Autocorrelation Analysis: A New and Improved Method for Measuring Branch Predictability

Jian Chen, Lizy K. John
Department of Electrical and Computer Engineering
The University of Texas at Austin, Austin, Texas, USA
chenjian@mail.utexas.edu, ljohn@ece.utexas.edu

ABSTRACT

Branch taken rate and transition rate have been proposed as metrics to characterize the branch predictability. However, these two metrics may misclassify branches with regular history patterns as hard-to-predict branches, causing an inaccurate and ambiguous view of branch predictability. This study uses autocorrelation to analyze the branch history patterns and presents a new metric *Degree of Pattern Irregularity (DPI)* for branch classification. The proposed metric is evaluated with different branch predictors, and the results show that DPI significantly improves the quality and the accuracy of branch classification over traditional taken rate and transition rate.

Categories and Subject Descriptors: C.4 [**Performance of Systems**][Measurement techniques, Modeling techniques]
General Terms: Measurement, Performance
Keywords: Branch characterization, Autocorrelation

1. INTRODUCTION

Classifying branches in terms of their predictability has been applied in many areas of computer architecture, including branch prediction, predicated execution, etc. The existing metrics for characterizing branch behaviors include branch taken rate [2], which measures the taken frequency of a branch, and branch transition rate [3], which captures the frequency of a branch switching between taken and not taken. These metrics characterize the branch predictability based on their values: branches with very high or very low taken rate/transition rate are easy to predict, and branches with near 50% taken rate/transition rate are hard to predict. Although simple, these metrics may misclassify some of the easy-to-predict branches as hard-to-predict. For instance, a branch with regular history pattern "110110110..." is indeed easy to predict since a 3-bit history length is sufficient to make a perfect prediction for this branch. However, this branch has 0.667 taken rate and 0.667 transition rate, and is therefore misclassified as a hard-to-predict branch. To address this limitation, this work proposes to characterize branch predictability with a novel metric called *Degree of Pattern Irregularity* (DPI), which measures the degree of the branch behavior deviating from regular pattern by *autocorrelation analysis*. We show that DPI significantly improves the accuracy of hard-to-predict branch classification compared with taken rate and transition rate.

2. BRANCH AUTOCORRELATION

Autocorrelation is widely applied in signal processing and pattern recognition to find repeating patterns buried under noise. For a real-value discrete sequence of n elements $\{h(i)\}_{i=0}^{n}$, the autocorrelation of this sequence is $R_{hh}(j) = \sum_{i=0}^{n} h(i)h(i-j)$, where $j \in [0, n]$. In order to prevent undefined values outside the window $[0, n]$ from polluting the calculation, the sequence $\{h(i)\}_{i=0}^{n}$ is typically extended periodically to the left, creating a rotation effect in the window $[0, n]$ as the sequence slides to the right. Therefore, the autocorrelation holds the following two properties [1]: (a) It reaches its maximum value at the origin; (b) If the discrete sequence is periodic, its autocorrelation is also periodic with the same period. Considering the fact that branch history only consist of "0" (as not taken) and "1" (as taken), we have the following implication: for the autocorrelation of branch histories, the difference between the maximum value at the origin and the largest value off the origin reflects the amount of irregularity in the branch history. This can be understood by treating an irregular branch history as a regular branch history XORed with one or more bits deviating from the regular pattern. The number of these deviating bits is reflected on the difference between the two largest values of the autocorrelation. As shown in Figure 1(b), one bit highlighted with dark grey deviates from the periodic pattern, which causes the difference between the two largest autocorrelation values equivalent to one. The irregularity measured by such difference is one of the main sources of branch misprediction. Hence, the fraction of the irregularity over the number of the branch dynamic accesses is the direct indicator of branch predictability, which we refer to as the *Degree of Pattern Irregularity* (DPI). Note that a regular branch history is equivalent to a branch history with zero DPI, as shown in Figure 1(a).

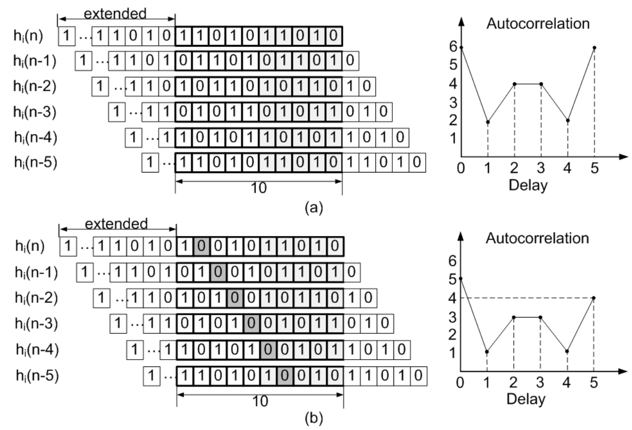

Figure 1: (a) Autocorrelation for regular branch history.(b) Autocorrelation for irregular branch history.

The complexity of the autocorrelation analysis involves two aspects: computation and storage. Since branch history only contains 0's and 1's, its autocorrelation only requires logic AND operations and bit-wise accumulate operations. Compared with the storage requirements of taken rate (1 bit per static branch) and transition rate

SIGMETRICS'11, June 7–11, 2011, San Jose, California, USA.
ACM 978-1-4503-0262-3/11/06.

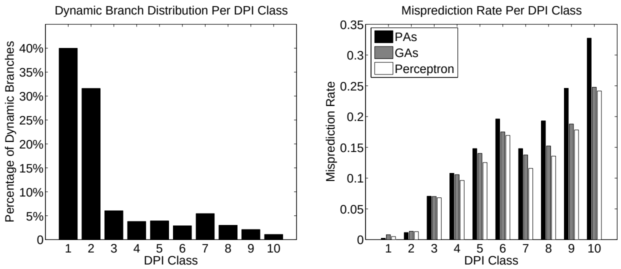

(a) (b)

Figure 2: Dynamic branch classification based on DPI.

(2 bits per static branch), autocorrelation analysis requires more storage space, yet its impact on the profiling speed is negligible as long as the history length is within a reasonable range.

3. EXPERIMENT AND RESULTS

We use PIN, a dynamic instrumentation tool on x86 platform, to instrument the workload and obtain the trace of conditional branches. This trace is then seamlessly fed to our detailed branch analyzer, which is able to perform autocorrelation analysis on each static branch and simulate different types of branch predictors simultaneously. The workloads of the experiment are composed of all programs from SPEC CPU2006 benchmark suite, with each compiled to x86-ISA at base configurations. To reduce the simulation time, we use PinPoints to identify the representative simulation points. For each program, we simulate the dominant simulation points that covers 90% of the total weights, and each simulation point contains 100 million instructions.

We evaluate the proposed metric by using three different types of branch predictors to ensure the generality. These three branch predictors are: a per-address history predictor (PAs), a global two-level predictor (GAs) and a global neural network predictor (Perceptron) [4], each with history length of 16. For PAs and GAs, the size of Pattern History Table (PHT) is set to 64K entries, and the branch history table (BHT) of PAs has 1024 entries. To be consistent with PAs and GAs, the Perceptron predictor also contains 64K entries for the weights with each 8-bit wide. In this work, we only consider the conditional branches.

Branch Classification: In Figure 2(a), we classify the branches into 10 groups in terms of their DPI values. Class 1 has DPI value 0, representing the branches with regular history pattern. Class 2 to 6 have DPI values in the ranges of (0,0.01], (0.01,0.02], ..., (0.04,0.05], respectively; and class 7 to 10 have DPI values with the ranges of (0.05-0.10], (0.10-0.15], (0.15-0.20], (0.20-1] respectively. As shown in the figure, 40.0% of the total dynamic conditional branches fall in class 1, and 31.6% of them fall in class 2. The occupancies of the other classes are significantly lower, with each class less than 6.0%. Figure 2(b) further shows the misprediction rate of the branches in each DPI class for PAs, GAs, and Perceptron predictors. Notice that there is an overall trend that the misprediction rate increases as the branch DPI increases. This trend holds true for all three different types of branch predictors, which demonstrates that DPI is an appropriate metric for branch predictability. Moreover, this figure also shows that the misprediction rates of the branches in DPI class 1 and 2 are drastically smaller than those in the rest DPI classes, which means branches with DPI less than 0.01 are the easy-to-predict branches. As a result, DPI allows us to classify the branch predictability in a clear and coherent way: branches with DPI less than 0.01 are the easy-to-predict branches; whereas branches with DPI larger than 0.01 are the hard-to-predict branches.

Comparison with Conventional Metrics: Figure 3(a) shows the percentage of the branches classified as easy-to-predict branches among the branches with prediction rate larger than 95% (The easy-to-predict branches are classified by taken rate $\in [0, 0.05) \bigcup (0.95, 0.1]$, transition rate $\in [0, 0.1) \bigcup (0.9, 0.1]$, or DPI $\in [0, 0.01]$). As shown in this figure, DPI consistently yields larger percentage than transition rate or taken rate across all three types of branch predictors, meaning that DPI can identify more truly easy-to-predict branches than taken rate or transition rate. On the other hand, we also measure the percentages of the branches with prediction rate less than 95% over the branches classified as hard-to-predict. As shown in Figure 3(b), DPI improves the accuracy of the hard-to-predict branch classification by up to 17.7% over taken rate, and 15.0% over transition rate. The reason that DPI is superior in branch classification is that it has a broader view of branch history when characterizing the branch behaviors. In fact, taken rate examines the branch history *bit by bit*, and transition rate does it *two-bit by two-bit*; whereas DPI examines the branch history at a *broader pattern level*.

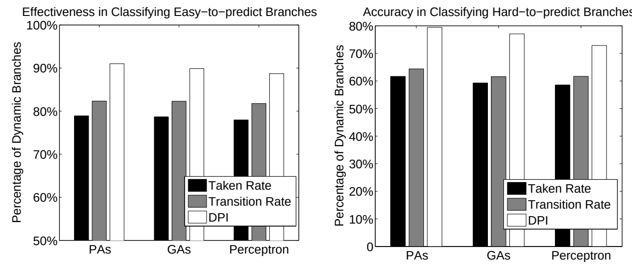

(a) (b)

Figure 3: Comparison of branch classification quality.

Applications: As an important extension to the existing metrics, the proposed DPI metric can be applied in the fields where the conventional branch classification metrics are used. These fields include, but not limited to: identifying hard-to-predict branches for predication, characterizing control flow for benchmark cloning and synthesizing [5].

4. CONCLUSIONS

Based on the autocorrelation analysis of branch history patterns, this paper presents a new metric *Degree of Pattern Irregularity* (DPI) for branch predictability characterization. Unlike existing taken rate or transition rate metrics, DPI directly measures the regularity of the patterns in per-address branch history, and hence is able to identify more easy-to-predict branches and significantly improve the accuracy of the classification of hard-to-predict branches. Our experiments show that DPI improves the accuracy of hard-to-predict branch classification by up to 17.7% over taken rate and 15.0% over transition rate. Overall, this metric examines the branch history at a broader *pattern level*, and is an important extension to the existing metrics in branch classification.

5. REFERENCES

[1] E. O. Brigham. *The Fast Fourier Transform*, chapter 13. 1974.

[2] P.-Y. Chang and et al. Branch classification: a new mechanism for improving branch predictor performance. In *MICRO '94*, pages 22–31, 1994.

[3] M. Haungs, et al. Branch transition rate: a new metric for improved branch classification analysis. In *HPCA '00*, pages 241 –250, 2000.

[4] D. Jimenez and C. Lin. Dynamic branch prediction with perceptrons. In *HPCA '01*, pages 197 –206, 2001.

[5] A. Joshi, et al. Automated microprocessor stressmark generation. In *HPCA '08*, pages 229 –239, 2008.