

Selective Topics in Micro-Architecture Design: Branch Prediction

Lei Chen

IEEE Sianal Processing Society



Outline

- Processor, Pipelining and the Branch Problem
- The "Solution": Branch Prediction
 - Design Examples
- The Journal of Instruction-Level Parallelism (JILP) Championship Branch Prediction Competition
- Branch Prediction and Others
- Branch Prediction Research and Design in Academia vs. Industry
- Issues in Branch Prediction and its Future
- Conclusions

The Basic 5-stage Pipeline

Pattterson and Hennessy, "Computer Organization and Design"



The Actual Processor



- Instruction Fetch
- Decode and Rename
 Writeback and
- Instruction Issue

- Instruction Execution
- Writeback and Commit

Two Real Processors

http://bwrc.eecs.berkeley.edu/CIC/die_photos/pentium.gif/, http://www.theinquirer.net/img/3911/nove3.jpg.





IEEE CTS ComSoc/SP Talk

The Branch Problem

- History
 - ENIAC
 - First electronic computer
 - Stretch (IBM 7030)
 - Pipeline processing



- More critical in superscalar and deeplypipelined processors
 - more branch instrs, higher mispred penalty
 - Performance bottleneck

The "Solution": Branch Prediction

- An educated guess of branch direction
 Taken or not-taken
- Static
 - fixed before program execution
 - adv: little hardware cost
 - disadv: low accuracy, application-specific
- Dynamic
 - adv: higher accuracy
 - disadv: hardware cost, power

Another "Solution": Branch Predication

- Allow instructions dependent on branches to conditionally update architecture state
- Basic idea:
 - Conditional branches set *predicate* registers
 - Dependant instructions continue execution but update architecture state based on the associated *predicate* register value
- Good for short branches, no misprediction but fixed overhead

Static Branch Prediction

- Always taken
 - Sun SuperSparc
- Always not taken
 - MIPS R4000, Intel i486
- Instruction specific
 - PowerPC 601
- Backward-taken-forward-not-taken:
 - HP PA-7x00
- Profile-based

Dynamic Branch Prediction

- 1-bit
 - Alpha 21064 instruction cache
- Bimodal (2-bit): 00, 01, 10, 11
 - Alpha 21064A, Alpha 21164 and PowerPC 604
- History-based
 - Gselect and Gshare
 - use global history
 - Local
 - use history from each branch
 - Path-based
 - further distinguish the paths leading to a branch

Branch Prediction Accuracy

Linley Gwennap, "New Algorithm Improves Branch Prediction"



IEEE CTS ComSoc/SP Talk

2-level Adaptive Branch Prediction



Local Branch Predictor



Combined Branch Predictor



14

Multi-Hybrid Predictor

Branch Target Buffer



June 15, 2009

The Aliasing Problem

- Finite indexing tables
- Destructive aliasing
 - different bias
- Non-destructive aliasing
 - same bias
- De-aliasing branch predictors
 Agree, Bi-Mode, Skewed, Filter, and YAGS

Agree Branch Predictor



IEEE CTS ComSoc/SP Talk

Bi-Mode Branch Predictor



Skew Branch Predictor



Filter Branch Predictor



2Bc-gskew Branch Predictor (EV8)



Perceptron Branch Predictor



Branch Prediction Using Register Values

- PREG prediction
 - Using profile information to find the branches correlated with data values
- Branch difference predictor
 - Maintains the history of differences between branch source registers
 - backing predictor, rare event predictor
- Branch prediction through value prediction

Branch Prediction through Value Prediction



Branch Prediction Using Data Dependence Information

- ARVI branch predictor
 - Available Register Value Information
 - Data dependence -> branch calculation
- Dynamic Dataflow-based Identification of Correlated Branches
 - Identify correlated branches from a large global history
- Multiple-cycle overriding schemes

Intuition Behind ARVI



Branch Outcome

Branch Dependence Chain

• The outstanding data dependence chain leading to the branch



June 15, 2009

IEEE CTS ComSoc/SP Talk

DDT Configuration



















Register Set Extractor (RSE)



ARVI Operations



IEEE CTS ComSoc/SP Talk

ARVI Access Timing



IEEE CTS ComSoc/SP Talk

Outline

- Processor, Pipelining and the Branch Problem
- The "Solution": Branch Prediction
 - Design Examples
- The Journal of Instruction-Level Parallelism (JILP) Championship
 Branch Prediction Competition
- Branch Prediction and Others
- Branch Prediction Research and Design in Academia vs. Industry
- Issues in Branch Prediction and its Future
- Conclusions

JILP Championship Branch Prediction Competition (CBP)

- Objective: Compare branch predictors in a common framework
- Held in conjunction with MICRO
- CBP-1: 2004, CBP-2: 2006
- Result:
 - Most are adapted perceptron predictor
 - A great deal of design engineering
 - Submissions from all around the world

JILP Championship Branch Prediction Competition (CBP)

- Different branches need to be predicted by different schemes
 - Bias and profiling
 - Loop, call/return
- History length
 - Global
 - Variable length
- Update scheme is very important
- No other information utilized
 - static information about the instructions, data values, and memory addresses.

Outline

- Processor, Pipelining and the Branch Problem
- The "Solution": Branch Prediction
 - Design Examples
- The Journal of Instruction-Level Parallelism (JILP) Championship Branch Prediction Competition
- Branch Prediction and Others
- Branch Prediction Research and Design in Academia vs. Industry
- Issues in Branch Prediction and its Future
- Conclusions

Branch Prediction and Others

- Data compression and information theory
 - Partial matching, Markov predictor
 - Entropy
- Signal processing
 - FAB: Fourier Analysis Branch
- Analog circuits
 - Digital-to-analog conversion for neural networks -> our brains?
- Hacking Intel branch prediction
- Data encryption
 - A spy-process running simultaneously with an RSA-process -> resource sharing
 - Square-and-Multiply Algorithm (SM) -> branches -> measure mispredictions -> key

Branch Prediction and Others Branch Prediction under Scrutiny for Possible Security Flaw

Branch predictors allow processors to execute the next instructions without waiting for the previous ones to be resolved. These predictors are crucial for achieving high performance. But a recently published research suggests that this CPU feature might as well encompass some security weakness. A branch prediction specialist at Irisa/INRIA research center, André Seznec explains what it is all about.

FOR I=1 TO M SEQUENCE X // 1,000s aycles IF KEY[1]=1 THEN SEQUENCE Y // 1000s and ENDFOR 2000s and

IEEE CTS ComSoc/SP Talk

Outline

- The Base Processor and Base Pipeline
- The Branch Problem
- The "Solution": Branch Prediction
 - Design Examples
- The Journal of Instruction-Level Parallelism (JILP) Championship Branch
 Prediction Competition
- Branch Prediction and Others
- Branch Prediction Research and Design in Academia vs. Industry
- Issues in Branch Prediction and its Future
- Conclusions

Branch Prediction in Academia

ISCA Branch Prediction Papers



Branch Prediction in Industry

USPTO Granted "Branch Prediction" Patents



June 15, 2009

IEEE CTS ComSoc/SP Talk

Branch Prediction in Industry

USPTO Granted "Branch Prediction" Patents



IEEE CTS ComSoc/SP Talk

Issues in Branch Prediction and Its Future

- Aliasing still a problem*
 - Time and space
- Global Branch History Length*
 - Longer and variable (CBP)
- Branch Predictor Delay
 - Overriding scheme
- Power Issues*
- Systematic design approach
 - Target prediction

Conclusions

- The branch problem is due to pipeline processing and it will become more critical
- Static and dynamic branch prediction schemes have been proposed to solve the branch problem
- Branch prediction is still unsolved in industry and requires careful systematic engineering



Selective Topics in Micro-Architecture Design: Branch Prediction

Lei Chen

IEEE Sianal Processing Society

