

## Constructive Computer Architecture:

# Branch Prediction

Arvind

Computer Science & Artificial Intelligence Lab.  
Massachusetts Institute of Technology

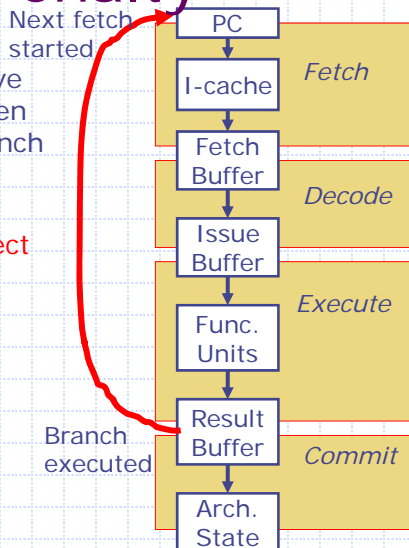
October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-1

## Control Flow Penalty

- ◆ Modern processors may have > 10 pipeline stages between next PC calculation and branch resolution !
- ◆ How much work is lost if pipeline doesn't follow correct instruction flow?
- ◆ What fraction of executed instructions are branch instructions?



October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-2

# How frequent are branches? ARM Cortex 7

Blem et al [HPCA 2013] Spec INT 2006

ARM Cortex-A9; ARMv7 ISA					
Benchmark	Total Instructions	branch %	load %	store %	other %
astar	1.47E+10	16.0	55.6	13.0	15.4
bzip2	2.41E+10	8.7	34.6	14.4	42.2
gcc	5.61E+09	10.2	19.1	11.2	59.5
gobmk	5.75E+10	10.7	25.4	7.2	56.8
hmmer	1.56E+10	5.1	41.8	18.1	35.0
h264	1.06E+11	5.5	30.4	10.4	53.6
libquantum	3.97E+08	11.5	8.1	11.7	68.7
omnetpp	2.67E+09	11.7	19.3	8.9	60.1
perlbench	2.69E+09	10.7	24.6	9.3	55.5
sjeng	1.34E+10	11.5	39.3	13.7	35.5
Average		8.2	31.9	10.9	49.0

Every 12<sup>th</sup> instruction is a branch

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-3

# How frequent are branches? X86

Blem et al [HPCA 2013] Spec INT 2006

core i7; x86 ISA					
Benchmark	Total Instructions	branch %	load %	store %	other %
astar	5.71E+10	6.9	19.5	6.9	66.7
bzip2	4.25E+10	11.1	31.2	11.8	45.9
hmmer	2.57E+10	5.3	30.5	9.4	54.8
gcc	6.29E+09	15.1	22.1	14.1	48.7
gobmk	8.93E+10	12.1	21.7	13.4	52.7
h264	1.09E+11	7.1	46.8	18.5	27.6
libquantum	4.18E+08	13.2	39.3	6.8	40.7
omnetpp	2.55E+09	16.4	28.6	21.4	33.7
perlbench	2.91E+09	17.3	25.9	16.0	40.8
sjeng	2.11E+10	14.8	22.8	11.0	51.4
Average		9.4	31.0	13.4	46.2

Every 10<sup>th</sup> or 11<sup>th</sup> instruction is a branch

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-4

# How frequent are branches? ARM Cortex 7

Blem et al [HPCA 2013] Spec FP 2006

Benchmark	ARM Cortex-A9; ARMv7 ISA				
	Total Instructions	branch %	load %	store %	other %
bwaves	3.84E+11	13.5	1.4	0.5	84.7
cactusADM	1.02E+10	0.5	51.4	17.9	30.1
leslie3D	4.92E+10	6.2	2.0	3.7	88.1
milc	1.38E+10	6.5	38.2	13.3	42.0
tonto	1.30E+10	10.0	40.5	14.1	35.4
Average		12.15	4.68	1.95	81.22

Every 8<sup>th</sup> instruction is a branch

# How frequent are branches? X86

Blem et al [HPCA 2013] Spec FP 2006

Benchmark	core i7; x86 ISA				
	Total Instructions	branch %	load %	store %	other %
bwaves	3.41E+10	3.2	51.4	16.8	28.7
cactusADM	1.05E+10	0.4	55.3	18.6	25.8
leslie3D	6.25E+10	4.9	35.3	12.8	46.9
milc	3.29E+10	2.2	32.2	13.8	51.8
tonto	4.88E+09	7.1	27.2	12.4	53.3
Average		3.6	39.6	14.4	42.4

Every 27<sup>th</sup> instruction is a branch

## Observations

- ◆ Control transfer happens every 8<sup>th</sup> to 30<sup>th</sup> instruction
- ◆ Static vs dynamic predictors: Does the prediction depend upon the execution history?
- ◆ Processors often use more than one predictor and it takes considerable effort to
  - Integrate a prediction scheme in the pipeline
  - Understand the interactions between various schemes
  - Understand the performance implications
- ◆ There is a plethora of branch prediction schemes – their importance grows with the depth of processor pipeline

we will start with the basics ...

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-7

## RISC V Branches & Jumps

Each instruction fetch depends on some information from the preceding instruction:

1. Is the preceding instruction a taken branch?
2. If so, what is the target address?

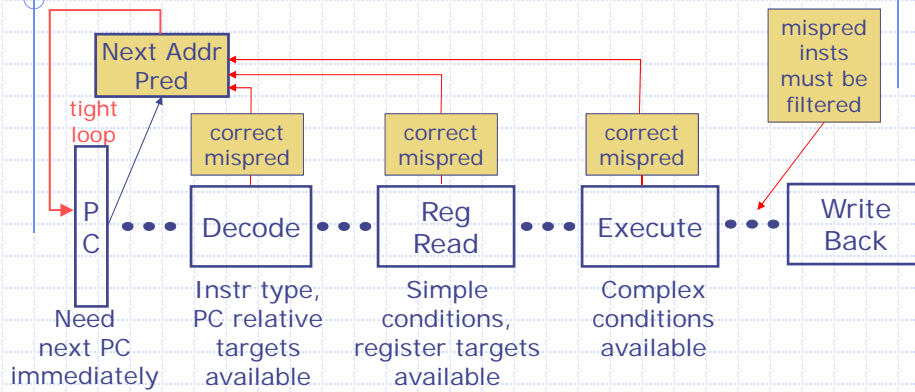
<i>Instruction</i>	<i>Direction known after</i>	<i>Target known after</i>
JAL		
JALR		
BEQ/BNE ...		

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-8


# Overview of control prediction



Given (pc, ppc), a misprediction can be corrected (used to redirect the pc) as soon as it is detected. In fact, pc can be redirected as soon as we have a "better" prediction. However, for forward progress it is important that a correct pc should never be redirected.

# Static Branch Prediction


- ◆ Since most instructions do not result in a control transfer, pc+4 is a good predictor
- ◆ Overall probability a branch is taken is ~60-70% but:



- ◆ ISA can attach preferred direction semantics to branches, e.g., Motorola MC88110
  - bne0 (preferred taken)      beq0 (not taken)
- ◆ ISA can allow arbitrary choice of statically predicted direction, e.g., HP PA-RISC, Intel IA-64
  - reported as ~80% accurate

# Dynamic Branch Prediction

*learning based on past behavior*



## ◆ Temporal correlation

- The way a branch resolves may be a good predictor of the way it will resolve at the next execution

## ◆ Spatial correlation

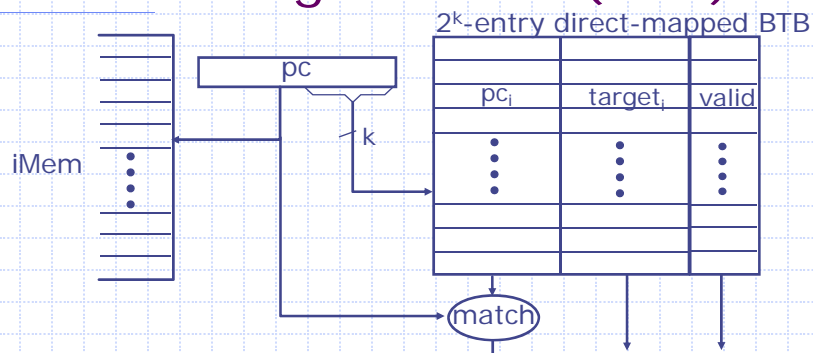
- Several branches may resolve in a highly correlated manner (a preferred path of execution)

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-11

# Next Address Predictor: Branch Target Buffer (BTB)



## ◆ BTB remembers recent targets for a set of *control instructions*

- Fetch: looks for the pc and the associated target in BTB; if pc in not found then ppc is pc+4
- Execute: checks prediction, if wrong kills the instruction and updates BTB (only for branches and jumps)

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-12

## Next Addr Predictor interface

*Predictor training:* On a pc misprediction, information about redirecting the pc is passed to the fetch stage. However for training the branch predictors information has to be passed even when there is no misprediction.

```
interface AddrPred;  
  method Addr nap(Addr pc);  
  method Action update(Redirect redirect);  
endinterface  
  
                <pc, target-pc, taken/not-taken>
```

Next two implementations:

- a) Simple PC+4 predictor
- b) Predictor using BTB

## Simple PC+4 predictor

```
module mkPcPlus4(AddrPred);  
  method Addr nap(Addr pc);  
    return pc + 4;  
  endmethod  
  
  method Action update(Redirect rd);  
  endmethod  
endmodule
```

## BTB predictor

```
module mkBtb(AddrPred);
  RegFile#(BtbIndex, Addr) ppcArr <- mkRegFileFull;
  RegFile#(BtbIndex, BtbTag) entryPcArr <- mkRegFileFull;
  Vector#(BtbEntries, Reg#(Bool))
    validArr <- replicateM(mkReg(False));
  function BtbIndex getIndex(Addr pc)=truncate(pc>>2);
  function BtbTag getTag(Addr pc) = truncateLSB(pc);
  method Addr nap(Addr pc);
    BtbIndex index = getIndex(pc);
    BtbTag tag = getTag(pc);
    if(validArr[index] && tag == entryPcArr.sub(index))
      return ppcArr.sub(index);
    else return (pc + 4);
  endmethod
  method Action update(Redirect redirect); ...
endmodule
```

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-15

## BTB predictor update method

redirect input contains a pc, the correct next pc and whether the branch was taken or not (to avoid making entries for not-taken branches)

```
method Action update(Redirect redirect);
  if(redirect.taken)
    begin
      let index = getIndex(redirect.pc);
      let tag = getTag(redirect.pc);
      validArr[index] <= True;
      entryPcArr.upd(index, tag);
      ppcArr.upd(index, redirect.nextPc);
    end
  else if(tag == entryPcArr.sub(index))
    validArr[index] <= False;
  endmethod
```

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-16



## Integrating BTB in the 2-Stage pipeline

```
module mkProc(Proc);
  Reg#(Addr)      pc <- mkRegU;
  RFile           rf <- mkRFile;
  IMemory         iMem <- mkIMemory;
  DMemory         dMem <- mkDMemory;
  Fifo#(Decode2Execute) d2e <- mkFifo;
  Reg#(Bool)      fEpoch <- mkReg(False);
  Reg#(Bool)      eEpoch <- mkReg(False);
  Fifo#(Addr)     redirect <- mkFifo;
  AddrPred        btb <- mkBtb

  Scoreboard#(1) sb <- mkScoreboard;
  rule doFetch ...
  rule doExecute ...
endmodule
```

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-17

## 2-Stage pipeline: doExecute rule Passing information for predictor update

```
rule doExecute;
  let x = d2e.first;
  ...
  if(epoch == eEpoch) begin
    let eInst = exec(dInst, rVal1, rVal2, pc, ppc);
    ...
    if(eInst.mispredict) begin
      redirect.enq(eInst.addr); eEpoch <= !eEpoch; end
    end
  end
  d2e.deq; sb.remove;
endrule
```

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-18

## 2-Stage pipeline: doFetch rule Updating the Branch Predictor

```

rule doFetch;
  let inst = iMem.req(pc);
  if(redirect.notEmpty) begin
    fEpoch <= !fEpoch; pc <= redirect.first;
    redirect.deq;      end

  else begin
    let ppc = nextAddrPredictor(pc);
    let dInst = decode(inst);
    ...
    end
endrule

```

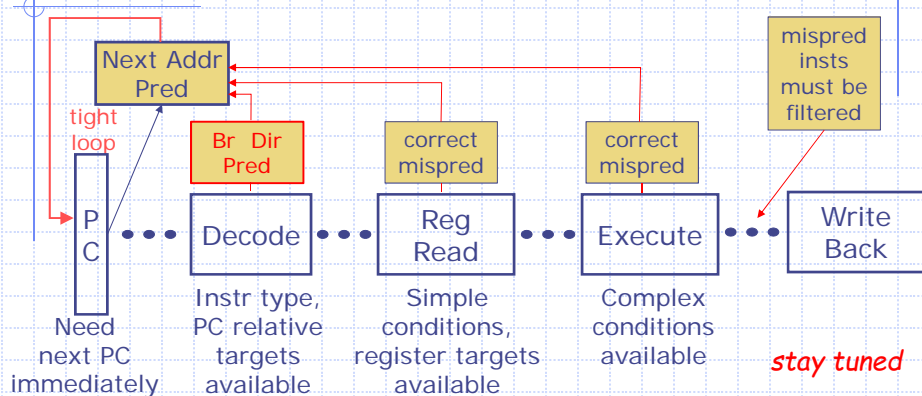
Update BTB but change pc only on a mispredict

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-19

## Multiple Predictors: BTB + Branch Direction Predictors



◆ Suppose we maintain a table of how a particular Br has resolved before. At the decode stage we can consult this table to check if the incoming (pc, ppc) pair matches our prediction. If not redirect the pc

October 24, 2016

<http://csg.csail.mit.edu/6.175>

L15-20