Journal of Instruction-Level Parallelism 7 (2005) 1-4 Submitted 1/05; published 4/05

Introduction to JILP's Special Edition for Finalists of the Championship Branch Prediction (CBP1) Competition

Chris Wilkerson CHRIS.WILKERSON@INTEL.COM
Jared Stark JARED.W.STARK@INTEL.COM
Jones Farm 3, Hillsboro, OR

1. Background

CBP is a branch prediction competition in which the contestants were given a fixed storage budget to implement their best branch prediction algorithms on a common evaluation framework distributed by the CBP steering committee. This first CBP competition was held in 2004. It was formally announced at ISCA-31 and concluded at MICRO-37. This issue of JILP contains papers describing the predictors of all the finalists in this competition.

The goal of the CBP competition was to evaluate and compare the performance of branch prediction algorithms in a common framework. The competition's evaluation process was intended to be simple and transparent, enabling the dissemination of results and techniques to the larger computer design community and allowing independent verification of the competition's results. The performance and cost metrics were selected to be as simple and quantitative as possible.

2. Framework & Traces

The framework was designed to be general enough that all proposed predictors could be evaluated on it while at the same time providing room for innovation. The framework consists of two sets of traces and a driver.

The traces were x86 instruction traces, were 30 million instructions long, and included both user and system activity. There were two sets of traces, a distributed and an undistributed set. The distributed set was distributed to the contestants, and was used in the initial round of the competition to select the finalists. The undistributed set was not distributed to the contestants, and was used in the final round of the competition to rank the finalists and determine the champion. (More information on the competition rounds is provided below.) There were 20 traces per set: 5 integer, 5 floating-point, 5 multimedia, and 5 server.

The driver read one instruction at a time from a trace and provided it to a contestant's predictor. It provided both static and dynamic information. The static information included instruction PCs, types (e.g., ALU, load, store, or branch), and register source and destination specifiers. The dynamic information included instruction results, load/store addresses, and branch outcomes. In a real processor, this information is not immediately known to the predictor due to the processor's pipeline. The CBP competition accounted for this by providing the dynamic information to the predictor on a delayed basis; e.g., an instruction's result wasn't available to the predictor until N instructions after the predictor had been given the instruction.

3. The Competition

Contestants were given a budget of (64K + 256) bits to implement their predictors. Other costs, such as power, latency, and design effort, are important when designing a predictor, but are difficult to quantitatively assess in a reasonable amount of time. For these costs, the CBP selection committee reviewed the predictors and disqualified those for which these costs were judged to be absurd; that is, they were so high the predictor could never be useful in a real processor design. However, none of the submitted predictors were found to be absurd, so none were disqualified.

Predictors were evaluated only on the direction prediction (taken/not-taken) for conditional branches. Mispredict rates were measured in mispredicts per 1000 instructions, measured over an entire trace. Predictor performance was measured as the (arithmetic) average of the mispredict rates over all traces in a trace set.

The competition was held in two rounds. In the initial round, six finalists were selected from a total of eighteen submissions using predictor performance on the distributed trace set as the selection criterion. Although we originally planned to select only five finalists we found the performance of the top six submissions to be highly clustered. There was roughly a 10% performance difference between the first place and sixth place submission and the sixth place submission performed 10% better than the seventh place submission. In light of this, the selection committee opted to select the top six submissions as finalists.

4. The Finalists

Each of the finalists selected in the initial round were reviewed by the selection committee to determine if the submission adhered to the spirit of the competition. All 6 submissions were accepted in this round, although the selection committee decided to create an additional "Best Practices Award" to give special recognition to a submission that the committee felt balanced practicality and performance. Finally the 6 finalist submissions were evaluated using the undistributed trace set and ranked according to their performance. It's worth noting that although the ordering of the finalists changed when run on the undistributed trace set, when all submissions were run on the undistributed trace set we found that the same six finalists would have made the cut.

We were very pleased with the quality of the submissions. The average CBP finalist performed 45% better than a gshare of the same size. Five of the six finalists used perceptrons or were perceptron inspired; Pierre Michaud's PPM-like predictor was the only exception. Each finalist introduced new techniques to improve prediction accuracy. Each technique was, for the most part, only used in one of the finalists' predictors, but could be adapted to work with other predictors. It's also worth noting that despite the fact that the framework permitted the use of data values in the predictor design only 1 submission (and none of the finalists) used data values. In fact, all six of the finalists achieved higher performance through the use of longer histories.

The following table shows the finalists, their final ranking, and their final scores (i.e., average mispredict rates over the undistributed trace set, in mispredicts per 1000 instructions.) Gshare is also shown for comparison. More information about the finalists can be found by following the "Workshop Agenda and Results" link on the CBP web site, http://www.jilp.org/cbp/, including links to their presentations, predictor writeups, and predictor source code.

Table 1: List of Finalists

Rank	Finalist	Average Score
1.	***Adaptive Information Processing: An Effective Way to Improve Perceptron Predictors*** Hongliang Gao and Huiyang Zhou, *University of Central Florida*	2.574
2.	***The O-GEHL Branch Predictor***, André Seznec, *IRISA/INRIA*	2.627
3.	***The Frankenpredictor***, Gabriel Loh, *Georgia Institute of Technology*	2.700
4.	***Idealized Piecewise Linear Branch Prediction***, Daniel A. Jiménez, *Rutgers University*	2.742
5.	***A PPM-like, Tag-based Predictor***, Pierre Michaud, *IRISA/INRIA*	2.777
6.	***A 2bcgskew Predictor Fused by a Redundant History Skewed Perceptron Predictor***, Veerle Desmet, Hans Vandierendonck, and Koen De Bosschere, *Ghent University*	2.807
	GSHARE	4.520

Overall we feel that the hard work of the participants made the CBP a great success. We'd also like to express our appreciation to the steering committee members: Dan Connors, *Univ. of Colorado;* Tom Conte, *North Carolina State Univ;* Konrad Lai, *MRL, Intel;* Yale Patt, *Univ. of Texas at Austin;* Jim Smith, *Univ. of Wisconsin;* Jared Stark, *MRL, Intel;* Mateo Valero, *Univ. Politecnica Catalunya;* and Chris Wilkerson, *MRL, Intel.* In particular we'd like to thank Tom Conte, Konrad Lai, and Dan Connors.

We'd also like to thank the members of the selection committee: Dan Connors, *Univ. of Colorado;* Tom Conte, *North Carolina State Univ;*Phil Emma, *IBM Research;* Konrad Lai, *MRL,Intel;* Scott McFarling, *Microsoft (now at Intel);* Chuck Moore, *AMD;* Yale Patt, *Univ. of Texas at Austin;* Jim Smith, *Univ. of Wisconsin;* Jared Stark, *MRL, Intel;* Mateo Valero, *Univ. Politecnica Catalunya;* and Chris Wilkerson, *MRL, Intel.*

We met our goals of evaluating the world's best branch predictors in a quantitative, transparent process. The innovation in the submissions surpassed our expectations. Given the success of this model, we may organize a CBP-2 to be held in 2006. In addition, we're interested in applying this model to other areas of research. We hope to hold a prefetching competition in the second half of 2005.