



Microprogramming

Krste Asanovic
Laboratory for Computer Science
M.I.T.

<http://www.csg.lcs.mit.edu/6.823>



The DLX ISA

Processor State

32 32-bit GPRs, R0 always contains a 0
32 single precision FPRs, may also be viewed as
16 double precision FPRs
FP status register, used for FP compares & exceptions
PC, the program counter
some other special registers

Data types

8-bit byte, 2-byte half word
32-bit word for integers
32-bit word for single precision floating point
64-bit word for double precision floating point

Load/Store style instruction set

data addressing modes- immediate & indexed
branch addressing modes- PC relative & register indirect
Byte addressable memory- big-endian mode


All instructions are 32 bits

(See Chapter 2, H&P for full description)



Microarchitecture

Implementation of an ISA



Structure: How components are connected.

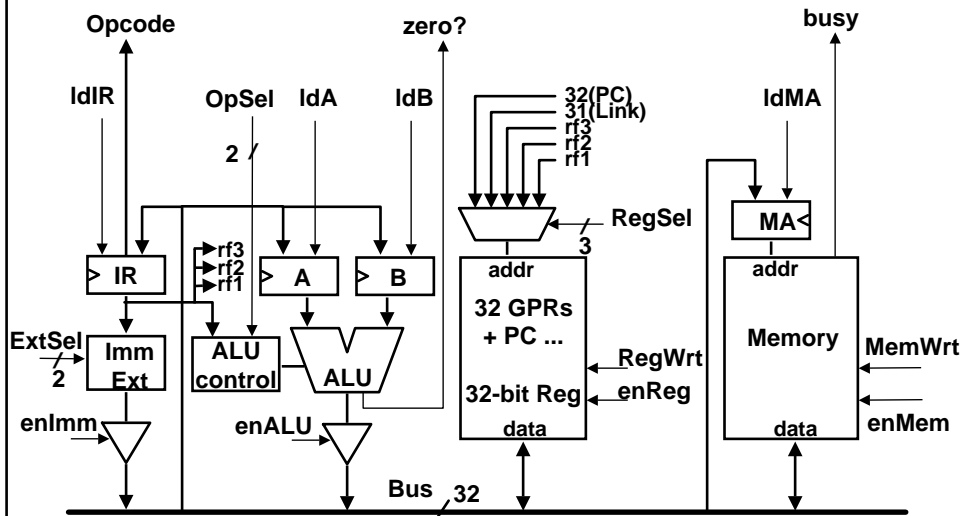
Static

Sequencing: How data moves between components

Dynamic



A Bus-based Datapath for DLX



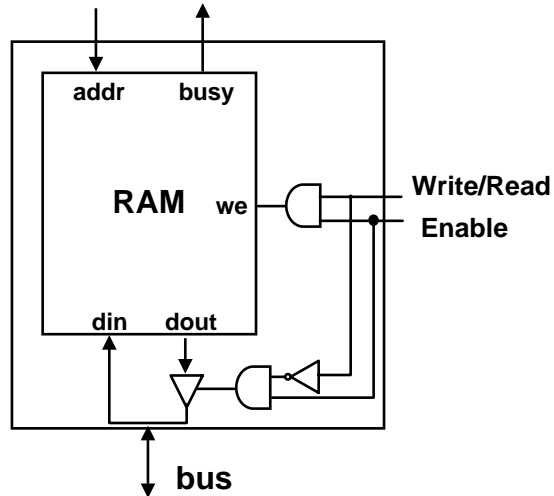
Microinstruction: register to register transfer (17 control signals)

MA ← PC means RegSel = PC; enReg=yes; IdMA= yes

B ← Reg[rf2] means RegSel = rf2; enReg=yes; IdB = yes



Memory Module



We will assume that Memory operates asynchronously and is slow as compared to Reg-to-Reg transfers



Instruction Execution

Execution of a DLX instruction involves

1. instruction fetch
2. decode and register fetch
3. ALU operation
4. memory operation (optional)
5. write back to register file (optional)

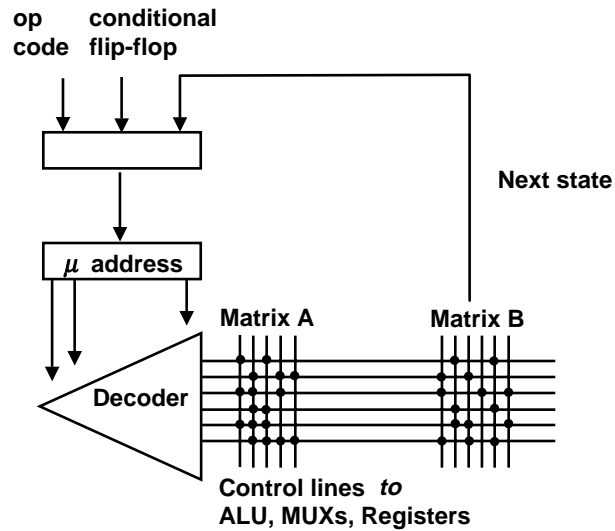
and the computation of the address of the *next instruction*




Microcontrol Unit

Maurice Wilkes, 1954

Embed the control logic state table in a memory array




DLX ALU Instructions



Register-Register form:

$$\text{Reg}[\text{rf3}] \leftarrow \text{function}(\text{Reg}[\text{rf1}], \text{Reg}[\text{rf2}])$$



Register-Immediate form:

$$\text{Reg}[\text{rf2}] \leftarrow \text{function}(\text{Reg}[\text{rf1}], \text{SignExt}(\text{immediate}))$$



Microprogram Fragments

instr fetch: $MA \leftarrow PC$
 $IR \leftarrow \text{Memory}$
 $A \leftarrow PC$
 $PC \leftarrow A + 4$
 dispatch on OPcode


can be treated as a macro

ALU: $A \leftarrow \text{Reg}[\text{rf1}]$
 $B \leftarrow \text{Reg}[\text{rf2}]$
 $\text{Reg}[\text{rf3}] \leftarrow \text{func}(A,B)$
do instruction fetch

ALUi: $A \leftarrow \text{Reg}[\text{rf1}]$
 $B \leftarrow \text{Imm}$ *sign extention ...*
 $\text{Reg}[\text{rf2}] \leftarrow \text{Opcode}(A,B)$
do instruction fetch



DLX Load/Store Instructions



Load/store byte, halfword, word to/from GPR:


LB, LBU, SB, LH, LHU, SH, LW, SW

byte and half-word can be sign or zero extended

Load/store single and double FP to/from FPR:

LF, LD, SF, SD

- Byte addressable machine
- Memory access must be data aligned
- A single addressing mode
 (base) + displacement
- Big-endian byte ordering






DLX Control Instructions

Conditional branch on GPR



Unconditional register-indirect jumps



Unconditional PC-relative jumps



- PC-offset are specified in bytes
- jump-&-link (JAL) stores PC+4 into the link register (R31)
- (*Real DLX has delayed branches – ignored this lecture*)



Microprogram Fragments (*cont.*)

LW: $A \leftarrow \text{Reg}[\text{rf1}]$
 $B \leftarrow \text{Imm}$
 $\text{MA} \leftarrow A + B$
 $\text{Reg}[\text{rf2}] \leftarrow \text{Memory}$
 do instruction fetch

J: $A \leftarrow \text{PC}$
 $B \leftarrow \text{Imm}$
 $\text{PC} \leftarrow A + B$
 do instruction fetch

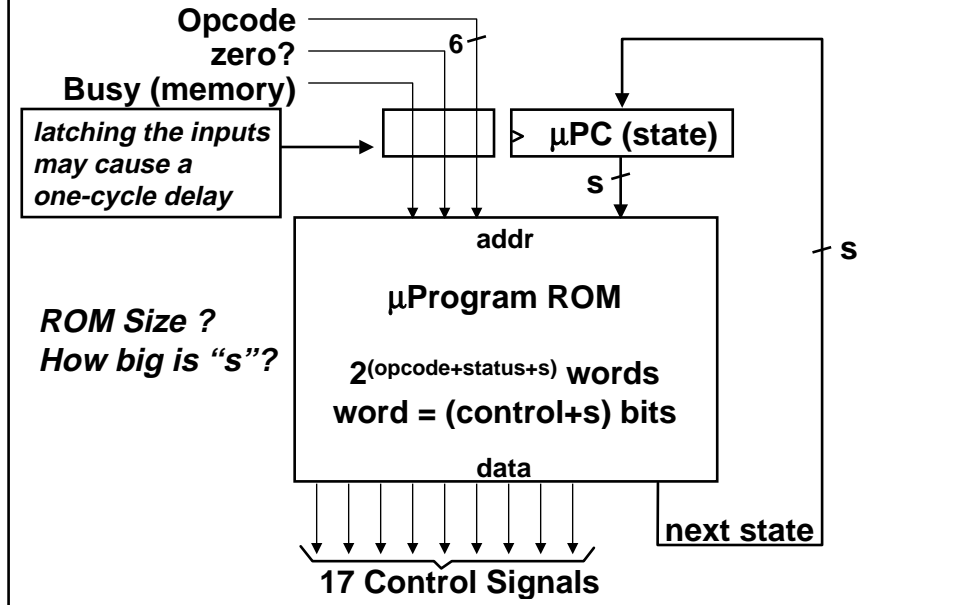
beqz: $A \leftarrow \text{Reg}[\text{rf1}]$
 If zero?(A) then go to bz-taken
 do instruction fetch

bz-taken: $A \leftarrow \text{PC}$
 $B \leftarrow \text{Imm}$
 $\text{PC} \leftarrow A + B$
 do instruction fetch



DLX Microcontroller: *first attempt*

Krste Asanovic
February 20, 2001
6.823, L4-13



Microprogram in the ROM *worksheet*

Krste Asanovic
February 20, 2001
6.823, L4-14

State	Op	zero?	busy	Control points	next-state
fetch ₀	*	*	*	MA ← PC	fetch ₁
fetch ₁	*	*	yes	fetch ₁
fetch ₁	*	*	no	IR ← Memory	fetch ₂
fetch ₂	*	*	*	A ← PC	fetch ₃
fetch ₃	*	*	*	PC ← A + 4	?
ALU ₀	*	*	*	A ← Reg[rf1]	ALU ₁
ALU ₁	*	*	*	B ← Reg[rf2]	ALU ₂
ALU ₂	*	*	*	Reg[rf3] ← func(A,B)	fetch ₀



Microprogram in the ROM

State	Op	zero?	busy	Control points	next-state
fetch ₀	*	*	*	MA ← PC	fetch ₁
fetch ₁	*	*	yes	fetch ₁
fetch ₁	*	*	no	IR ← Memory	fetch ₂
fetch ₂	*	*	*	A ← PC	fetch ₃
fetch ₃	ALU	*	*	PC ← A + 4	ALU ₀
fetch ₃	ALUi	*	*	PC ← A + 4	ALUi ₀
fetch ₃	LW	*	*	PC ← A + 4	LW ₀
fetch ₃	SW	*	*	PC ← A + 4	SW ₀
fetch ₃	J	*	*	PC ← A + 4	J ₀
fetch ₃	JAL	*	*	PC ← A + 4	JAL ₀
fetch ₃	JR	*	*	PC ← A + 4	JR ₀
fetch ₃	JALR	*	*	PC ← A + 4	JALR ₀
fetch ₃	beqz	*	*	PC ← A + 4	beqz ₀
...					
ALU ₀	*	*	*	A ← Reg[rf1]	ALU ₁
ALU ₁	*	*	*	B ← Reg[rf2]	ALU ₂
ALU ₂	*	*	*	Reg[rf3] ← func(A,B)	fetch ₀

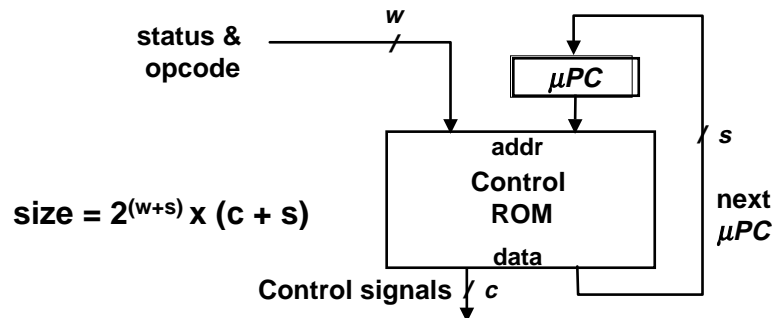


Microprogram in the ROM *Cont.*

State	Op	zero?	busy	Control points	next-state
ALUi ₀	*	*	*	A ← Reg[rf1]	ALUi ₁
ALUi ₁	sExt	*	*	B ← sExt ₁₆ (Imm)	ALUi ₂
ALUi ₁	uExt	*	*	B ← uExt ₁₆ (Imm)	ALUi ₂
ALUi ₂	*	*	*	Reg[rf3] ← Op(A,B)	fetch ₀
...					
J ₀	*	*	*	A ← PC	J ₁
J ₁	*	*	*	B ← sExt ₂₆ (Imm)	J ₂
J ₂	*	*	*	PC ← A+B	fetch ₀
...					
beqz ₀	*	*	*	A ← Reg[rf1]	beqz ₁
beqz ₁	*	yes	*	A ← PC	beqz ₂
beqz ₁	*	no	*	fetch ₀
beqz ₂	*	*	*	B ← sExt ₁₆ (Imm)	beqz ₃
beqz ₃	*	*	*	PC ← A+B	fetch ₀
...					



Size of Control Store



$$\text{size} = 2^{(w+s)} \times (c + s)$$

DLX

$$w = 6+2$$

$$c = 17$$

$$s = ?$$

no. of steps per opcode = 4 to 6 + fetch-sequence

no. of states \approx (4 steps per op-group) \times op-groups
+ common sequences

$$= 4 \times 8 + 10 \text{ states} = 42 \text{ states}$$

$$\Rightarrow s = 6$$

$$\text{Control ROM} = 2^{(8+6)} \times 23 \text{ bits} \approx 48 \text{ Kbytes}$$



Reducing the Size of Control Store

Control store has to be *fast* \Rightarrow *expensive*

Reduce the ROM height (= address bits)

\Rightarrow *reduce inputs by extra external logic*
each input bit doubles the size of the control store

\Rightarrow *reduce states by grouping opcodes*
find common sequences of actions

\Rightarrow *condense input status bits*
combine all exceptions into one, i.e.,
exception/no-exception

Reduce the ROM width

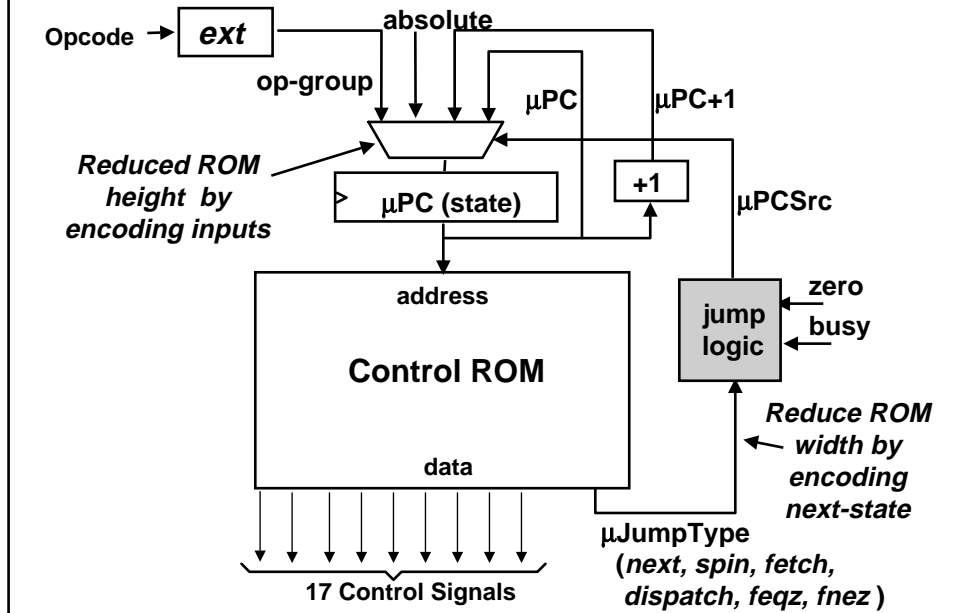
\Rightarrow *restrict the next-state encoding*

Next, Dispatch on opcode, Wait for memory, ...

\Rightarrow *encode control signals (vertical microcode)*



DLX Microcontroller-2



Jump Logic

$\mu\text{PCSrc} = \text{Case } \mu\text{JumpTypes}$

- next** \Rightarrow $\mu\text{PC}+1$
- spin** \Rightarrow if (busy) then μPC else $\mu\text{PC}+1$
- fetch** \Rightarrow absolute
- dispatch** \Rightarrow op-group
- feqz** \Rightarrow if (zero) then absolute else $\mu\text{PC}+1$
- fnez** \Rightarrow if (zero) then $\mu\text{PC}+1$ else absolute



DLX-Controller-2 worksheet

Krste Asanovic
February 20, 2001
6.823, L4-21

State	Control points	next-state
fetch ₀	$MA \leftarrow PC$	
fetch ₁	$IR \leftarrow \text{Memory}$	
fetch ₂	$A \leftarrow PC$	
fetch ₃	$PC \leftarrow A + 4$	
...		
ALU ₀	$A \leftarrow \text{Reg}[\text{rf1}]$	
ALU ₁	$B \leftarrow \text{Reg}[\text{rf2}]$	
ALU ₂	$\text{Reg}[\text{rf3}] \leftarrow \text{func}(A,B)$	
...		
J ₀	$A \leftarrow PC$	
J ₁	$B \leftarrow \text{sExt}_{26}(\text{Imm})$	
J ₂	$PC \leftarrow A+B$	
...		
BEQZ ₀	$A \leftarrow \text{Reg}[\text{rf1}]$	
BEQZ ₁		
BEQZ ₂	$A \leftarrow PC$	
BEQZ ₃	$B \leftarrow \text{sExt}_{16}(\text{Imm})$	
BEQZ ₄	$PC \leftarrow A+B$	



Instruction Fetch & ALU: DLX-Controller-2

Krste Asanovic
February 20, 2001
6.823, L4-22

State	Control points	next-state
fetch ₀	$MA \leftarrow PC$	next
fetch ₁	$IR \leftarrow \text{Memory}$	spin
fetch ₂	$A \leftarrow PC$	next
fetch ₃	$PC \leftarrow A + 4$	dispatch
...		
ALU ₀	$A \leftarrow \text{Reg}[\text{rf1}]$	next
ALU ₁	$B \leftarrow \text{Reg}[\text{rf2}]$	next
ALU ₂	$\text{Reg}[\text{rf3}] \leftarrow \text{func}(A,B)$	fetch
...		
ALUi ₀	$A \leftarrow \text{Reg}[\text{rf1}]$	next
ALUi ₁	$B \leftarrow \text{sExt}_{16}(\text{Imm})$	next
ALUi ₂	$\text{Reg}[\text{rf3}] \leftarrow \text{Op}(A,B)$	fetch



Load & Store: *DLX-Controller-2*

State	Control points	next-state
LW ₀	A ← Reg[rf1]	next
LW ₁	B ← sExt ₁₆ (Imm)	next
LW ₂	MA ← A+B	next
LW ₃	Reg[rf2] ← Memory	spin
LW ₄		fetch
SW ₀	A ← Reg[rf1]	next
SW ₁	B ← sExt ₁₆ (Imm)	next
SW ₂	MA ← A+B	next
SW ₃	Memory ← Reg[rf2]	spin
SW ₄		fetch



Branches: *DLX-Controller-2*

State	Control points	next-state
BEQZ ₀	A ← Reg[rf1]	next
BEQZ ₁		fnez
BEQZ ₂	A ← PC	next
BEQZ ₃	B ← sExt ₁₆ (Imm)	next
BEQZ ₄	PC ← A+B	fetch
BNEZ ₀	A ← Reg[rf1]	next
BNEZ ₁		feqz
BNEZ ₂	A ← PC	next
BNEZ ₃	B ← sExt ₁₆ (Imm)	next
BNEZ ₄	PC ← A+B	fetch



Jumps: *DLX-Controller-2*

State	Control points	next-state
J_0	$A \leftarrow PC$	next
J_1	$B \leftarrow sExt_{26}(Imm)$	next
J_2	$PC \leftarrow A+B$	fetch
JR_0	$PC \leftarrow Reg[rf1]$	fetch
JAL_0	$Reg[31] \leftarrow PC$	next
JAL_1	$A \leftarrow PC$	next
JAL_2	$B \leftarrow sExt_{26}(Imm)$	next
JAL_3	$PC \leftarrow A+B$	fetch
$JALR_0$	$Reg[31] \leftarrow PC$	next
$JALR_1$	$PC \leftarrow Reg[rf1]$	fetch



Nanocoding

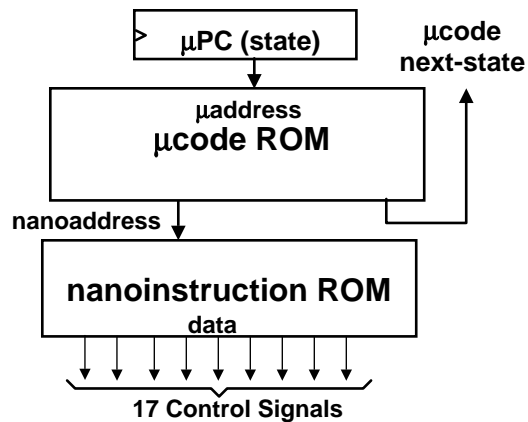
Exploits recurring control signal patterns in μ code, e.g.,

$ALU_0 \ A \leftarrow Reg[rf1]$

...

$ALUi_0 \ A \leftarrow Reg[rf1]$

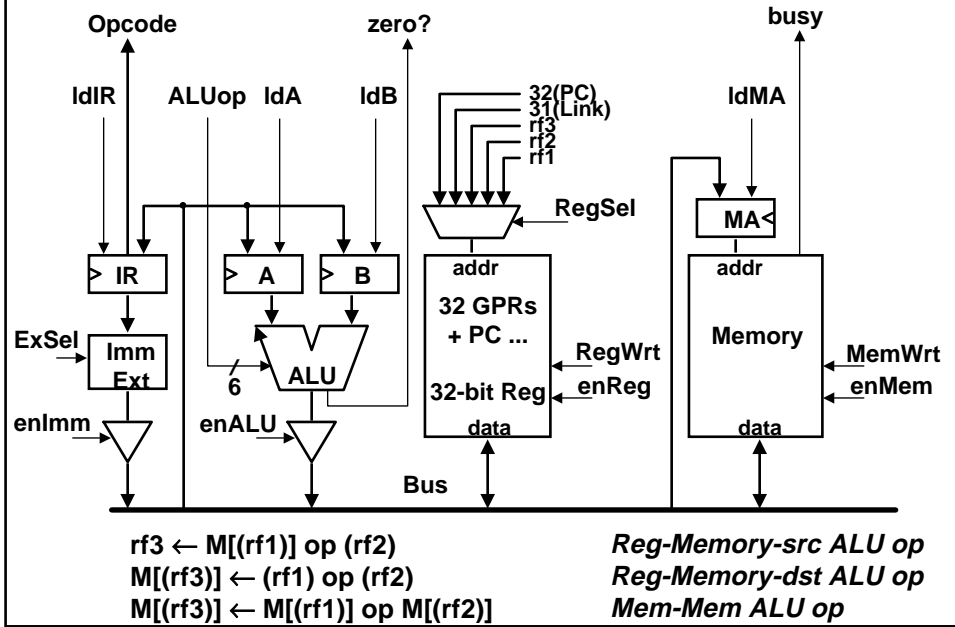
...



- Place common control signal patterns into *nanoinstruction* table
- μ code word contains pointer to nanoinstruction
- Higher latency to get control signals out



Implementing Complex Instructions



Mem-Mem ALU Instructions: DLX-Controller-2

Mem-Mem ALU op $M[(rf3)] \leftarrow M[(rf1)] \text{ op } M[(rf2)]$

ALUMM ₀	MA ← Reg[rf1]	next
ALUMM ₁	A ← Memory	spin
ALUMM ₂	MA ← Reg[rf2]	next
ALUMM ₃	B ← Memory	spin
ALUMM ₄	MA ← Reg[rf3]	next
ALUMM ₅	Memory ← func(A,B)	spin
ALUMM ₆		fetch

Complex instructions usually do not require datapath modifications in a microprogrammed implementation -- only extra space for the control program

Implementing these instructions using a hardwired controller is difficult without datapath modifications



Microprogramming in the Seventies

Thrived because

- Significantly faster ROMs than DRAMs were available
- For complex instruction sets, datapath and controller were *cheaper and simpler*
- *New instructions*, e.g., floating point, could be supported without datapath modifications
- *Fixing bugs* in the controller was easier
- ISA compatibility across various models could be achieved cheaply



Writable Control Store (WCS)

Implement control store with SRAM not ROM

- MOS SRAM memories now almost as fast as control store (core memories/DRAMs were 10x slower)
- Bug-free microprograms difficult to write

WCS provided as option on several minicomputers

- Allowed users to change microcode for each process

WCS *failed*

- Little or no programming tools support
- Hard to fit software into small space
- Microcode control tailored to original ISA, less useful for others
- Large WCS part of processor state - expensive context switches
- Protection difficult if user can change microcode
- Virtual memory required *restartable* microcode



Performance Issues

Microprogrammed control
⇒ multiple cycles per instruction

Cycle time ?

$$t_C > \max(t_{\text{reg-reg}}, t_{\text{ALU}}, t_{\mu\text{ROM}}, t_{\text{RAM}})$$

Given complex control, t_{ALU} & t_{RAM} can be broken into multiple cycles. However, $t_{\mu\text{ROM}}$ cannot be broken down. Hence

$$t_C > \max(t_{\text{reg-reg}}, t_{\mu\text{ROM}})$$

Suppose $10 * t_{\mu\text{ROM}} < t_{\text{RAM}}$
*good performance, relative to the single-cycle
hardwired implementation, can be achieved
even with a CPI of 10*



VLSI & Microprogramming

By late seventies

- technology assumption about ROM & RAM speed became invalid
- micromachines became more complicated
 - to overcome slower ROM, micromachines were pipelined
 - complex instruction sets led to the need for subroutine and call stacks in μ code.
 - need for fixing bugs in control programs was in conflict with read-only nature of μ ROM
⇒ *WCS (B1700, QMachine, Intel432, ...)*
- introduction of caches and buffers, especially for instructions, made multiple-cycle execution of reg-reg instructions unattractive



Modern Usage

Microprogramming is far from extinct

Played a crucial role in micros of the Eighties,
Motorola 68K (first microcoded microprocessor)
Intel 386 and 486

Microcode is present in most modern CISC micros in an assisting role (*e.g. AMD Athlon, Intel Pentium-4*)

- Most instructions are executed directly, i.e., with hard-wired control
- Infrequently-used and/or complicated instructions invoke the microcode engine

Patchable microcode common for post-fabrication bug fixes, e.g. Intel Pentiums load μ code patches at bootup