Branch Prediction

03-1

This Set: How to predict branch direction targets.

Review Material

McFarling 93, a concise description of basic branch prediction techniques.

EE 4720 Lecture Set 12

References at end of set.

Branch Prediction Techniques

03-2

Basic Techniques (Covered in EE 4720)

Bimodal (Two-Bit Counter, don't confuse with bi-mode)

gshare (Global History $\oplus$ PC)

Local History

Advanced Techniques

Variation on one or more basic techniques.

Performance of Basic Techniques

03-3

Bimodal (Don't confuse with bi-mode.)

Accuracy: 93% on SPEC89. (McFarling 93)

Advantage: Fast warmup.

Limitation: Only works for highly biased branches.

gshare (Global History $\oplus$ PC)

Accuracy: 96% on SPEC89. (McFarling 93)

Advantage: Handles a wide variety of branches.

Limitations: Slow warmup, large table needed to avoid collisions.

03-4

Local History

Accuracy: 97% on SPEC89. (McFarling 93)

Advantage: Can precisely predict short loops.

Limitation: Cannot predict one branch based on another.

Advanced Techniques

03-5

Each makes better use of basic techniques.

Hybrid Predictor: Choose best predictor for each branch.

Skewed: Avoid collisions by storing counter in multiple places.

Bias Filtering: Avoid collisions by accounting for bias.

History Filtering: Use a subset of history appropriate for branch.

Hybrid Branch Predictors

03-6

Primary Reference: McFarling 93. Early work on hybrid predictors, also provides good summary of basic predictors.

Other references: Chang 94, Evers 96.

Hybrid Branch Predictor:

A branch predictor that predicts a branch by choosing a prediction made by one of several predictors. The choice might be based on a chooser table, or by using the majority prediction.

Status: Used in existing processors, including Alpha 21264.

Hybrid Branch Predictors

03-7

Problem: Some branches best predicted by local predictor, some by bimodal, etc.

Solution: Use several predictors.

Details:

Maintain two or more predictors.

Each one is updated for every branch.

Use a *chooser* table (also called a *metapredictor* to pick the best predictor . . .
. . . or have (odd number of) predictors vote on a direction.

(To be covered in detail on blackboard.)

Typical Configuration

Use gshare and local predictors.

Skewed Predictors

03-8

Primary Reference: Michaud 97.

Problem: Collisions in PHT.

Solution: Predict using counters stored in several places within PHT.

Details:

Use three PHTs.

Hash (combine) global history with PC three different ways ...
... creating three different indices (index 1, index 2, index 3, one per PHT) ...
... and use these to lookup one counter in each PHT.

Use majority prediction.

Bias Filtering

03-9

Problem: Collisions by highly biased branches.

Solutions:

Use separate PHTs for taken and not-taken branches. (Bi-Mode, Lee 97)

Use PHTs only for branches not predicted by bimodal predictor (YAGS, Eden 98).

If bit in instruction cache set invert prediction. (Agree Sprangle 97)

Details covered in class.

History Filtering

03-10

Problem: Global history has irrelevant outcomes.

Solutions:

Adjust length of global history. Juan 98

Take a weighted combination of outcomes, with weights chosen for branch.

Details to be added. (13 September 2003, 16:06:57 CDT)

Branch Prediction References

03-11

Hybrid predictors.

McFarling 93: Scott McFarling, "Combining branch predictors," Digital Equipment Corporation WRL Technical Note TN-36, June 1993.

Chang 94: Po-Yung Chang, Eric Hao, Tse-Yu Yeh, and Yale Patt, "Branch classification: a new mechanism for improving branch predictor performance," in *Proceedings of the Proceedings of the 27th annual international symposium on microarchitecture*, November 1994, pp. 22–31.

Evers 96: Marius Evers, Po-Yung Chang, and Yale N. Patt, "Using hybrid branch predictors to improve branch prediction accuracy in the presence of context switches," in *Proceedings of the Proceedings of the 23th annual international symposium on computer architecture*, May 1996, pp. 3-11.

Skewed predictor.

Michaud 97: Pierre Michaud, Andr Seznec, and Richard Uhlig, "Trading conflict and capacity aliasing in conditional branch predictors," in *Proceedings of the Proceedings of the 24th annual international symposium on computer architecture*, 1997, pp. 292-303.

Bias Filtering Predictors

Lee 97: Chih-Chieh Lee, I-Cheng K. Chen, and Trevor N. Mudge, "The bi-mode branch predictor," in *Proceedings of the Proceedings of the Thirtieth Annual IEEE/ACM International Symposium on Microarchitecture*, December 1997, pp. 4–13.

03-12

eden: A. N. Eden and T. Mudge, "The YAGS branch prediction scheme," *International Symposium on Microarchitecture,* December 1998, pp. 69–77.

Sprangle 97: Eric Sprangle, Robert S. Chappell, Mitch Alsup, and Yale N. Patt, "The agree predictor: a mechanism for reducing negative branch history interference," in *Proceedings of the International Symposium on Computer Architecture,* June 1997, pp. 284–291.

History Filtering

Juan 98: Toni Juan, Sanji Sanjeevan, and Juan J. Navarro, "Dynamic history-length fitting: a third level of adaptivity for branch prediction," in *Proceedings of the International Symposium on Computer Architecture,* June 1998, pp. 155–166.