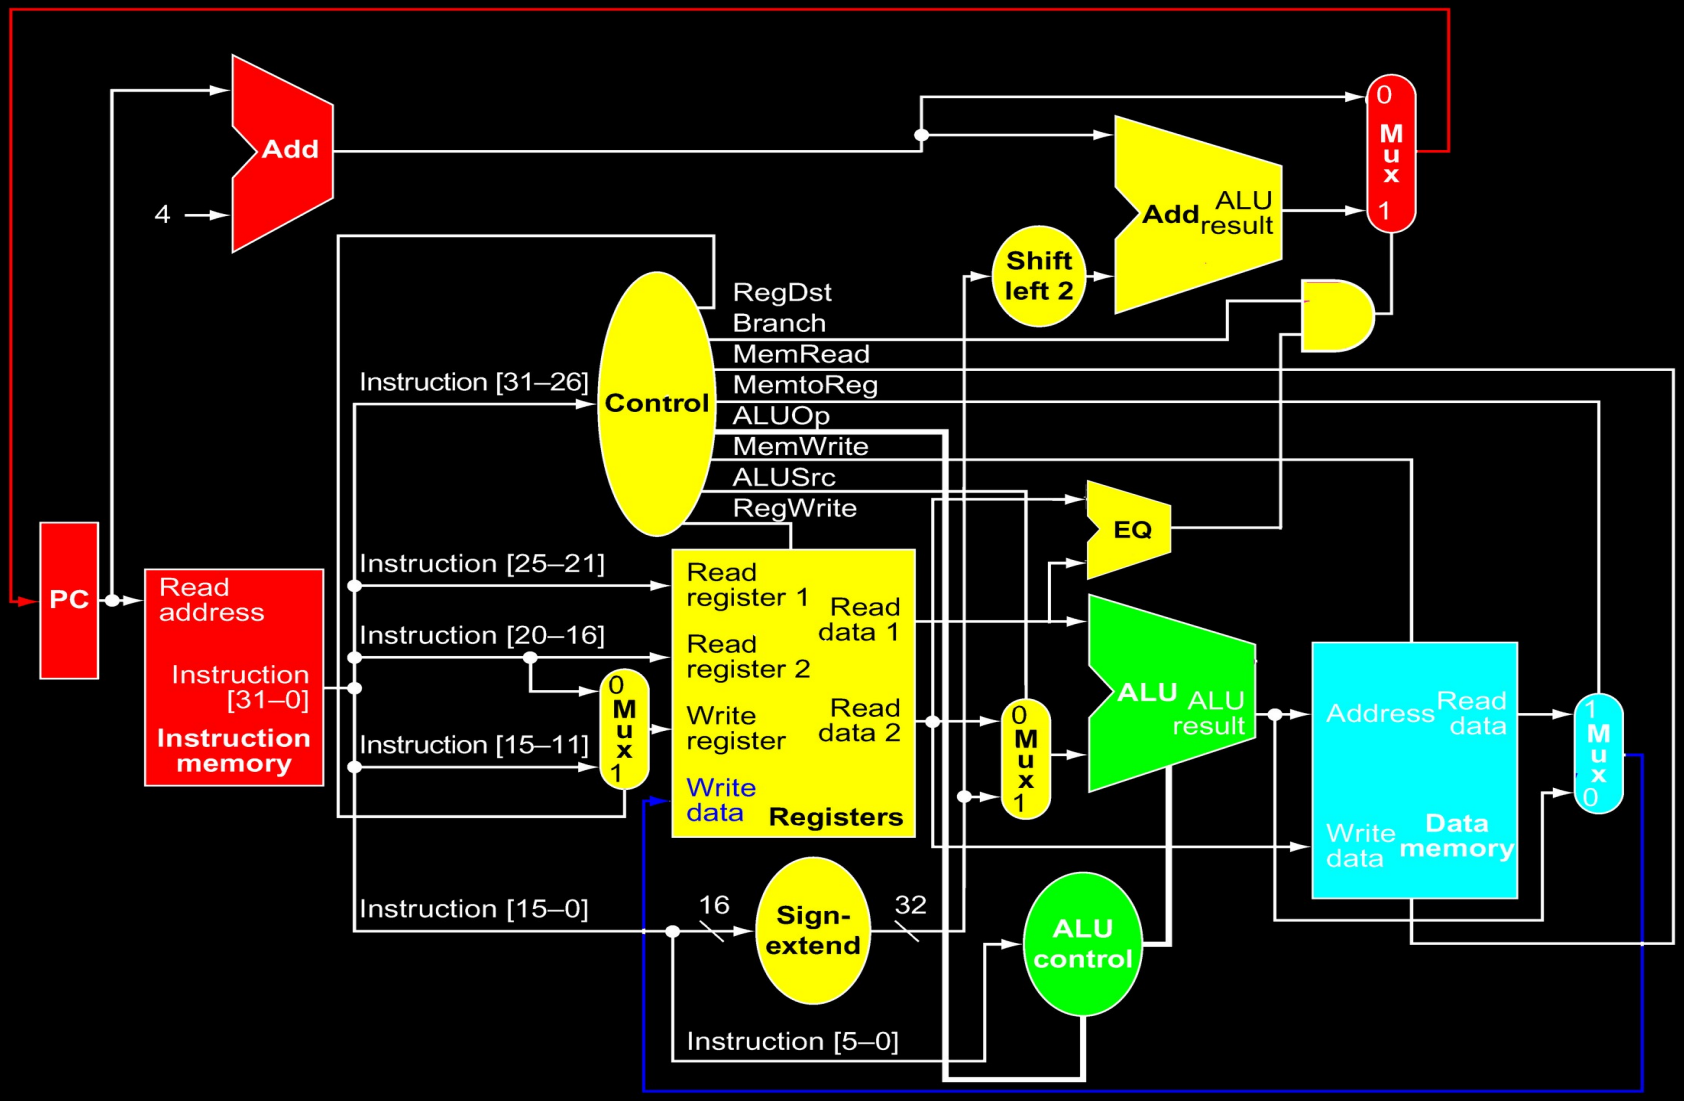

IF: Instruction Fetch EX: Execute WB: Write Back
ID: Instruction Decode MEM: Memory Access

Branch Prediction. Do I Feel Lucky?

Guess wrong and some instructions are gonna die
*(well, we actually say they're **squashed**)*

- **Always not taken** – the easiest guess
- **Always taken** – more often right for do loops
- **Always BOTH not taken and taken**
  - Was tried using dual pipe of Pentium Pro
  - Always wastes 50% of instruction fetch!
- **Forward not, backward taken**

Those who cannot remember the past are condemned to repeat it. – *G. Santayana*

- **BHB: Branch History Buffer**
  - Jump & branch taken/not-taken "history," but actually records a prediction state, *not* history
  - Indexed by PC; can be merged with BTB

Simplest is 4-state, just 2 bits/entry:

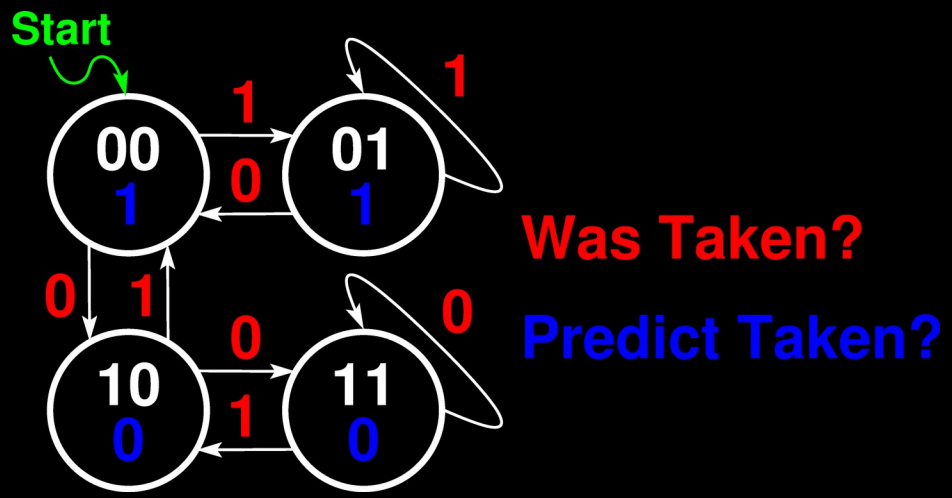

Fancier Prediction Methods

- Two types of prediction history recorded:
  - History of *this particular* branch/jump
  - History of *last K* branches/jumps anywhere

- More complex state machines can recognize *much* more complex patterns; also can combine multiple predictors as a tournament predictor

- Some systems also allow compiler hints by marking usually taken vs. usually not taken

A Real Processor: AMD Athlon

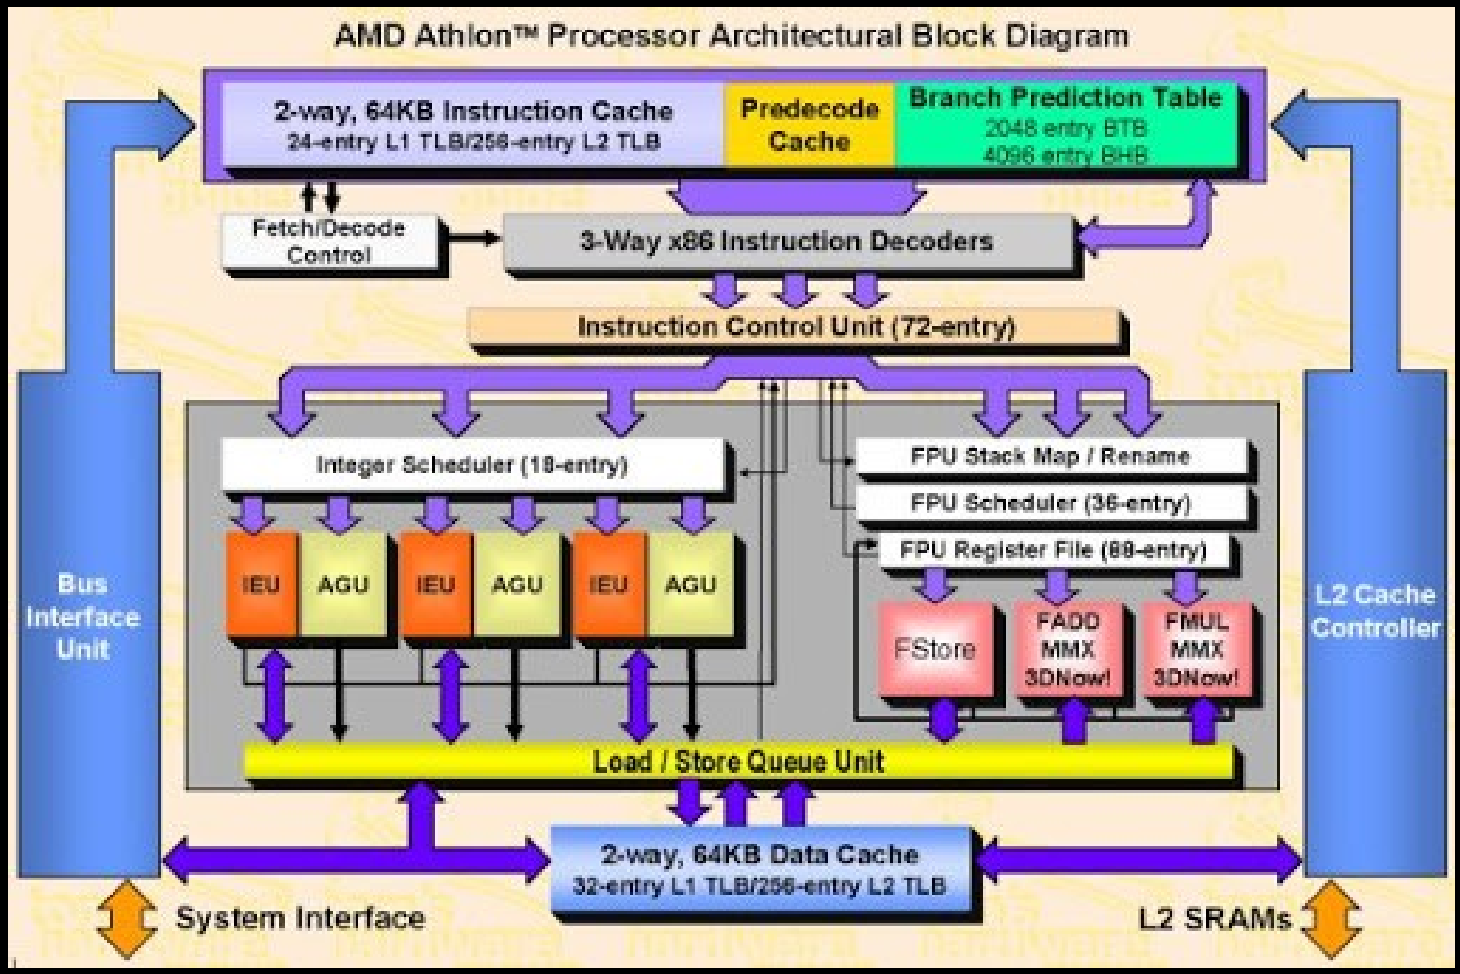

Note 2048-entry BTB and 4096-entry BHB

Verilog Implementation

- Like you'd expect:
  - Can reuse basic single-cycle design
  - Each stage becomes it's own **`always`**
  - Need multiple copies of some signals, one for each stage that uses them
- Not like you'd expect:
  - Some things don't follow pipe flow
  - Some non-stages should be separate things

Owner Computes

- For example, who updates the PC?
  - IF sets PC = PC + 4
  - ID sets PC = branch target
  - EX, MEM, or WB forces PC = PC because data dependence blocks the pipe
- How to coordinate access to shared data?
  - Multiple readers is fine, with one writer
  - Could use locks, semaphores, etc.
  - Easiest is only one entity can update each: the **owner** picks the value and is only writer

A Pipelined Verilog Version

- Organized as parallel-executing chunks for:
  - IF stage: reads and writes **`IF_`**
  - ID stage: reads **`IF_`**, writes **`ID_`**
  - EX stage: reads **`ID_`**, writes **`EX_`**
  - MEM stage: reads **`EX_`**, writes **`MEM_`**
  - WB stage: reads **`MEM_`**
  - Are we **`running`**?
  - Are mispredicted instructions **`squashed`**?
  - Are we **`blocked`** by data dependence or forwarding values?
  - Simulation tracing support

A Pipelined Verilog Version

```
// Testbench options
`define RUNTIME 200     // How long simulator can run
`define CLKDEL  1       // Clock transition delay
`define TRACE   1       // enable simulation trace

// Types
`define WORD    [31:0] // size of a data word
`define ADDR    [31:0] // size of a memory address
`define INST    [31:0] // size of an instruction
`define REG     [4:0]  // size of a register number
`define REGCNT  [31:0] // register count
`define MEMCNT  [511:0] // memory count inplemented
`define OPCODE  [5:0]  // 6-bit opcodes
`define ALUOP   [3:0]  // Simplified ALU codes

// Fields
`define OP      [31:26] // opcode field
`define RS      [25:21] // rs field
`define RT      [20:16] // rt field
`define RD      [15:11] // rd field
`define IMM     [15:0]  // immediate/offset field
`define SHAMT   [10:6]  // shift ammount
`define FUNCT   [5:0]   // function code (opcode extension)
`define JADDR   [25:0]  // jump address field
`define JPACK(R,D,J)        begin R`OP=0; R`JADDR=J; end
`define IPACK(R,O,S,T,I)    begin R`OP=O; R`RS=S; R`RT=T; R`IMM=I; end
`define RPACK(R,S,T,D,SH,FU)    begin R`OP=`RTYPE; R`RS=S; R`RT=T; R`RD=D; R`SHAMT=SH; R`FUNCT=FU; end
`define NOP     `OR     // Null OPeration is or $0,$0,$0

// Instruction encoding
`define RTYPE   6'h00   // OP field for all RTYPE instructions
`define BEQ     6'h04   // OP field
`define ADDIU   6'h09   // OP field
`define SLTIU   6'h0b   // OP field
`define ANDI    6'h0c   // OP field
`define ORI     6'h0d   // OP field
`define XORI    6'h0e   // OP field
`define LUI     6'h0f   // OP field
`define LW      6'h23   // OP field
`define SW      6'h2b   // OP field
`define ADDU    6'h21   // FUNCT field
`define SUBU    6'h23   // FUNCT field
`define AND     6'h24   // FUNCT field
`define OR      6'h25   // FUNCT field
`define XOR     6'h26   // FUNCT field
`define SLTU    6'h2b   // FUNCT field

// Simplified ALU codes, default to lui
`define ALUADD  4'b0000
`define ALUSUB  4'b0001
`define ALUADD  4'b0010
`define ALUSUB  4'b0110
`define ALUSLT  4'b0111
`define ALULUI  4'b1000
`define ALUXOR  4'b1111

// Generic multi-cycle processor
module processor(halt, reset, clk);
output reg halt;
input reset, clk;
reg `WORD `MEMCNT;
reg `WORD r `REGCNT;

// Initialize register file and memory
initial begin
    r[1] = 22; r[2] = 1; r[3] = 42;
    r[4] = 601; r[5] = 13813; r[0] = -1;
    `RPACK(m[0], 2, 3, 1, 0, `ADDU)
    `RPACK(m[1], 2, 3, 1, 0, `SLTU)
    `RPACK(m[2], 3, 5, 1, 0, `AND)
    `RPACK(m[3], 5, 3, 1, 0, `OR)
    `RPACK(m[4], 3, 5, 1, 0, `XOR)
    `RPACK(m[5], 4, 3, 1, 0, `SUBU)
    `IPACK(m[6], `ADDIU, 3, 1, -1)
    `IPACK(m[7], `SLTIU, 5, 1, 12345)
    `IPACK(m[8], `ANDI, 3, 1, 3)
    `IPACK(m[9], `ORI, 3, 1, 3)
    `IPACK(m[10], `XORI, 3, 1, 3)
    `IPACK(m[11], `LUI, 0, 1, 1)
    `IPACK(m[12], `LW, 2, 1, 1023)
    `IPACK(m[13], `SW, 0, 2, 1024)
    `IPACK(m[14], `ADDIU, 6, 6, 1)
    `IPACK(m[15], `BEQ, 0, 6, -2)
    m[16] = 0;
    m[250] = 22;
end

// IF registers
reg `ADDR IF_pc;
reg `INST IF_ir;

// ID registers
reg RegDst, Branch, MemRead, MemWrite, ALUSrc, RegWrite, Bad, ID_MemRead, ID_MemWrite, ID_Bad;
reg `ALUOP ALUOp, ID_ALUOp;
reg `WORD s, t, imm, ID_s, ID_t, ID_src;
reg `REG ID_rd;
reg `ADDR target;
reg squash;

// EX registers
reg EX_MemRead, EX_MemWrite, EX_Bad;
reg `WORD alu, EX_alu, EX_t;
reg `REG EX_rd;

// MEM registers
reg MEM_Bad;
reg `WORD v, MEM_v;
reg `REG MEM_rd;
```

```
// Running state?
wire running;
assign running = (!!halt) && (!reset);

// Squash instruction fetched on a mispredicted branch
wire `INST squashed;
assign squashed = squash ? `NOP : IF_ir;

// Are we blocked by a dependence on rs or rt?
// Also handles all forwarding
wire EX_depx, EX_dept, MEM_depx, MEM_dept, WB_depx, WB_dept, canfwdx, canfwdt, dep, blocked;
wire `WORD fwdx, fwdt;
assign ID_needs = (IF_ir `RS != 0);
assign ID_needt = (IF_ir `RT != 0);
assign EX_depx = (IF_ir `RS == EX_rd);
assign EX_dept = (IF_ir `RT == EX_rd);
assign MEM_depx = (IF_ir `RS == MEM_rd);
assign MEM_dept = (IF_ir `RT == MEM_rd);
assign WB_depx = (IF_ir `RS == WB_rd);
assign WB_dept = (IF_ir `RT == WB_rd);
assign canfwdx = ID_needs && ((EX_depx && !MemRead) || ((!EX_depx) && (MEM_depx || WB_depx)));
assign canfwdt = ID_needt && ((EX_dept && !MemRead) || ((!EX_dept) && (MEM_dept || WB_dept)));
assign fwdx = (EX_depx ? alu : (MEM_depx ? v : MEM_v));
assign fwdt = (EX_dept ? alu : (MEM_dept ? v : MEM_v));
assign dep = ((ID_needs && (EX_depx || MEM_depx || WB_depx)) ||
              (ID_needt && (EX_dept || MEM_dept || WB_dept)));
assign blocked = (dep && ((ID_needs && !canfwdx) || (ID_needt && !canfwdt)));

// IF: Instruction Fetch stage
always @(posedge clk) if (running && !blocked) begin
  IF_ir <= m[squash ? target : IF_pc >> 2];
  IF_pc <= (squash ? target : IF_pc) + 4;
end

// ID: Instruction Decode stage
always @(posedge clk) if (running && !ID_Bad) begin
  if (blocked) begin
    ID_s <= 0;
    ID_t <= 0;
    ID_src <= 0;
    ID_rd <= 0; // not writing
    ID_MemRead <= 0; // not reading
    ID_ALUOp <= 0;
    ID_MemWrite <= 0; // not storing
    ID_Bad <= 0;
  end else begin
    // Brute force decode of instruction
    case (squashed `OP)
      `RTYPE: begin
        case (squashed `FUNCT)
          `ADDU:   begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUADD; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `SUBU:   begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUSUB; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `AND:    begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUAND; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `OR:     begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `XOR:    begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUXOR; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `SLTU:   begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUSLT; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          default: begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=0; RegWrite=0; Bad=1; end
        endcase
      end
      `BEQ:    begin RegDst=0; Branch=1; MemRead=0; ALUOp=`ALUSUB; MemWrite=0; ALUSrc=0; RegWrite=0; Bad=0; end
      `ADDIU:  begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUADD; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `SLTIU:  begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUSLT; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `ANDI:   begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUAND; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `ORI:    begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `XORI:   begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUXOR; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `LUI:    begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALULUI; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `LW:     begin RegDst=0; Branch=0; MemRead=1; ALUOp=`ALUADD; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `SW:     begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUADD; MemWrite=1; ALUSrc=1; RegWrite=0; Bad=0; end
      default: begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=0; RegWrite=0; Bad=1; end
    endcase

    s = (canfwdx ? fwdx : r[squashed `RS]);
    t = (canfwdt ? fwdt : r[squashed `RT]);
    imm = {{16{squashed[15]}}, squashed `IMM};
    target <= IF_pc + {imm[29:0], 2'b00};

    squash <= (Branch && (s == t));
    ID_s <= s;
    ID_t <= t;
    ID_src <= (ALUSrc ? imm : t);
    ID_rd <= (RegWrite ? (RegDst ? squashed `RD : squashed `RT) : 0); // $0 if not writing
    ID_MemRead <= MemRead;
    ID_ALUOp <= ALUOp;
    ID_MemWrite <= MemWrite;
    ID_Bad <= Bad;
  end
end

// EX: EXecute stage
always @(posedge clk) if (running) begin
  case (ID_ALUOp)
    `ALUADD: alu = ID_s & ID_src;
    `ALUOR:  alu = ID_s | ID_src;
    `ALUADD: alu = ID_s + ID_src;
    `ALUSUB: alu = ID_s - ID_src;
    `ALUSLT: alu = ID_s < ID_src;
    `ALUXOR: alu = ID_s ^ ID_src;
    default: alu = {ID_src <<< 16};
  endcase

  EX_alu <= alu;
  EX_t <= ID_t;
  EX_rd <= ID_rd;
  EX_MemRead <= ID_MemRead;
  EX_MemWrite <= ID_MemWrite;
  EX_Bad <= ID_Bad;
end
```

Color key: initialization
IF squashed ID blocked EX MEM WB debugging

```
// MEM: data MEMory access stage
always @(posedge clk) if (running) begin
  if (EX_MemRead) v = m[EX_alu >> 2]; else v = EX_alu;
  if (EX_MemWrite) m[EX_alu >> 2] <= EX_t;

  MEM_v <= v;
  MEM_rd <= EX_rd;
  MEM_Bad <= EX_Bad;
end

// WB: register Write Back stage
always @(posedge clk) if (running) begin
  if (MEM_rd) r[MEM_rd] <= MEM_v;
  if (MEM_Bad) halt <= 1;
end

// Reset
always @(posedge clk) if (reset) begin
  halt<=0;
  r[0]<=0;
  squash <= 0; target <= 0;
  IF_ir<= `NOP; IF_pc<=0;
  ID_s<=0; ID_t<=0; ID_src<=0; ID_rd<=0; ID_MemRead<=0; ID_ALUOp<=0; ID_MemWrite<=0; ID_Bad<=0;
  EX_alu<=0; EX_t<=0; EX_rd<=0; EX_MemRead<=0; EX_MemWrite<=0; EX_Bad<=0;
  MEM_v<=0; MEM_rd<=0; MEM_Bad<=0;
end

`ifdef TRACE
// State-by-state trace
always @(posedge clk) if (running) begin
  `define Fs      (canfwdx ? (EX_depx ? "EX " : (MEM_depx ? "MEM " : "WB ")) : "")
  `define Ft      (canfwdt ? (EX_dept ? "EX " : (MEM_dept ? "MEM " : "WB ")) : "")
  $display("\n  running=%1d blocked=%1d s=%x$%1d t=%x$%1d squash=%1d", running, blocked, `Fs, IF_ir `RS, `Ft, IF_ir `RT, squash);
  $display("IF  ir=%x pc=%1d", IF_ir, IF_pc);
  case (IF_ir `OP)
    `RTYPE: begin
      case (IF_ir `FUNCT)
        `ADDU:   $display("IF  addu $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `SUBU:   $display("IF  subu $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `AND:    $display("IF  and $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `OR:     $display("IF  or $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `XOR:    $display("IF  xor $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `SLTU:   $display("IF  sltu $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        default: $display("IF  OP=%x RS=%1d RT=%1d RD=%1d SHAMT=%1d FUNCT=%x", IF_ir `OP, IF_ir `RS, IF_ir `RT, IF_ir `RD, IF_ir `SHAMT, IF_ir `FUNCT);
      endcase
    end
    `BEQ:    $display("IF  beq $%1d,$%1d,offset=$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `ADDIU:  $display("IF  addiu $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `SLTIU:  $display("IF  sltiu $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `ANDI:   $display("IF  andi $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `ORI:    $display("IF  ori $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `XORI:   $display("IF  xori $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `LUI:    $display("IF  lui $%1d,$%1d", IF_ir `RS, IF_ir `IMM);
    `LW:     $display("IF  lw $%1d,$%1d($%1d)", IF_ir `RT, IF_ir `IMM, IF_ir `RS);
    `SW:     $display("IF  sw $%1d,$%1d($%1d)", IF_ir `RT, IF_ir `IMM, IF_ir `RS);
    default: $display("IF  OP=%x RS=%1d RT=%1d IMM=%x", IF_ir `OP, IF_ir `RS, IF_ir `RT, IF_ir `IMM);
  endcase
  if (ID_Bad) $display("ID  illegal instruction");
  else $display("ID  s=%1d t=%1d src=%1d rd=%1d MemRead=%b ALUOp=%b MemWrite=%b", ID_s, ID_t, ID_src, ID_rd, ID_MemRead, ID_ALUOp, ID_MemWrite);
  if (EX_Bad) $display("EX  illegal instruction");
  else $display("EX  alu=%1d t=%1d rd=%1d MemRead=%b MemWrite=%b", EX_alu, EX_t, EX_rd, EX_MemRead, EX_MemWrite);
  if (MEM_Bad) $display("MEM illegal instruction");
  else begin
    if (EX_MemWrite) $display("MEM m[%1d]=%1d v=%1d rd=%1d", EX_alu, EX_t, MEM_v, MEM_rd);
    else $display("MEM v=%1d rd=%1d", MEM_v, MEM_rd);
  end
  if (MEM_rd) $display("WB  r[%1d]=%1d", MEM_rd, MEM_v);
end
`endif
endmodule

// Testbench
module bench;
wire halt;
reg reset = 1;
reg clk = 0;

processor PE(halt, reset, clk);

initial begin
  # `CLKDEL clk = 1;
  # `CLKDEL clk = 0;
  reset = 0;
  while ($time < `RUNTIME) && !halt) begin
    # `CLKDEL clk = 1;
    # `CLKDEL clk = 0;
  end
end
endmodule
```

IF: Instruction Fetch stage

- Really simple…
  - Memory is `WORD, not byte, so address>>2
  - The only thing setting IF_ir and IF_pc

```
// IF: Instruction Fetch stage
always @(posedge clk) if (running && !blocked) begin
  IF_ir <= m[(squash ? target : IF_pc) >> 2];
  IF_pc <= (squash ? target : IF_pc) + 4;
end
```

ID: Instruction Decode stage

- Decodes the instruction
- Reads from register file `r[]`
- Computes `beq` comparison & `target`

```
// ID: Instruction Decode stage
always @(posedge clk) if (running && !ID_Bad) begin
  if (blocked) … else begin
    case (squashed `OP)
      `RTYPE: begin
        case (squashed `FUNCT)
          `ADDU:   begin RegDst=1; Branch=0; MemRead=0;
                   ALUOp=`ALUADD; MemWrite=0; ALUSrc=0;
                   RegWrite=1; Bad=0; end
    … endcase … endcase
    … ID_s <= s; ID_t <= t; … ID_ALUOp <= ALUOp; …
end end
```

EX: Execute stage

- Contains the ALU for integers & addresses

```
// EX: EXecute stage
always @(posedge clk) if (running) begin
  case (ID_ALUOp)
    `ALUAND: alu = ID_s & ID_src;
    `ALUOR:  alu = ID_s | ID_src;
    `ALUADD: alu = ID_s + ID_src;
    `ALUSUB: alu = ID_s - ID_src;
    `ALUSLT: alu = ID_s < ID_src;
    `ALUXOR: alu = ID_s ^ ID_src;
    default: alu = (ID_src << 16);
  endcase
  EX_alu <= alu;
  EX_t <= ID_t;
  EX_rd <= ID_rd;
  EX_MemRead <= ID_MemRead;
  EX_MemWrite <= ID_MemWrite;
  EX_Bad <= ID_Bad;
end
```

MEM: MEMory access stage

- Does a memory read or write
- Uses **`EX_MemRead`** for both read and mux **`v`**

```
// MEM: data MEMory access stage
always @(posedge clk) if (running) begin
  if (EX_MemRead) v = m[EX_alu >> 2]; else v = EX_alu;
  if (EX_MemWrite) m[EX_alu >> 2] <= EX_t;

  MEM_v <= v;
  MEM_rd <= EX_rd;
  MEM_Bad <= EX_Bad;
end
```

WB: Write Back stage

- Writes result into register…
  - Register $0 is read only
  - Not writing? Say we write to register 0…
- An instruction isn't done until it's here, so this is where halting really happens

```
// WB: register Write Back stage
always @(posedge clk) if (running) begin
  if (MEM_rd) r[MEM_rd] <= MEM_v;
  if (MEM_Bad) halt <= 1;
end
```

Running?

- Enables normal operation of stages
- Not running normally if:
  - Halted
  - There's a reset in progress; reset writes into stuff it doesn't own, so we need owners off

```
// Running state?
wire running;
assign running = ((!halt) && (!reset));
```

Squashed?

- For **beq**, we'll predict not taken, so **IF_pc+4**
- If we were right, no bubble
- If wrong, need to squash fetched instructions
  - No side-effects to undo yet
  - Just convert into **`NOP** to prevent future s-e

```
`define NOP `OR   // Null OPeration is or $0,$0,$0

// Squash instruction fetched on a mispredicted branch
wire `INST squashed;
assign squashed = (squash ? `NOP : IF_ir);
```

Data Dependence / Forwarding

- Need `$rs`, `$rt` unless either is `$0`
- Desired value could be in EX, MEM, or WB
- Forward gives priority to EX over MEM over WB
- Can't forward `lw` value from EX stage!

```
// Are we blocked by a dependence on rs or rt?
// Also handles all forwarding
…
assign canfwds = ID_needs && ((EX_deps && !MemRead) ||
                 ((!EX_deps) && (MEM_deps || WB_deps)));
assign canfwdt = ID_needt && ((EX_dept && !MemRead) ||
                 ((!EX_dept) && (MEM_dept || WB_dept)));
assign fwds = (EX_deps ? alu : (MEM_deps ? v : MEM_v));
assign fwdt = (EX_dept ? alu : (MEM_dept ? v : MEM_v));
assign dep = ((ID_needs && (EX_deps||MEM_deps||WB_deps)) ||
          (ID_needt && (EX_dept || MEM_dept || WB_dept)));
assign blocked = (dep && ((ID_needs && !canfwds) ||
                          (ID_needt && !canfwdt)));
```

Simulation Tracing

- Yet another (complex!) parallel-executing thing:

```
// State-by-state trace
`ifdef   TRACE
always @(posedge clk) if (running) begin
  …
  $display("IF ir=%x pc=%1d", IF_ir, IF_pc);
  case (IF_ir `OP)
    `RTYPE: begin
      case (IF_ir `FUNCT)
        `ADDU:  $display("IF  addu $%1d,$%1d,$%1d",
                  IF_ir `RD, IF_ir `RS, IF_ir `RT); …
  endcase endcase
  if (ID_Bad) $display("ID  illegal instruction");
    else $display("ID  s=%1d t=%1d src=%1d rd=%1d
      MemRead=%b ALUOp=%b MemWrite=%b", ID_s, ID_t,
      ID_src, ID_rd, ID_MemRead, ID_ALUOp, ID_MemWrite);
   …
end
`endif
endmodule
```

A Pipelined Verilog Version

```
// Testbench options
`define RUNTIME 200     // How long simulator can run
`define CLKDEL  1       // Clock transition delay
`define TRACE   1       // enable simulation trace

// Types
`define WORD    [31:0] // size of a data word
`define ADDR    [31:0] // size of a memory address
`define INST    [31:0] // size of an instruction
`define REG     [4:0]  // size of a register number
`define REGCNT  [31:0] // register count
`define MEMCNT  [511:0] // memory count inplemented
`define OPCODE  [5:0]  // 6-bit opcodes
`define ALUOP   [3:0]  // Simplified ALU codes

// Fields
`define OP      [31:26] // opcode field
`define RS      [25:21] // rs field
`define RT      [20:16] // rt field
`define RD      [15:11] // rd field
`define IMM     [15:0]  // immediate/offset field
`define SHAMT   [10:6]  // shift ammount
`define FUNCT   [5:0]   // function code (opcode extension)
`define JADDR   [25:0]  // jump address field
`define JPACK(R,D,J)        begin R`OP=0; R`JADDR=J; end
`define IPACK(R,O,S,T,I)    begin R`OP=O; R`RS=S; R`RT=T; R`IMM=I; end
`define RPACK(R,S,T,D,SH,FU)    begin R`OP=`RTYPE; R`RS=S; R`RT=T; R`RD=D; R`SHAMT=SH; R`FUNCT=FU; end
`define NOP     `OR     // Null OPeration is or $0,$0,$0

// Instruction encoding
`define RTYPE   6'h00   // OP field for all RTYPE instructions
`define BEQ     6'h04   // OP field
`define ADDIU   6'h09   // OP field
`define SLTIU   6'h0b   // OP field
`define ANDI    6'h0c   // OP field
`define ORI     6'h0d   // OP field
`define XORI    6'h0e   // OP field
`define LUI     6'h0f   // OP field
`define LW      6'h23   // OP field
`define SW      6'h2b   // OP field
`define ADDU    6'h21   // FUNCT field
`define SUBU    6'h23   // FUNCT field
`define AND     6'h24   // FUNCT field
`define OR      6'h25   // FUNCT field
`define XOR     6'h26   // FUNCT field
`define SLTU    6'h2b   // FUNCT field

// Simplified ALU codes, default to lui
`define ALUADD  4'b0000
`define ALUSUB  4'b0001
`define ALUADD  4'b0010
`define ALUSUB  4'b0110
`define ALUSLT  4'b0111
`define ALULUI  4'b1000
`define ALUXOR  4'b1111

// Generic multi-cycle processor
module processor(halt, reset, clk);
output reg halt;
input reset, clk;
reg `WORD `MEMCNT;
reg `WORD r `REGCNT;

// Initialize register file and memory
initial begin
    r[1] = 22; r[2] = 1; r[3] = 42;
    r[4] = 601; r[5] = 13813; r[0] = -1;
    `RPACK(m[0], 2, 3, 1, 0, `ADDU)
    `RPACK(m[1], 2, 3, 1, 0, `SLTU)
    `RPACK(m[2], 3, 5, 1, 0, `AND)
    `RPACK(m[3], 5, 3, 1, 0, `OR)
    `RPACK(m[4], 3, 5, 1, 0, `XOR)
    `RPACK(m[5], 4, 3, 1, 0, `SUBU)
    `IPACK(m[6], `ADDIU, 3, 1, -1)
    `IPACK(m[7], `SLTIU, 5, 1, 12345)
    `IPACK(m[8], `ANDI, 3, 1, 3)
    `IPACK(m[9], `ORI, 3, 1, 3)
    `IPACK(m[10], `XORI, 3, 1, 3)
    `IPACK(m[11], `LUI, 0, 1, 1)
    `IPACK(m[12], `LW, 2, 1, 1023)
    `IPACK(m[13], `SW, 0, 2, 1024)
    `IPACK(m[14], `ADDIU, 6, 6, 1)
    `IPACK(m[15], `BEQ, 0, 6, -2)
    m[16] = 0;
    m[250] = 22;
end

// IF registers
reg `ADDR IF_pc;
reg `INST IF_ir;

// ID registers
reg RegDst, Branch, MemRead, MemWrite, ALUSrc, RegWrite, Bad, ID_MemRead, ID_MemWrite, ID_Bad;
reg `ALUOP ALUOp, ID_ALUOp;
reg `WORD s, t, imm, ID_s, ID_t, ID_src;
reg `REG ID_rd;
reg `ADDR target;
reg squash;

// EX registers
reg EX_MemRead, EX_MemWrite, EX_Bad;
reg `WORD alu, EX_alu, EX_t;
reg `REG EX_rd;

// MEM registers
reg MEM_Bad;
reg `WORD v, MEM_v;
reg `REG MEM_rd;
```

```
// Running state?
wire running;
assign running = (!!halt) && (!reset);

// Squash instruction fetched on a mispredicted branch
wire `INST squashed;
assign squashed = squash ? `NOP : IF_ir;

// Are we blocked by a dependence on rs or rt?
// Also handles all forwarding
wire EX_depx, EX_dept, MEM_depx, MEM_dept, WB_depx, WB_dept, canfwdx, canfwdt, dep, blocked;
wire `WORD fwdx, fwdt;
assign ID_needs = (IF_ir `RS != 0);
assign ID_needt = (IF_ir `RT != 0);
assign EX_depx = (IF_ir `RS == EX_rd);
assign EX_dept = (IF_ir `RT == EX_rd);
assign MEM_depx = (IF_ir `RS == MEM_rd);
assign MEM_dept = (IF_ir `RT == MEM_rd);
assign WB_depx = (IF_ir `RS == WB_rd);
assign WB_dept = (IF_ir `RT == WB_rd);
assign canfwdx = ID_needs && ((EX_depx && !MemRead) || ((!EX_depx) && (MEM_depx || WB_depx)));
assign canfwdt = ID_needt && ((EX_dept && !MemRead) || ((!EX_dept) && (MEM_dept || WB_dept)));
assign fwdx = (EX_depx ? alu : (MEM_depx ? v : MEM_v));
assign fwdt = (EX_dept ? alu : (MEM_dept ? v : MEM_v));
assign dep = ((ID_needs && (EX_depx || MEM_depx || WB_depx)) ||
              (ID_needt && (EX_dept || MEM_dept || WB_dept)));
assign blocked = (dep && ((ID_needs && !canfwdx) || (ID_needt && !canfwdt)));

// IF: Instruction Fetch stage
always @(posedge clk) if (running && !blocked) begin
  IF_ir <= m[squash ? target : IF_pc >> 2];
  IF_pc <= (squash ? target : IF_pc) + 4;
end

// ID: Instruction Decode stage
always @(posedge clk) if (running && !ID_Bad) begin
  if (blocked) begin
    ID_s <= 0;
    ID_t <= 0;
    ID_src <= 0;
    ID_rd <= 0; // not writing
    ID_MemRead <= 0; // not reading
    ID_ALUOp <= 0;
    ID_MemWrite <= 0; // not storing
    ID_Bad <= 0;
  end else begin
    // Brute force decode of instruction
    case (squashed `OP)
      `RTYPE: begin
        case (squashed `FUNCT)
          `ADDU:   begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUADD; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `SUBU:   begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUSUB; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `AND:    begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUAND; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `OR:     begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `XOR:    begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUXOR; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          `SLTU:   begin RegDst=1; Branch=0; MemRead=0; ALUOp=`ALUSLT; MemWrite=0; ALUSrc=0; RegWrite=1; Bad=0; end
          default: begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=0; RegWrite=0; Bad=1; end
        endcase
      end
      `BEQ:    begin RegDst=0; Branch=1; MemRead=0; ALUOp=`ALUSUB; MemWrite=0; ALUSrc=0; RegWrite=0; Bad=0; end
      `ADDIU:  begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUADD; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `SLTIU:  begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUSLT; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `ANDI:   begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUAND; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `ORI:    begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `XORI:   begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUXOR; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `LUI:    begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALULUI; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `LW:     begin RegDst=0; Branch=0; MemRead=1; ALUOp=`ALUADD; MemWrite=0; ALUSrc=1; RegWrite=1; Bad=0; end
      `SW:     begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUADD; MemWrite=1; ALUSrc=1; RegWrite=0; Bad=0; end
      default: begin RegDst=0; Branch=0; MemRead=0; ALUOp=`ALUOR;  MemWrite=0; ALUSrc=0; RegWrite=0; Bad=1; end
    endcase

    s = (canfwdx ? fwdx : r[squashed `RS]);
    t = (canfwdt ? fwdt : r[squashed `RT]);
    imm = {{16{squashed[15]}}, squashed `IMM};
    target <= IF_pc + {imm[29:0], 2'b00};

    squash <= (Branch && (s == t));
    ID_s <= s;
    ID_t <= t;
    ID_src <= (ALUSrc ? imm : t);
    ID_rd <= (RegWrite ? (RegDst ? squashed `RD : squashed `RT) : 0); // $0 if not writing
    ID_MemRead <= MemRead;
    ID_ALUOp <= ALUOp;
    ID_MemWrite <= MemWrite;
    ID_Bad <= Bad;
  end
end

// EX: EXecute stage
always @(posedge clk) if (running) begin
  case (ID_ALUOp)
    `ALUADD: alu = ID_s & ID_src;
    `ALUOR:  alu = ID_s | ID_src;
    `ALUADD: alu = ID_s + ID_src;
    `ALUSUB: alu = ID_s - ID_src;
    `ALUSLT: alu = ID_s < ID_src;
    `ALUXOR: alu = ID_s ^ ID_src;
    default: alu = {ID_src <<< 16};
  endcase

  EX_alu <= alu;
  EX_t <= ID_t;
  EX_rd <= ID_rd;
  EX_MemRead <= ID_MemRead;
  EX_MemWrite <= ID_MemWrite;
  EX_Bad <= ID_Bad;
end
```

Color key: initialization
IF squashed ID blocked EX MEM WB debugging

```
// MEM: data MEMory access stage
always @(posedge clk) if (running) begin
  if (EX_MemRead) v = m[EX_alu >> 2]; else v = EX_alu;
  if (EX_MemWrite) m[EX_alu >> 2] <= EX_t;

  MEM_v <= v;
  MEM_rd <= EX_rd;
  MEM_Bad <= EX_Bad;
end

// WB: register Write Back stage
always @(posedge clk) if (running) begin
  if (MEM_rd) r[MEM_rd] <= MEM_v;
  if (MEM_Bad) halt <= 1;
end

// Reset
always @(posedge clk) if (reset) begin
  halt<=0;
  r[0]<=0;
  squash <= 0; target <= 0;
  IF_ir<= `NOP; IF_pc<=0;
  ID_s<=0; ID_t<=0; ID_src<=0; ID_rd<=0; ID_MemRead<=0; ID_ALUOp<=0; ID_MemWrite<=0; ID_Bad<=0;
  EX_alu<=0; EX_t<=0; EX_rd<=0; EX_MemRead<=0; EX_MemWrite<=0; EX_Bad<=0;
  MEM_v<=0; MEM_rd<=0; MEM_Bad<=0;
end

`ifdef TRACE
// State-by-state trace
always @(posedge clk) if (running) begin
  `define Fs      (canfwdx ? (EX_depx ? "EX " : (MEM_depx ? "MEM " : "WB ")) : "")
  `define Ft      (canfwdt ? (EX_dept ? "EX " : (MEM_dept ? "MEM " : "WB ")) : "")
  $display("\n  running=%1d blocked=%1d s=%x$%1d t=%x$%1d squash=%1d", running, blocked, `Fs, IF_ir `RS, `Ft, IF_ir `RT, squash);
  $display("IF  ir=%x pc=%1d", IF_ir, IF_pc);
  case (IF_ir `OP)
    `RTYPE: begin
      case (IF_ir `FUNCT)
        `ADDU:   $display("IF  addu $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `SUBU:   $display("IF  subu $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `AND:    $display("IF  and $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `OR:     $display("IF  or $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `XOR:    $display("IF  xor $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        `SLTU:   $display("IF  sltu $%1d,$%1d,$%1d", IF_ir `RD, IF_ir `RS, IF_ir `RT);
        default: $display("IF  OP=%x RS=%1d RT=%1d RD=%1d SHAMT=%1d FUNCT=%x", IF_ir `OP, IF_ir `RS, IF_ir `RT, IF_ir `RD, IF_ir `SHAMT, IF_ir `FUNCT);
      endcase
    end
    `BEQ:    $display("IF  beq $%1d,$%1d,offset=$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `ADDIU:  $display("IF  addiu $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `SLTIU:  $display("IF  sltiu $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `ANDI:   $display("IF  andi $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `ORI:    $display("IF  ori $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `XORI:   $display("IF  xori $%1d,$%1d,$%1d", IF_ir `RT, IF_ir `RS, IF_ir `IMM);
    `LUI:    $display("IF  lui $%1d,$%1d", IF_ir `RS, IF_ir `IMM);
    `LW:     $display("IF  lw $%1d,$%1d($%1d)", IF_ir `RT, IF_ir `IMM, IF_ir `RS);
    `SW:     $display("IF  sw $%1d,$%1d($%1d)", IF_ir `RT, IF_ir `IMM, IF_ir `RS);
    default: $display("IF  OP=%x RS=%1d RT=%1d IMM=%x", IF_ir `OP, IF_ir `RS, IF_ir `RT, IF_ir `IMM);
  endcase
  if (ID_Bad) $display("ID  illegal instruction");
  else $display("ID  s=%1d t=%1d src=%1d rd=%1d MemRead=%b ALUOp=%b MemWrite=%b", ID_s, ID_t, ID_src, ID_rd, ID_MemRead, ID_ALUOp, ID_MemWrite);
  if (EX_Bad) $display("EX  illegal instruction");
  else $display("EX  alu=%1d t=%1d rd=%1d MemRead=%b MemWrite=%b", EX_alu, EX_t, EX_rd, EX_MemRead, EX_MemWrite);
  if (MEM_Bad) $display("MEM illegal instruction");
  else begin
    if (EX_MemWrite) $display("MEM m[%1d]=%1d v=%1d rd=%1d", EX_alu, EX_t, MEM_v, MEM_rd);
    else $display("MEM v=%1d rd=%1d", MEM_v, MEM_rd);
  end
  if (MEM_rd) $display("WB  r[%1d]=%1d", MEM_rd, MEM_v);
end
`endif
endmodule

// Testbench
module bench;
wire halt;
reg reset = 1;
reg clk = 0;

processor PE(halt, reset, clk);

initial begin
  # `CLKDEL clk = 1;
  # `CLKDEL clk = 0;
  reset = 0;
  while ($time < `RUNTIME) && !halt) begin
    # `CLKDEL clk = 1;
    # `CLKDEL clk = 0;
  end
end
endmodule
```

Three Verilog Implementations

- **Multi-cycle** MIPS, multiple CPI:

  `http://aggregate.org/EE380/multiv.html`

- **Single-cycle** MIPS, 1 CPI, but slow clock:

  `http://aggregate.org/EE380/onebeq.html`

- **Pipelined** MIPS, fast clock, ~1 CPI throughput:

  `http://aggregate.org/EE380/pipe.html`

Out-Of-Order Implementation?

- *Beyond the scope of this class*, but it's dataflow
- CDC6600 Scoreboard
  - Need both operands in registers
  - Instruction waits @ function unit
- Tomasulo Scheduling
  - Copies operands when *each one* is available
  - Instruction waits @ reservation station
  - Common data bus broadcasts each result
- Now rename registers, no serial bottleneck