### Control Dependences (branch/jump) in Pipelines

- 'Data Dependence' = next instruction uses data (register/memory) from previous
- 'Control Dependence' = which is the next instruction depends on the previous
- Control Dependences arise from 'Control Transfer Instructions (CTI)'
- Control Transfer Instructions (CTI) are: Jump and Branch Instructions
- 'Jumps' are Unconditional CTI's: they always transfer control
- 'Branches' are Conditional CTI's: whether or not they transfer control depends on the result of a data comparison that they have to perform

*Statistics* (rough numbers, in a majority of programs, but NOT always so):

- Branches are about 15–16% of all ('dynamically') executed instructions in a program
  - about 2/3 of executed branches are 'taken' (successful) =  $\sim 10\%$  of all instr.
  - about 1/3 of executed branches are not taken (unsuccessful) =  $\sim 5\%$  of all instr.
  - most backwards branches appear in loops, and they are about 90% taken
- Jumps are about 4–5% of all executed instructions in a program
  - procedure calls are about 1%, and returns another ~1%, of all executed instr.

Copyright 2020 University of Crete – https://www.csd.uoc.gr/~hy225/20a/copyright.html (slides 1–9); Elsevier (slides 10–11)

## Branch Taken example

add

fetch

► 44.

+4

36:

fetch

► 40:

- In modern processors, branch latency is quite long
- In our simple pipeline, branch latency is 2 cycles (read registers; compare) (with MIPS-style comparisons (beq/bne only) it could even be 1 cycle)

DM

branch! (2 or more cycles)

sd

• Example here with 3–cycle branch latency

**ALU** 

fetch



► 48: and fetch • need to abort +4 noop noop noop noop ► 52: fetch speculative execution +4 - 56 before it causes fetch **ALU** DM WB ld permanent damage: before DM and WB stages 76: fetch ALU DM xor

WB

- In this example, each taken branch causes the loss of 3 extra clock cycles
- About 2/3 of all executed branches are taken, so this is a heavy loss

## Branch Not-taken example

- A not-taken branch is equivalent to a noop instruction
- In the simple fetch-next-below policy that we used up to now, not-taken branches cost NO extra clock cycle



36: add ...

48: and ...

44: sd

40: beq ..., goto72

• Can we do any better for the majority of branches (taken branches – and jumps) ?? 3

## Branch Target known in 1 Cycle

- Branch Prediction:
- Simplest possible prediction, here: branches always taken
  ~65% accuracy: about 2/3 of executed branches are taken





• In this example, each taken branch causes the loss of 1 extra clock cycle

#### Branch with failed Prediction example

• Simplest possible prediction, again: branches always taken ~65% accuracy: about 2/3 of executed branches are taken (reason: loop branches (backwards) ~90% taken)

DM

add

**ALU** 

36:

fetch



5

**Speculative** Execution beg pc+2\*Imm → 40: rs1 cmp rs2 fetch **Continued Execution** +4► 44: fetch **ALU** WB sd DM +4 +48 noop noop noop ld fetch 72 *What happens* +4 noop noop noop noop ► 76: fetch *if this prediction* +4 80 ended up being wrong: ALU DM WB fetch and 48: if this branch is eventually not taken +452: fetch ALU DM or

WB

• In this example, each non-taken branch causes the loss of 2 extra clock cycles

• ~5% jumps + (~15% branches \* 2/3 taken) ~= 15% good prediction, versus ~5% bad prediction

#### Branch Target Buffer (BTB)

 A small table – a cache, like a hash table – containing pairs of (instruction) addresses for which there is statistical evidence that their next–PC is something other than PC+4

PC of a jump or branch–likely instruction;

Target PC to which this instruction usually went, in the past.



	'
260	200
40	72
88	120
180	160

- A 'best approximation' not necessarily correct information
- Branches that are believed not-taken are NOT entered into the BTB
- Like IM –the Instruction Cache– this will oftentimes 'overflow': old pairs are removed to make room for more recent ones
- May be complemented with a small hardware stack:
  - on every call (jal ra,...), push the return address;
  - on every return (jr ra), pop an address and predict jumpin to that one
- In parallel with each Fetch, search the fetched instruction's PC value in the BTB



(\*) when IRvalid==0, treat IR as containing a noop instruction

#### When the BTB prediction is Correct

• When a matching BTB entry is found, use its Prediction; else, fetch from PC+4

- - -





• When Prediction is Correct, NO extra clock cycles are lost!

#### When the BTB prediction is Wrong

• Prediction says: After fetching from 40, fetch from 72

**ALU** 

fetch

BTB

**^**76:

36:

fetch

BTB

260

40

88

180

(BTB)

40:

\_ \_ \_

200

72

120

160

add

fetch

**≁**72:

• But this time, the branch ends up going the other way: to 44

DM

ld

fetch

BTB



• When Mispredicted, branches cost 3 extra clock cycles in this pipeline 9

BTB



- Mispredict as taken on last iteration of inner loop
- Then mispredict as not taken on first iteration of inner loop next time around



# **2-Bit Predictor**

#### Only change prediction on two successive mispredictions





