

# Complex Pipelines and Branch Prediction

Ryan Lee

(slides adapted from prior 6.823 offerings)

# Since Last Time...

## 1. Complex Pipelines

- Superscalar execution
- Out-of-order (OoO) processing
  - Scoreboarding
  - OoO: Issue, completion, retiring
  - Register renaming

## 2. Branch Prediction

# Dependence vs. hazard


- Dependence is a property of programs
- Whether a dependence results in a hazard is a property of pipeline organizations

# Data hazard types

- RAW
- WAR
- WAW

I1: ADDI f0, f0, 0  
I2: ADDI f3, f0, 3  
I3: ADDI f4, f0, 4  
I4: ADDI f0, f5, 1  
I5: XOR f6, f6, f6  
I6: ADDI f0, f7, 1

Reads/Writes to f0




# Scoreboard

- A data structure that detects hazards dynamically
- Applicable to both in-order and out-of-order issue
- Why do we need this?
  - Many execution units
  - Variable execution latency
  - Dynamic instruction scheduling

# Scoreboard

- Can have many implementations!
- Example: In-order issue
  - WAR cannot happen (if value is latched to functional unit at issue)



- Can be simplified as Busy[FU#] and WP[reg#] (if WAW resolved conservatively)

# Scoreboard

- What strategy does it use to resolve RAW?
  - Stall
- How about bypass?
  - Less beneficial since the register write can happen right after execution finishes
  - Can still be incorporated to allow register read and write to happen in the same cycle

# Static vs. dynamic scheduling

- Reorder instructions to avoid hazards
- Static scheduling: programmer/compiler
- Dynamic scheduling: architectures
  - No need to re-compile!
  - Can handle unknown dependences and execution latencies



# Out-of-order execution

- Register renaming: an approach to resolve WAR and WAW hazards (caused by name dependences)
- Design tradeoffs
  - Data-in-ROB vs. unified-register-file
  - Centralized vs. distributed
  - ROB vs. issue queue + commit queue


# Branch Prediction

Control Flow Dependences. How to handle them?

- Stall: Delay until we know the next PC
- Speculate: Guess next value
- Do something else: Multi-threading

# Branch Predictors

- 1-bit predictor
- 2-bit predictor




# Branch Predictors


Two empirical observations

1. A branch's outcome can be correlated with other branches' outcomes
  - Global branch correlation
2. A branch's outcome can be correlated with past outcomes of the same branch
  - Local branch correlation

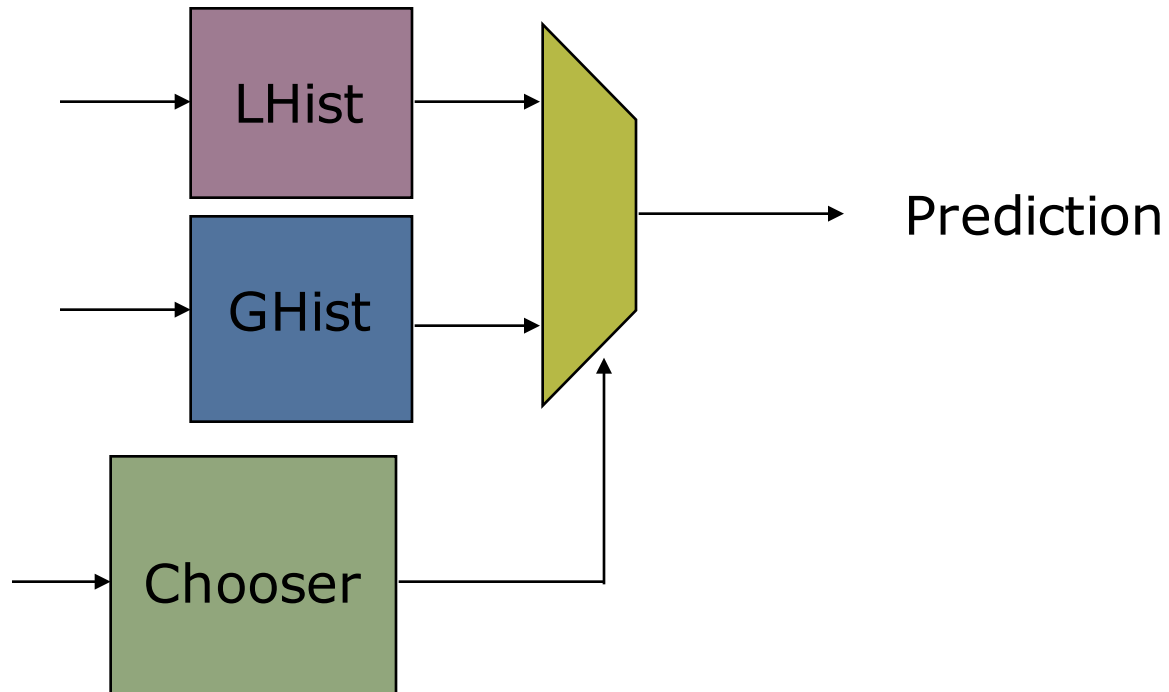
# History-based Prediction



# Two-level Predictor



# Tournament Predictors



# Lab 2 Due April 2

- Get going early
- Incrementally build more complex predictors
  - 2-bit predictor
  - Local history predictor
  - Tournament predictor
  - and so on
- Recommend researching more advanced predictors for full credit
  - TAGE, Perceptron, etc...
- Note on bug: running pintool with 1s will crash the pintool on Ubuntu 16.04 machines
  - Works fine with running lab benchmarks



Questions?