


Lecture 8: Branch Prediction, Dynamic ILP

- Topics: branch prediction, out-of-order processors
(Sections 2.3-2.6)

Control Hazards


- In the 5-stage in-order processor: assume always taken or assume always not taken; if the branch goes the other way, squash mis-fetched instructions (momentarily, forget about branch delay slots)
- Modern in-order and out-of-order processors: dynamic branch prediction; instead of a default not-taken assumption, either predict not-taken, or predict taken-to-X, or predict taken-to-Y
- Branch predictor: a cache of recent branch outcomes

Pipeline without Branch Predictor



In the 5-stage pipeline, a branch completes in two cycles →
If the branch went the wrong way, one incorrect instr is fetched →
One stall cycle per incorrect branch

Pipeline with Branch Predictor



In the 5-stage pipeline, a branch completes in two cycles →
If the branch went the wrong way, one incorrect instr is fetched →
One stall cycle per incorrect branch

Branch Mispredict Penalty

- Assume: no data or structural hazards; only control hazards; every 5th instruction is a branch; branch predictor accuracy is 90%
- Slowdown = $1 / (1 + \text{stalls per instruction})$
- Stalls per instruction = % branches x %mispreds x penalty
= 20% x 10% x 1
= 0.02
- Slowdown = $1/1.02$; if penalty = 20, slowdown = $1/1.4$

1-Bit Bimodal Prediction

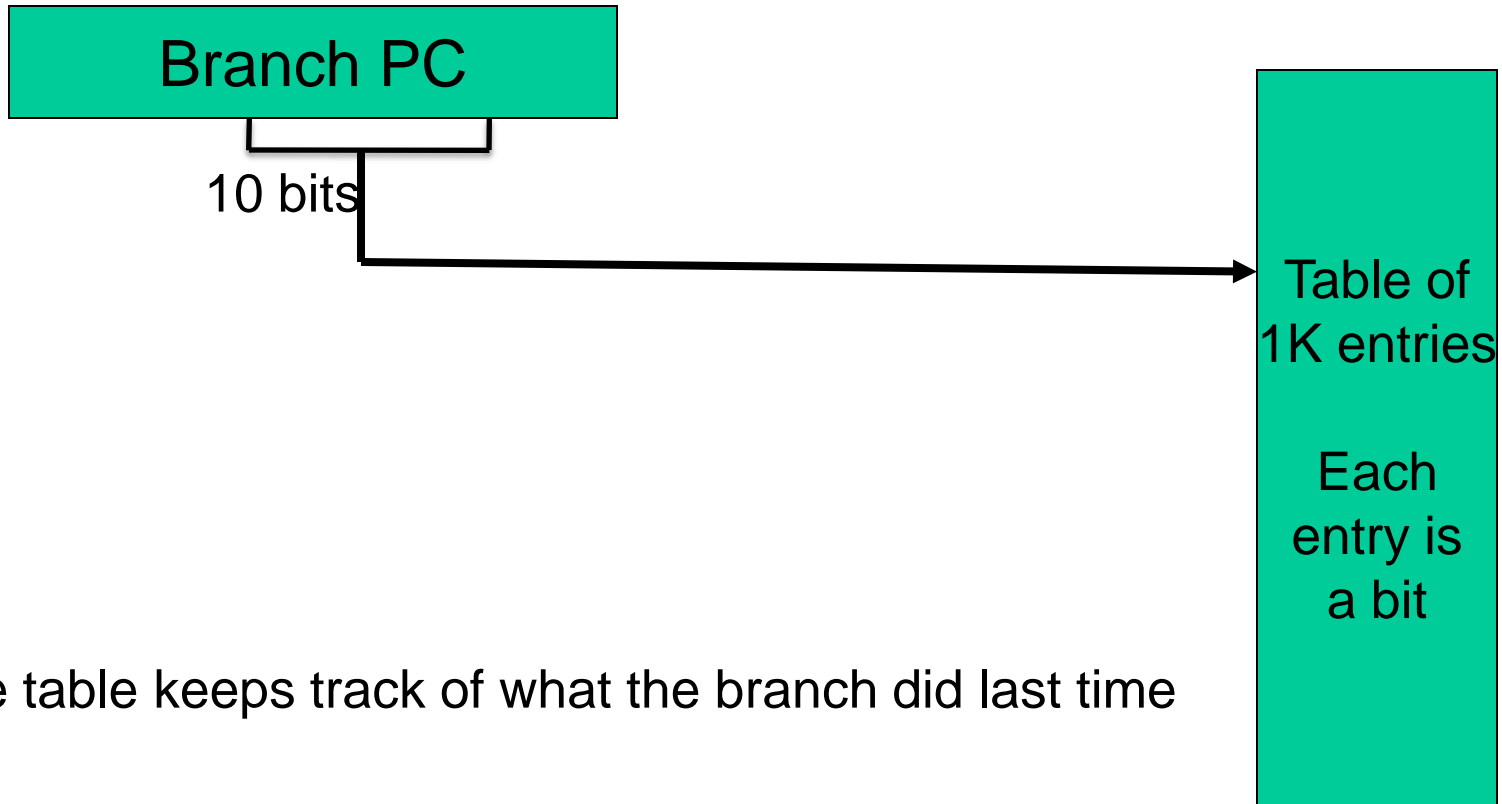
- For each branch, keep track of what happened last time and use that outcome as the prediction
- What are prediction accuracies for branches 1 and 2 below:

```
while (1) {  
    for (i=0;i<10;i++) {                branch-1  
        ...  
    }  
    for (j=0;j<20;j++) {                branch-2  
        ...  
    }  
}
```

2-Bit Bimodal Prediction


- For each branch, maintain a 2-bit saturating counter:
if the branch is taken: $\text{counter} = \min(3, \text{counter} + 1)$
if the branch is not taken: $\text{counter} = \max(0, \text{counter} - 1)$
- If ($\text{counter} \geq 2$), predict taken, else predict not taken
- Advantage: a few atypical branches will not influence the prediction (a better measure of “the common case”)
- Especially useful when multiple branches share the same counter (some bits of the branch PC are used to index into the branch predictor)
- Can be easily extended to N-bits (in most processors, $N \neq 2$)

Bimodal 1-Bit Predictor



The table keeps track of what the branch did last time

Bimodal 2-Bit Predictor




Correlating Predictors

- Basic branch prediction: maintain a 2-bit saturating counter for each entry (or use 10 branch PC bits to index into one of 1024 counters) – captures the recent “common case” for each branch
- Can we take advantage of additional information?
 - If a branch recently went 01111, expect 0; if it recently went 11101, expect 1; can we have a separate counter for each case?
 - If the previous branches went 01, expect 0; if the previous branches went 11, expect 1; can we have a separate counter for each case?

Hence, build **correlating predictors**

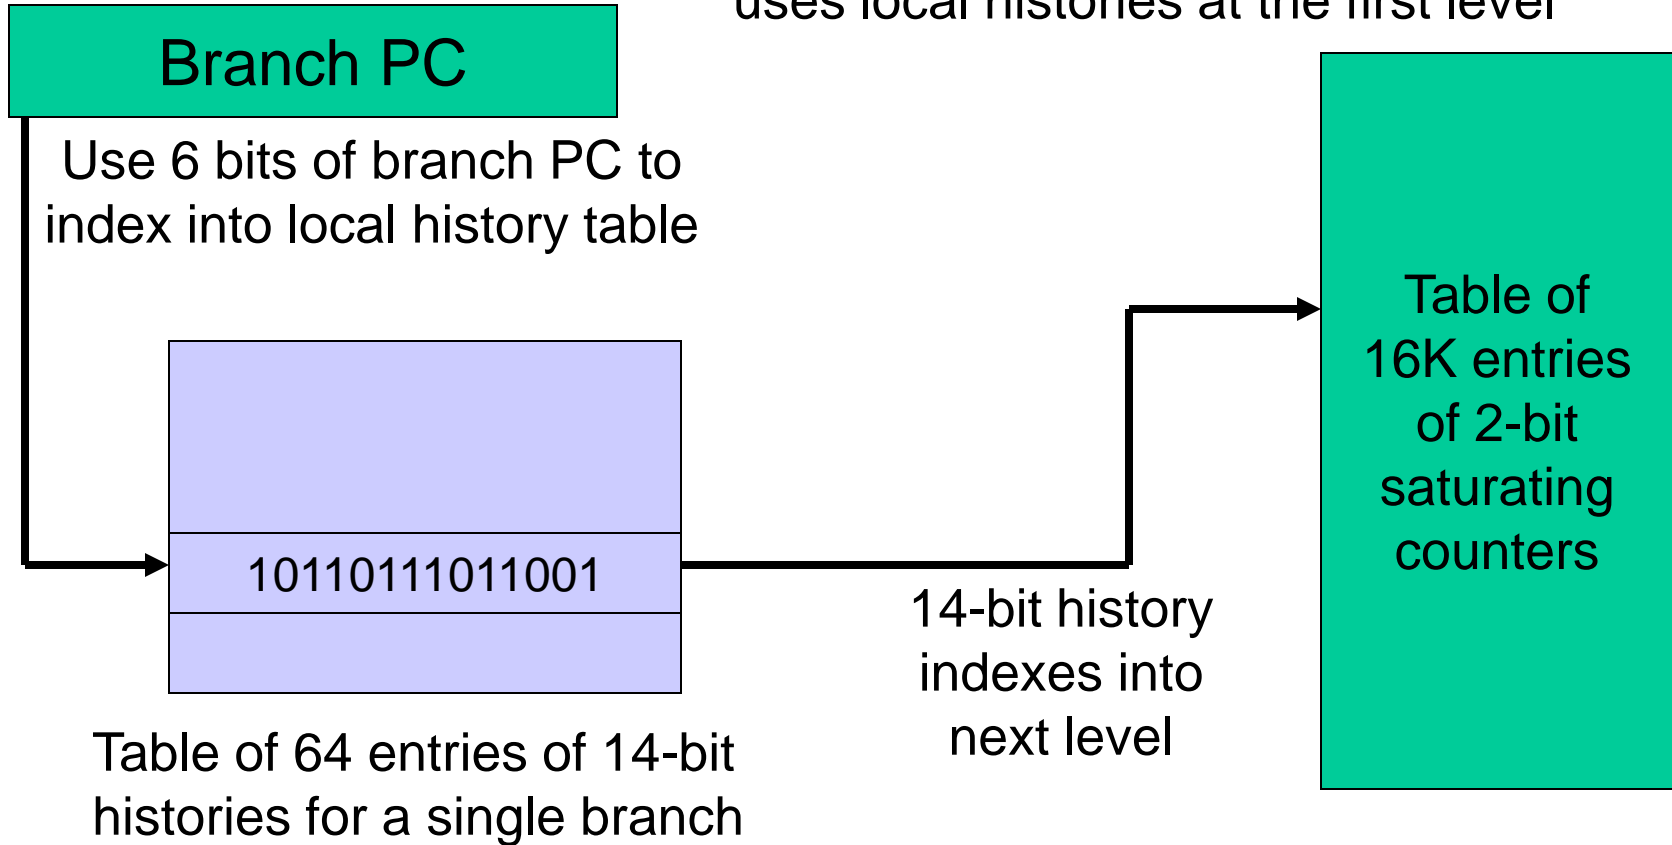
Global Predictor




The table keeps track of the common-case outcome for the branch/history combo

Local Predictor

Also a two-level predictor that only uses local histories at the first level



Local Predictor




The table keeps track of the common-case outcome for the branch/local-history combo

Local/Global Predictors

- Instead of maintaining a counter for each branch to capture the common case,
 - Maintain a counter for each branch and surrounding pattern
 - If the surrounding pattern belongs to the branch being predicted, the predictor is referred to as a local predictor
 - If the surrounding pattern includes neighboring branches, the predictor is referred to as a global predictor

Tournament Predictors

- A local predictor might work well for some branches or programs, while a global predictor might work well for others
- Provide one of each and maintain another predictor to identify which predictor is best for each branch



Alpha 21264:
1K entries in level-1
1K entries in level-2

4K entries
12-bit global history

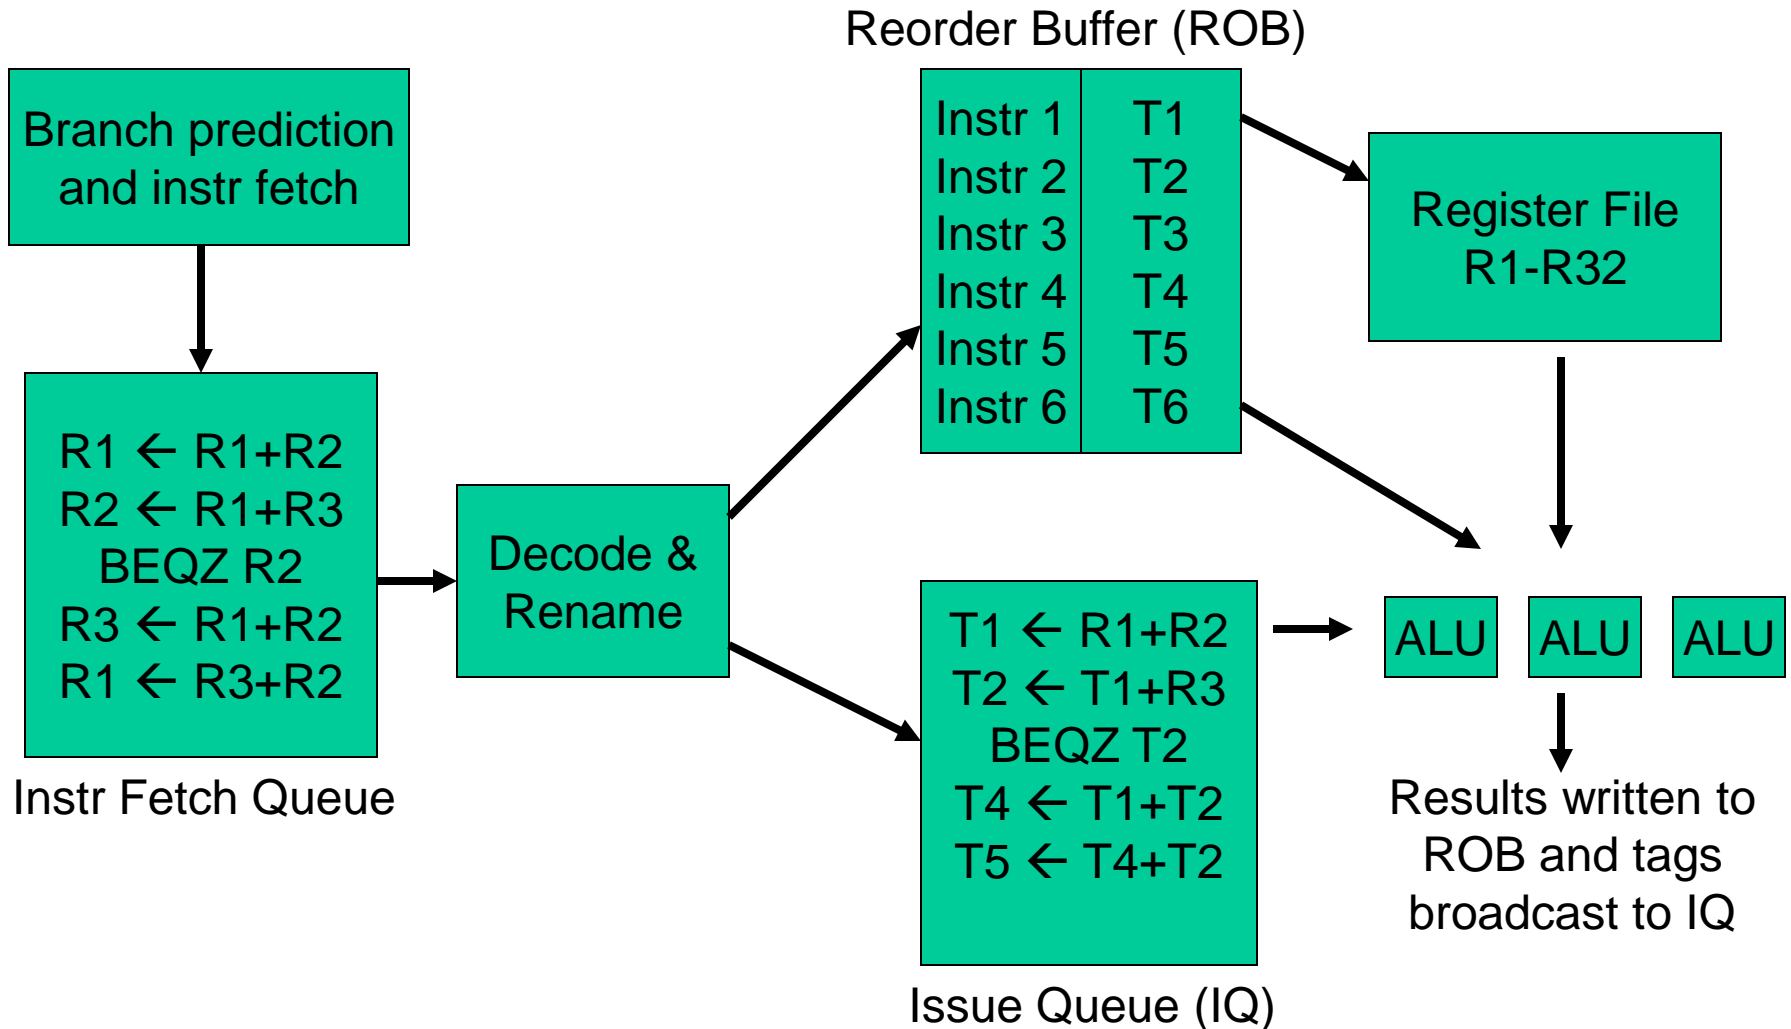
4K entries

Total capacity: ?

Branch Target Prediction

- In addition to predicting the branch direction, we must also predict the branch target address
- Branch PC indexes into a predictor table; indirect branches might be problematic
- Most common indirect branch: return from a procedure – can be easily handled with a stack of return addresses

An Out-of-Order Processor Implementation



Title

- Bullet