


Chapter Six

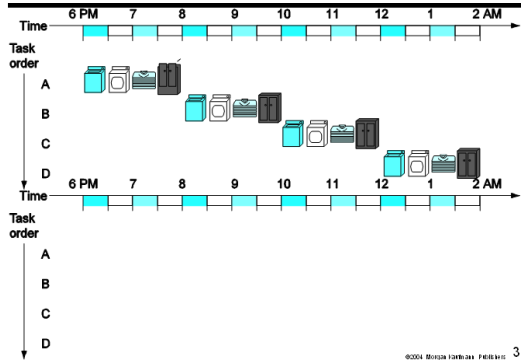
©2004 Morgan Kaufmann Publishers 1

Motivação



©2004 Morgan Kaufmann Publishers 2

Motivação: Você pode fazer melhor?




©2004 Morgan Kaufmann Publishers 3

Motivação: Quais são as restrições?

©2004 Morgan Kaufmann Publishers 4

Motivação



©2004 Morgan Kaufmann Publishers 5

Você consegue extrapolar para um processador?

- Período do clock de 200ps
- Esqueçam, temporariamente, as dependências entre instruções
- Quantas fases seriam?
- Quais seriam as fases?
 - Dica: lembre-se do controle do datapath multiciclo
 - Todas as fases gastam o mesmo tempo?

©2004 Morgan Kaufmann Publishers 6



Monte um diagrama da execução das instruções

- lw \$1, 100(\$2)
- lw \$3, 100(\$4)
- lw \$5, 100(\$6)
- lw \$7, 100(\$8)
- lw \$9, 100(\$10)

Tempo

Pipelining

- Improve performance by increasing instruction throughput

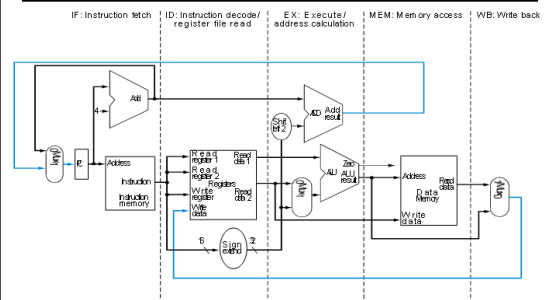


Ideal speedup is number of stages in the pipeline. Do we achieve this?

Pipelining


- What makes it easy
 - all instructions are the same length
 - just a few instruction formats
 - memory operands appear only in loads and stores
- What makes it hard?
 - structural hazards: suppose we had only one memory
 - control hazards: need to worry about branch instructions
 - data hazards: an instruction depends on a previous instruction
- We'll build a simple pipeline and look at these issues
- We'll talk about modern processors and what really makes it hard:
 - exception handling
 - trying to improve performance with out-of-order execution, etc.

Basic Idea




- What do we need to add to actually split the datapath into stages?

Pipelined Datapath




Can you find a problem even if there are no dependencies?
 What instructions can we execute to manifest the problem?

Corrected Datapath



Graphically Representing Pipelines

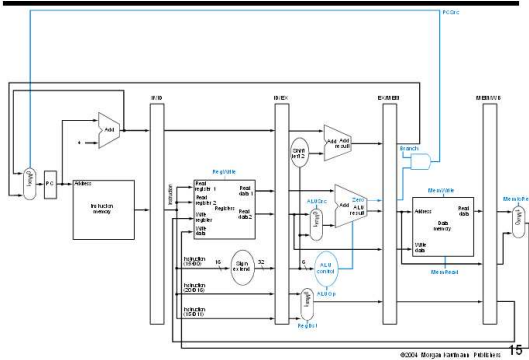


- Can help with answering questions like:
 - how many cycles does it take to execute this code?
 - what is the ALU doing during cycle 4?
 - use this representation to help understand datapaths

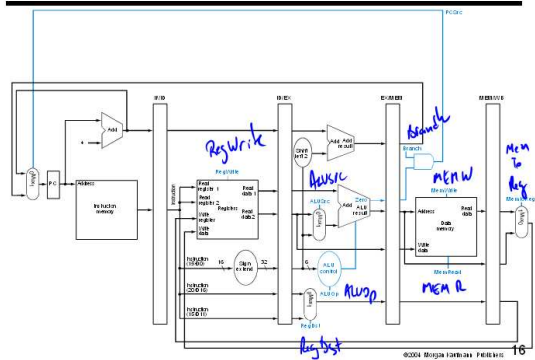
Pipeline control

- We have 5 stages. What needs to be controlled in each stage?
 - Instruction Fetch and PC Increment
 - Instruction Decode / Register Fetch
 - Execution
 - Memory Stage
 - Write Back
- How would control be handled in an automobile plant?
 - a fancy control center telling everyone what to do?
 - should we use a finite state machine?

Pipeline Control




Como você geraria os sinais de controle?



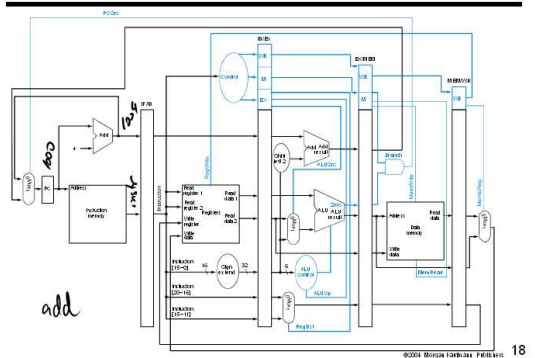
Pipeline Control

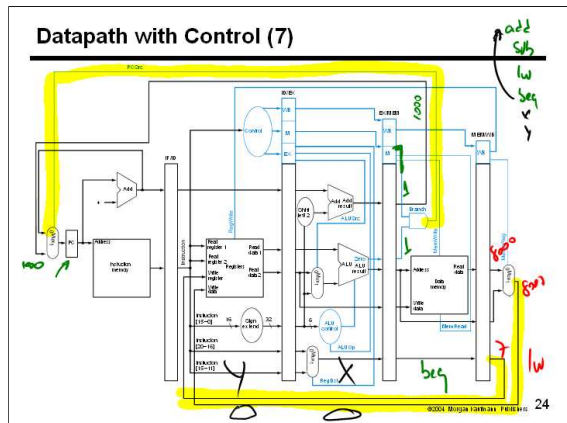
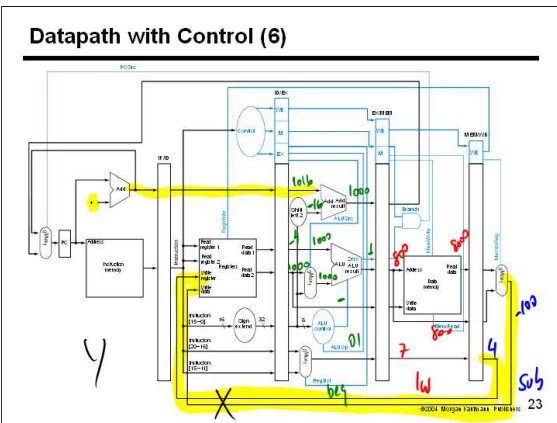
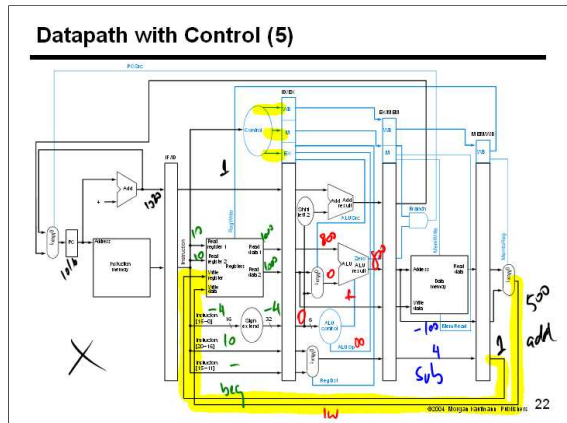
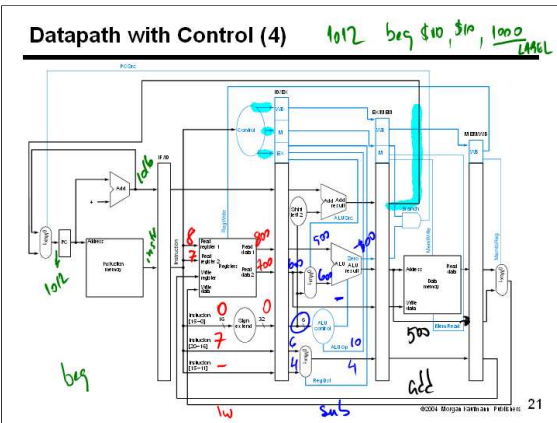
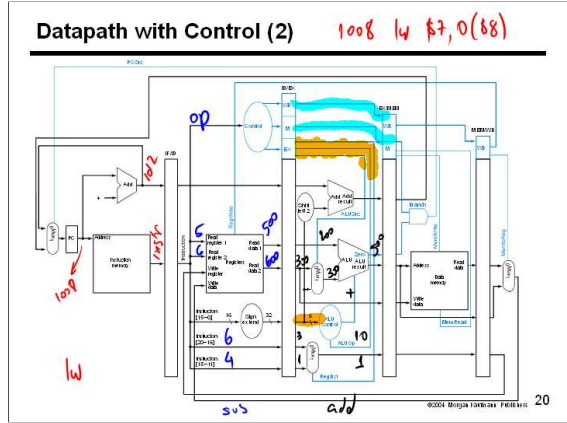
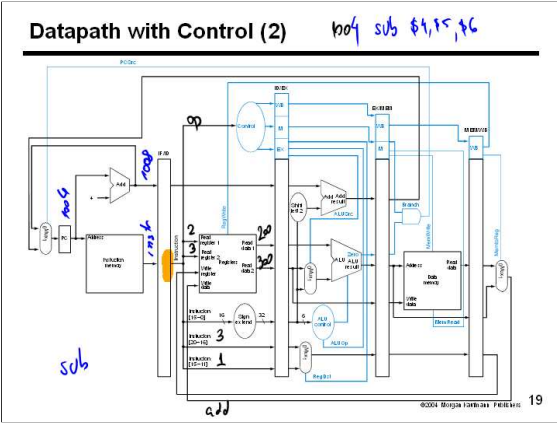
- Pass control signals along just like the data

Instruction	Execution/Address Calculation-stage control lines			Memory access stage control lines			Write-back stage control lines	
	Reg Dest	ALU Op	ALU Ono	Branch	Mem Read	Mem Write	Reg Write	Mem To Reg
R-format	1	1	0	0	0	0	1	0
lw	0	0	0	1	0	1	0	1
sw	X	0	0	1	0	0	0	X
beq	X	0	1	0	1	0	0	X

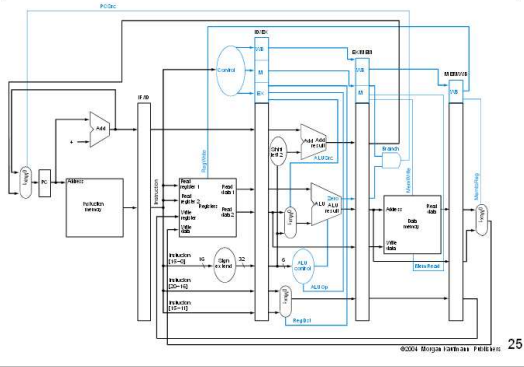


Datapath with Control (1) `lw $1, 100($0); add $1, $2, $3`

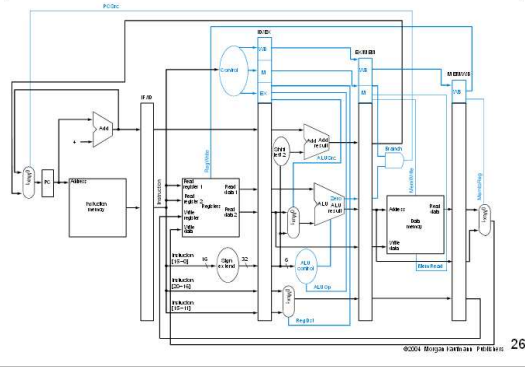




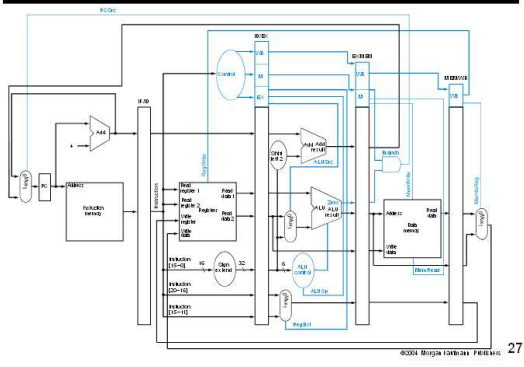
Datapath with Control (8)



Datapath with Control (9)




Datapath with Control (10)



E se existirem dependências de dados?

- Faça o diagrama de tempo abaixo ignorando as dependências
- Indique as dependências e mostre por que elas atrapalham




Como resolver por software?

- Modifique o código para que as dependências não atrapalhem
 - Não dá para trocar os registradores!!!
- Qual o impacto desta solução?

- add \$1, \$2, \$3
NOP / NOP
- add \$4, \$1, \$2
- add \$5, \$1, \$4
- add \$6, \$1, \$5

Dependencies

- Problem with starting next instruction before first is finished
 - dependencies that "go backward in time" are data hazards



Software Solution

- Have compiler guarantee no hazards
- Where do we insert the "nops"?

```


sub   $2, $1, $3
and   $12, $2, $5
or    $13, $6, $2
add   $14, $2, $2
sw    $15, 100($2)
    
```

→ NOP/NOP

- Problem: this really slows us down!

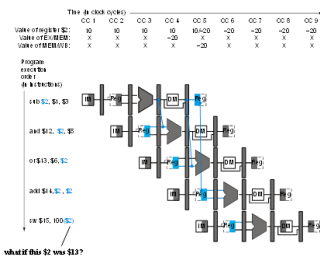
Solução por hardware

- Quando que os resultados das instruções ficam prontos?
 - Você consegue melhorar o circuito?



Forwarding

- Use temporary results, don't wait for them to be written
 - register file forwarding to handle read/write to same register
 - ALU forwarding



Forwarding

- De onde para onde?

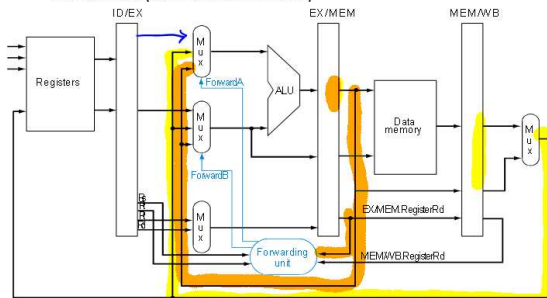
MEM → EX
WB → EX

- Funciona sempre?

Now lw \$3, 0(\$2) → NOP
add \$1, \$3, \$4


Forwarding

- The main idea (some details not shown)



Can't always forward

- Load word can still cause a hazard:
 - an instruction tries to read a register following a load instruction that writes to the same register.



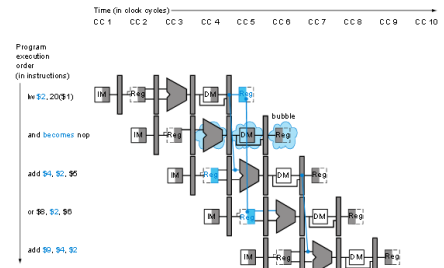
- Thus, we need a hazard detection unit to "stall" the load instruction

Como resolver quando não funciona?

- Software
- Hardware

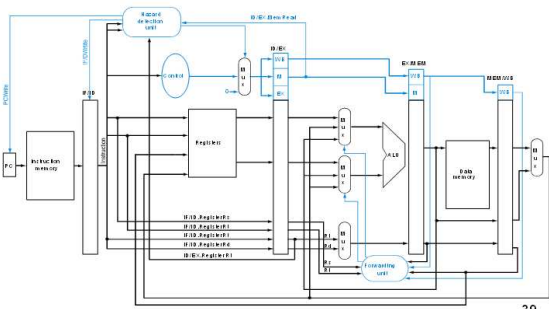
Stalling

- We can stall the pipeline by keeping an instruction in the same stage



Hazard Detection Unit

- Stall by letting an instruction that won't write anything go forward




E os saltos?

- O que acontece quando um salto ocorre?
- Quando você descobre que uma instrução precisa saltar?
- E aí?

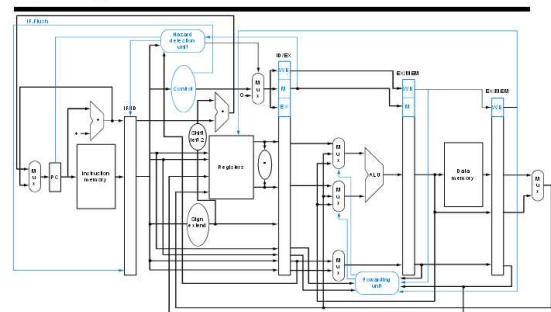
Branch Hazards

- When we decide to branch, other instructions are in the pipeline!



- We are predicting "branch not taken"
 - need to add hardware for flushing instructions if we are wrong

Flushing Instructions




Note: we've also moved branch decision to ID stage

Processadores advinham branches

- Você consegue descobrir como?
 - Dica: Lembre-se que os programas têm laços e as instruções estão em lugares fixos da memória

Branches

- If the branch is taken, we have a penalty of one cycle
- For our simple design, this is reasonable
- With deeper pipelines, penalty increases and static branch prediction drastically hurts performance
- Solution: dynamic branch prediction



A 2-bit prediction scheme

Branch Prediction

- Sophisticated Techniques:
 - A “branch target buffer” to help us look up the destination
 - Correlating predictors that base prediction on global behavior and recently executed branches (e.g., prediction for a specific branch instruction based on what happened in previous branches)
 - Tournament predictors that use different types of prediction strategies and keep track of which one is performing best.
 - A “branch delay slot” which the compiler tries to fill with a useful instruction (make the one cycle delay part of the ISA)
- Branch prediction is especially important because it enables other more advanced pipelining techniques to be effective!
- Modern processors predict correctly 95% of the time!

Improving Performance

- Try and avoid stalls! E.g., reorder these instructions:


```

            lw $t0, 0($t1)
            lw $t2, 4($t1)
            sw $t2, 0($t1)
            sw $t0, 4($t1)
            
```
- Dynamic Pipeline Scheduling
 - Hardware chooses which instructions to execute next
 - Will execute instructions out of order (e.g., doesn't wait for a dependency to be resolved, but rather keeps going!)
 - Speculates on branches and keeps the pipeline full (may need to rollback if prediction incorrect)
- Trying to exploit instruction-level parallelism

Advanced Pipelining

- Increase the depth of the pipeline
 - Start more than one instruction each cycle (multiple issue)
 - Loop unrolling to expose more ILP (better scheduling)
 - “Superscalar” processors
 - DEC Alpha 21264: 9 stage pipeline, 6 instruction issue
 - All modern processors are superscalar and issue multiple instructions usually with some limitations (e.g., different “pipes”)
 - VLIW: very long instruction word, static multiple issue (relies more on compiler technology)
- This class has given you the background you need to learn more!

Chapter 6 Summary

- Pipelining does not improve latency, but does improve throughput

