

Constructive Computer Architecture:

## Control Hazards

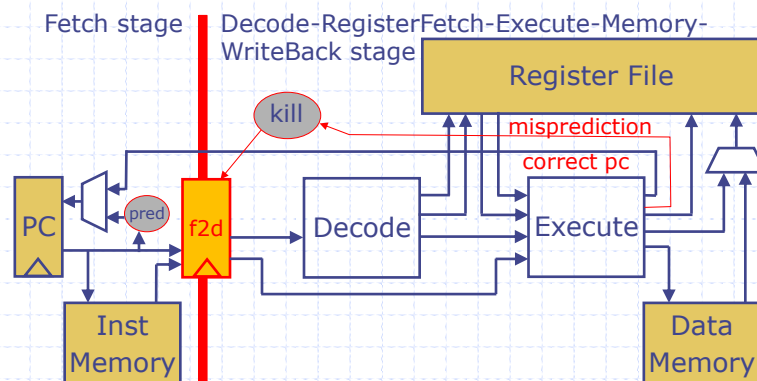
Arvind  
Computer Science & Artificial Intelligence Lab.  
Massachusetts Institute of Technology

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-1

## Two-stage Pipelined SMIPS



Fetch stage must predict the next instruction to fetch to have any pipelining

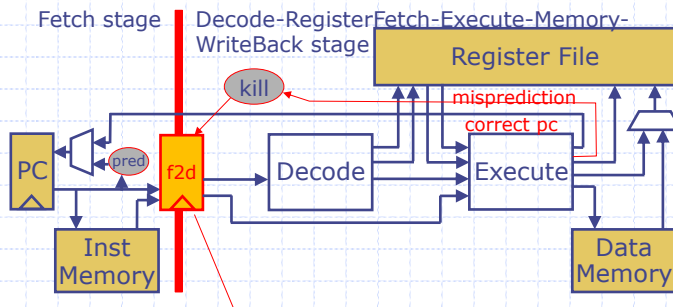
In case of a misprediction the Execute stage must kill the mispredicted instruction in f2d

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-2

## Two-stage Pipelined SMIPS



- ◆ f2d must contain a Maybe type value because sometimes the fetched instruction is killed
- ◆ Fetch2Decode type captures all the information that needs to be passed from Fetch to Decode, i.e.  
Fetch2Decode {pc:Addr, ppc: Addr, inst:Inst}

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-3

## Pipelining Two-Cycle SMIPS – single rule

```

rule doPipeline ;
  let instF = iMem.req(pc);
  let ppcF = nextAddr(pc); let nextPc = ppcF;
  let newf2d = Valid (Fetch2Decode{pc:pc,ppc:ppcF,
                                inst:instF});
  if(isValid(f2d)) begin
    let x = fromMaybe(?,f2d); let pcD = x.pc;
    let ppcD = x.ppc; let instD = x.inst;
    let dInst = decode(instD);
    ... register fetch ...;
    let eInst = exec(dInst, rVal1, rVal2, pcD, ppcD);
    ...memory operation ...
    ...rf update ...
    if (eInst.mispredict) begin nextPc = eInst.addr;
                              newf2d = Invalid; end
  end
  pc <= nextPc; f2d <= newf2d;
endrule
  
```

fetch

execute

these values are being redefined

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-4

# Inelastic versus Elastic pipeline

- ◆ The pipeline presented is inelastic, that is, it relies on executing Fetch and Execute together or atomically
- ◆ In a realistic machine, Fetch and Execute behave more asynchronously; for example memory latency or a functional unit may take variable number of cycles
- ◆ If we replace f2d register by a FIFO then it is possible to make the machine more elastic, that is, Fetch keeps putting instructions into f2d and Execute keeps removing and executing instructions from f2d

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-5

# An elastic Two-Stage pipeline

```
rule doFetch ;
  let instF = iMem.req(pc);
  let ppcF = nextAddr(pc); pc <= ppcF;
  f2d.enq(Fetch2Decode{pc:pc, ppc:ppcF, inst:instF});
endrule

rule doExecute;
  let x = f2d.first; let pcD = x.pc;
  let ppcD = x.ppc; let instD = x.inst;
  let dInst = decode(instD);
  ... register fetch ...;
  let eInst = exec(dInst, rVal1, rVal2, pcD, ppcD);
  ...memory operation ...
  ...rf update ...
  if (eInst.mispredict) begin
    pc <= eInst.addr; f2d.clear; end
  else f2d.deq;
endrule
```

Can these rules execute concurrently assuming the FIFO allows concurrent enq, deq and clear?

No because of a possible double write in pc

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-6

## An elastic Two-Stage pipeline: for concurrency make pc into an EHR

```

rule doFetch ;
  let instF = iMem.req(pc[0]);
  let ppcF = nextAddr(pc[0]); pc[0] <= ppcF;
  f2d.enq(Fetch2Decode{pc:pc[0],ppc:ppcF,inst:inst});
endrule

rule doExecute;
  let x = f2d.first; let pcD = x.pc;
  let ppcD = x.ppc; let instD = x.inst;
  let dInst = decode(instD);
  ... register fetch ...;
  let eInst = exec(dInst, rVal1, rVal2, pcD, ppcD);
  ...memory operation ...
  ...rf update ...
  if (eInst.mispredict) begin
    pc[1] <= eInst.addr; f2d.clear; end
  else f2d.deq;
endrule

```

These rules can execute concurrently assuming the FIFO has (enq CF deq) and (enq < clear)

Double writes in pc have been replaced by prioritized writes in pc

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-7

## Conflict-free FIFO with a Clear method



```

module mkCFFifo(Fifo#(2, t)) provisos (Bits#(t, tSz));
  Ehr#(3, t) da <- mkEhr(?);
  Ehr#(3, Bool) va <- mkEhr(False);
  Ehr#(3, t) db <- mkEhr(?);
  Ehr#(3, Bool) vb <- mkEhr(False);
  rule canonicalize if(vb[2] && !va[2]);
    da[2] <= db[2]; va[2] <= True; vb[2] <= False; endrule
  method Action enq(t x) if(!vb[0]);
    db[0] <= x; vb[0] <= True; endmethod
  method Action deq if(va[0]);
    va[0] <= False; endmethod
  method t first if(va[0]);
    return da[0]; endmethod
  method Action clear;
    va[1] <= False ; vb[1] <= False endmethod
endmodule

```

If there is only one element in the FIFO it resides in da

first CF enq  
deq CF enq  
first < deq  
enq < clear

Canonicalize must be the last rule to fire!

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-8

# Why canonicalize must be the last rule to fire


```
rule foo ;  
  f.deq; if (p) f.clear  
endrule
```

Consider rule foo. If p is false then canonicalize must fire after deq for proper concurrency.

If canonicalize uses EHR indices between deq and clear, then canonicalize won't fire when p is false

```
first CF enq  
deq  CF enq  
first < deq  
enq < clear
```

# Correctness issue



- ◆ Once Execute redirects the PC,
  - no wrong path instruction should be executed
  - the next instruction executed must be the redirected one
- ◆ This is true for the code shown because
  - Execute changes the pc and clears the FIFO atomically
  - Fetch reads the pc and enqueues the FIFO atomically

## Killing fetched instructions

- ◆ In the simple design with combinational memory we have discussed so far, the mispredicted instruction was present in the f2d. So the Execute stage can atomically
  - Clear the f2d
  - Set the pc to the correct target
- ◆ In highly pipelined machines there can be multiple mispredicted and partially executed instructions in the pipeline; it will generally take more than one cycle to kill all such instructions

Need a more general solution than clearing the f2d FIFO


October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-11

## Epoch: a method for managing control hazards

- ◆ Add an epoch register in the processor state
- ◆ The Execute stage changes the epoch whenever the pc prediction is wrong and sets the pc to the correct value
- ◆ The Fetch stage associates the current epoch with every instruction when it is fetched
- ◆ The epoch of the instruction is examined when it is ready to execute. If the processor epoch has changed the instruction is thrown away



October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-12

## An epoch based solution

```
rule doFetch ;
  let instF=iMem.req(pc[0]);
  let ppcF=nextAddr(pc[0]); pc[0]<=ppcF;
  f2d.eng (Fetch2Decode{pc:pc[0],ppc:ppcF,epoch:epoch,
  inst:instF});
endrule
rule doExecute;
  let x=f2d.first; let pcD=x.pc; let inEp=x.epoch;
  let ppcD = x.ppc; let instD = x.inst;
  if(inEp == epoch) begin
    let dInst = decode(instD); ... register fetch ...;
    let eInst = exec(dInst, rVal1, rVal2, pcD, ppcD);
    ...memory operation ...
    ...rf update ...
    if (eInst.mispredict) begin
      pc[1] <= eInst.addr; epoch <= epoch + 1; end
    end
  end
  f2d.deq; endrule
```

Can these rules execute concurrently?

yes

two values for epoch are sufficient

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-13

## Discussion


- ◆ Epoch based solution kills one wrong-path instruction at a time in the execute stage
- ◆ It may be slow, but it is more robust in more complex pipelines, if you have multiple stages between fetch and execute or if you have outstanding instruction requests to the iMem
- ◆ It requires the Execute stage to set the pc and epoch registers simultaneously which may result in a long combinational path from Execute to Fetch

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-14

## Decoupled Fetch and Execute



- ◆ In decoupled systems a subsystem reads and modifies only local state atomically
  - In our solution, pc and epoch are read by both rules
- ◆ Properly decoupled systems permit greater freedom in independent refinement of subsystems

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-15

## A decoupled solution using epochs



- ◆ Add fEpoch and eEpoch registers to the processor state; initialize them to the same value
- ◆ The epoch changes whenever Execute detects the pc prediction to be wrong. This change is reflected immediately in eEpoch and eventually in fEpoch via a message from Execute to Fetch
- ◆ Associate the fEpoch with every instruction when it is fetched
- ◆ In the execute stage, reject, i.e., kill, the instruction if its epoch does not match eEpoch

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-16





## The Fetch rule

```
rule doFetch;
  let instF = iMem.req(pc);
  if(!execRedirect.notEmpty)
  begin
    let ppcF = nextAddrPredictor(pc);
    pc <= ppcF;
    f2d.enq(Fetch2Execute{pc: pc, ppc: ppcF,
                        inst: instF, epoch: fEpoch});
  end
else
  begin
    fEpoch <= !fEpoch; pc <= execRedirect.first;
    execRedirect.deq;
  end
endrule
```

pass the pc and predicted pc to the execute stage

Notice: In case of PC redirection, nothing is enqueued into f2d

October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-19

## The Execute rule

```
rule doExecute;
  let instD = f2d.first.inst; let pcF = f2d.first.pc;
  let ppcD = f2d.first.ppc; let inEp = f2d.first.epoch;
  if(inEp == eEpoch) begin
    let dInst = decode(instD);
    let rVal1 = rf.rd1(validRegValue(dInst.src1));
    let rVal2 = rf.rd2(validRegValue(dInst.src2));
    let eInst = exec(dInst, rVal1, rVal2, pcD, ppcD);
    if(eInst.iType == Ld) eInst.data <-
      dMem.req(MemReq{op: Ld, addr: eInst.addr, data: ?});
    else if (eInst.iType == St) let d <-
      dMem.req(MemReq{op: St, addr: eInst.addr, data: eInst.data});
    if (isValid(eInst.dst))
      rf.wr(validRegValue(eInst.dst), eInst.data);
    if(eInst.mispredict) begin
      execRedirect.enq(eInst.addr); eEpoch <= !inEp;
    end
  end
  f2d.deq;
endrule
```

exec returns a flag if there was a fetch misprediction

Can these rules execute concurrently?  
yes, assuming CF FIFOs

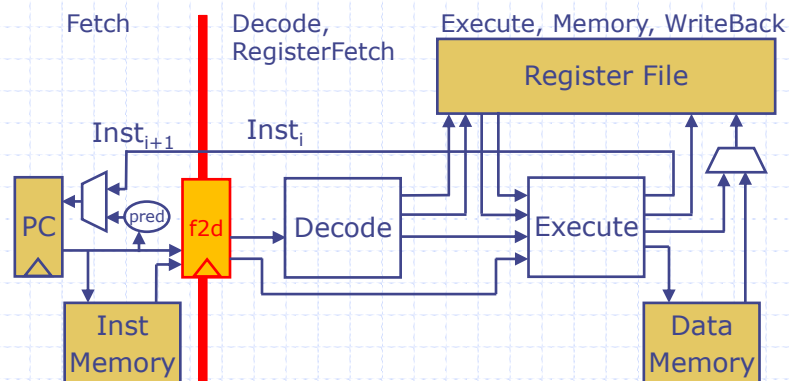
October 9, 2013

<http://csg.csail.mit.edu/6.S195>

L11-20

We will study more sophisticated branch prediction schemes later

Consider a different two-stage pipeline



Suppose we move the pipeline stage from Fetch to after Decode and Register fetch

What hazards will the pipeline have? Control? yes Any other?