Controlling the Power and Area of Neural Branch Predictors for Practical Implementation in High-Performance Processors*

Daniel A. Jiménez
Rutgers University
Department of Computer Science
djimenez@cs.rutgers.edu

Gabriel H. Loh
Georgia Institute of Technology
College of Computing
loh@cc.gatech.edu

Abstract

Neural-inspired branch predictors achieve very low branch misprediction rates. However, previously proposed implementations have a variety of characteristics that make them challenging to implement in future high-performance processors. In particular, the the original Perceptron branch predictor suffers from a long access latency, and the faster path-based neural predictor (PBNP) requires deep pipelining and additional area to support checkpointing for misprediction recovery.

The complexity of the PBNP predictor stems from the fact that the path history length, which determines the number of tables and pipeline stages, is equal to the history length, which is typically very long for high accuracy. We propose to decouple the path-history length from the outcome-history length through a new technique called modulo-path *history. By allowing a shorter path history, we can implement a PBNP with significantly fewer tables and pipeline stages while still exploiting a traditional long branch outcome history. The pipeline length reduction results in decreased power and implementation complexity. We also propose* folded modulo-path *history to allow the number of pipeline stages to differ from the path history length. We show that our modulo-path PBNP at 8KB can achieve prediction accuracy and overall performance within 0.8% (SPECint) of the original PBNP while simultaneously reducing predictor energy consumption by ∼29% per access and predictor die area by ∼35%. Our folded modulo-path history PBNP achieves performance within 1.3% of ideal, with a ∼37% energy reduction and ∼36% predictor area reduction.*

1. Introduction

After decades of academic and industrial research efforts focused on the branch prediction problem, pipeline flushes due to control flow mispredictions remain one of the primary bottlenecks in the performance of modern processors. A large amount of recent branch prediction research has centered around techniques inspired and derived from machine learning theory, with a particular emphasis on the *perceptron* algorithm [4, 5, 9, 10, 12, 15, 25]. These neural-based algorithms have been very successful in pushing the envelope of branch predictor accuracy.

Researchers have made a conscious effort to propose branch predictors that are highly amenable to pipelined and ahead-pipelined [22] organizations to minimize the impact of predictor latency on performance. There has been considerably less effort on addressing power consumption and implementation complexity of the neural predictors, both of which are now first-class design considerations in high-performance microprocessors. Reducing branch predictor power is not an easy problem because any reduction in the branch prediction accuracy can result in an overall increase in the *system* power consumption due to a corresponding increase in wrong-path instructions [3].

The goal of this work is to demonstrate that the complexity of neural branch predictors can be substantially reduced without altering the fundamental behaviors and characteristics of the prediction algorithm. In this paper, we will *not* re-argue the benefits of the conventional neural predictors as that has already been demonstrated in numerous previous works [4, 5, 9, 10, 12, 25].

In the rest of this paper, we will first review the design of neural-based branch predictors in Section 2. Section 3 details our new branch predictor organizations targeted at reducing complexity and overall energy consumption. Section 4 presents our experimental results demonstrating the impact of our proposal on processor performance and energy, and Section 5 concludes the paper.

2. Neural Branch Predictors

In this section, we review neural-based branch predictors. We then qualitatively describe the sources of complexity and power consumption that make conventional neural predictors difficult to implement in high-performance processors.

*A subset of the ideas presented in this paper appeared in an earlier workshop with unpublished proceedings.

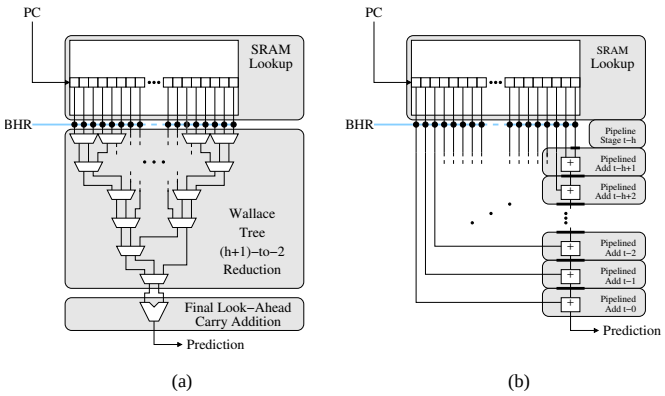

Figure 1. (a) The perceptron branch predictor and (b) the path-based neural predictor or PBNP.

2.1. Background

Traditional PHT-based (pattern history table) branch predictors such as GAs [26] and gshare [16] do not scale gracefully with longer branch history lengths. For h bits of branch history, a conventional PHT needs a table size exponential in h. The neural predictors are interesting because they can exploit deep correlations from very long history lengths with subexponential scaling.

The basic perceptron predictor [12] employs a vector of weights that learns correlations between the branch direction and the results of previous branches. Figure 1a shows how a table of weights is indexed by the program counter (PC) to choose a single vector of weights, which is then combined with the branch history register in a dot-product operation (each ● represents conditionally negating the weight depending on the direction of the corresponding branch history bit). Conceptually, a past branch that is strongly correlated with the outcome of the current branch will have a corresponding weight with a large magnitude. The perceptron trains (increments/decrements) the weights to predict only according to those branches in the history that have exhibited strong correlations to the branch under consideration. A Wallace tree reduces the $h + 1$ weights down to only two weights in $O(\log_{3/2} n)$ carry-save adder gate delays. A final carry-completing adder such as a look-ahead carry adder computes the final sum. The sign of this resulting sum indicates the final prediction. The main obstacle to implementing a perceptron branch predictor is the long latency required to read the weights and then perform the large dot-product operation. To reduce the latency of the predictor, it may be necessary to implement the adders with fast, leaky transistors that end up consuming more dynamic and static power.

The second-generation path-based neural predictor (PBNP) largely solves the latency problems of the original perceptron [9]. The central idea is that the i^{th} previous branch address (i.e. from the path history) can be used to look up the weight corresponding to the i^{th} oldest branch i cycles ahead of time. While this largely addresses the latency issues of the perceptron predictor, the PBNP still suffers from significant implemenation complexity. As shown in Figure 1b, the clever pipelining of the PBNP provides a much faster effective predictor latency: the critical path is now the table lookup and a single addition. The PBNP also requires a number of adders equal to the depth of the branch history, further increasing the hardware cost. The third-generation piecewise linear predictor has an organization that consists of k parallel PBNP pipelines that can each learn separate linearly separable functions for better prediction accuracy [10]. However, the k pipelines further increases the implementation overhead.

2.2. Power and Complexity

In this paper, we will focus on the path-based neural predictor (PBNP). The original perceptron predictor's long lookup latency makes it difficult to implement without making use of an overriding predictor organization [11] that just adds more power and complexity. The piecewise linear predictor has a very similar structure, and so we believe our findings based on the PBNP hold for the more complicated predictor as well.

During the lookup phase of the PBNP, each pipeline stage reads a weight corresponding to the exact same PC. This is due to the fact that the current PC_0 will be the next branch's PC_1 and next-next branch's PC_2 and so on. This allows an implementation where the weights are read in a single access using a single large SRAM row that contains all of the weights. During the update phase however, a single large access would force the update process to use a pipelined implementation as well. While at first glance this may seem desirable, this introduces considerable delay between update and lookup. For example a 30-stage update pipeline implies that even after a branch outcome has been determined, another 30 cycles must elapse before the PBNP has been fully updated to reflect this new information. This update delay can create a decrease in predictor accuracy. There are also some odd timing effects due to the fact that some weights of a branch will be updated before others.

An alternative organization uses h tables in parallel, one for each pipeline stage/history-bit position [9], as shown in Figure 2(a). This organization allows for a much faster update and better resulting accuracy and performance. The disadvantage of this organization is that there is now a considerable amount of area and power overhead to implement the row decoders for the h separate SRAM arrays. Furthermore, to support concurrent lookup and update of the predictor, each of these SRAM arrays needs to be dual-ported (one read port/one write port) which further increases the

area and power overhead of the SRAM row decoders. To use the PBNP, the branch predictor designer must choose between an increase in power and area or a decrease in prediction accuracy. Using a large number of small SRAM arrays makes it more difficult to derive energy savings through SRAM banking and sub-banking techniques [6].

On a branch misprediction, the PBNP pipeline must be reset to the state that corresponded to the mispredicting branch being the most recent branch in the branch and path history. To support this predictor state recovery, each branch must checkpoint all of the partial sums in the PBNP pipeline. On a branch misprediction, the PBNP restores all of the partial sums in the pipeline using this checkpointed state. For b-bit weights and a history length of h, a PBNP checkpoint requires approximately bh bits of storage. The total number of bits is slightly greater because the number of bits required to store a partial sum increases as the sum accumulates more weights. The total storage for all checkpoints corresponds to the maximum number of in-flight branches permitted in the processor. The checkpointing overhead represents additional area, power, and state that is often unaccounted for in neural predictor studies. This overhead increases with the history/path-length of the predictor since the PBNP must store one partial sum per predictor stage. The combination of these issues makes the conventional PBNP difficult to implement for high-performance.

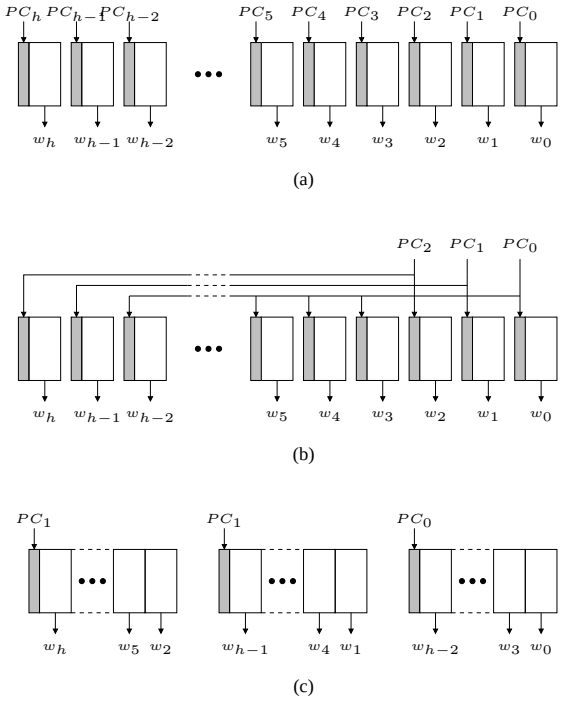

Figure 2. (a) Organization of the h tables of the PBNP, (b) logical organization of a PBNP using modulo path-history for $P = 3$, and (c) the corresponding physical organization of the same. The shaded portion represents the SRAM row decoder and related access logic.

3. Decoupling the Path and Branch History Lengths

In the original PBNP, the path history length is always equal to the branch history length. This is a result of using PC_i to compute the index for the weight of x_i. As described in the previous section, the pipeline depth directly increases the number of tables and the checkpointing overhead required. On the other hand, supporting a long history length requires the PBNP to be deeply pipelined.

3.1. Modulo-Path History

We propose *modulo path-history* where we decouple the branch history length from the path history length. We limit the path history to only the $P < h$ most recent branch addresses. Instead of using PC_i to compute the index for w_i, we use $PC_{i \mod P}$. In this fashion, we can reduce the degree of pipelining down to only P stages. Figure 2(b) shows the logical organization of a PBNP using modulo path-history (for $P = 3$). In this example, we only use a path length of three, but Figure 2(b) still appears to use $O(h)$ separate tables. Since every P^{th} weight is indexed with the same branch address, we can interleave the order of the weights in the tables such that only P tables are necessary. Figure 2(c) shows the physical organization where each table provides weights that correspond to h/P branch history outcomes, where each branch history outcome is separated by P bit positions.

By reducing the PBNP implementation to only use P distinct tables, we address several of the main sources of power and complexity as described in Section 2. Using only P tables reduces the duplicated row-decoder overhead. The reduction in the number of tables reduces the overall pipeline depth of the predictor which reduces the total bits of state that must be checkpointed (i.e. there are only P partial sums). The number of inter-stage latches and associated clocking overhead is also correspondingly reduced.

3.2. Folded Modulo-Path History

Reducing the path-history length of a path-based neural predictor may reduce prediction accuracy because the long path history may provide additional context for detecting correlations. That is, a branch prediction may be highly correlated to the k^{th} address in the branch history, but reducing the path length to $P < k$ by using modulo-path history eliminates this source of correlation. To recapture the correlation in the k^{th} path address, we would need to increase

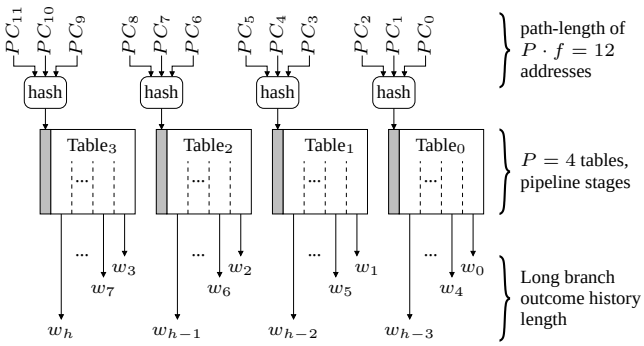

Figure 3. A neural predictor employing folded modulo-path history, where the path-length $\neq$ history-length $\neq$ number of tables.

the modulo-path history length to at least k (so indices are computed using $PC_{i \mod k}$ rather than $PC_{i \mod P}$). This unfortunately increases the predictor pipeline depth and the associated power and complexity.

We propose to decouple the predictor pipeline depth from the path-history length by using *folded modulo-path* history. Similar to the normal modulo-path history, our PBNP uses only P tables, but now we employ a path-history length that is $P \cdot f$ addresses long, where f is the *folding factor*. In a conventional PBNP we only use one branch address to compute an index for each table. With folded modulo-path history, we use f addresses hashed together. Folded modulo-path history can be considered the path-history analogue of the folded long branch outcome histories used in other predictors such as 2bc-gskew [21]. Figure 3 shows an example PBNP with $P = 4$ and $f = 3$ for a total path length of 12 while only using four tables.

To combine the f path addresses into a single index, we used a simple XOR and shift-based hash function. For each of the path addresses $PC_i, i \in \{0..f-1\}$ used to index a table, we hash the addresses by taking the exclusive-OR of $PC_i \ll i$. This is similar to the hash function used by Stark et al. for their path hashing [24].

Note that the folded-modulo-path history predictor needs a shift register to track the path-history, and this shift register will need to be checkpointed for misprediction recoveries. However, the size of this shift register is relatively small due to the reduced number of predictor stages enabled by the modulo history and each entry only needs to store enough bits to index into the perceptron table (e.g., 128-entry SRAM only require 7 bits from each branch address).

Modulo path-history is a unique way to manage the branch path history information. A PBNP can now choose between different lengths of branch and path history. Tarjan and Skadron proposed a "hashed" perceptron indexing scheme that removed the rigid relationship between history length and the number of tables [25]. Seznec's GEHL predictors use a similar hashing approach to map multiple bits of branch history to a single correlation weight [20]. Our work provides a different way of separating history length from table count, and also makes the contribution of separating the path length from either of these parameters as well.

3.3. Generality of the Techniques

In this paper, we focus on the path-based neural predictor. However, the proposed history-folding techniques can potentially be applied to other predictor organizations. Multi-table, ahead-pipelined predictors such as the Hashed-Perceptron [25], GEHL [20] or PPM [17] could all include folded history to incorporate additional information and context in their indexing functions.

4. Results

In this section, we present our experimental results to demonstrate the merits of our proposed branch predictor organizations.

4.1. Experimental Methodology

For our initial design space exploration, we used the in-order branch predictor simulator sim-bpred from the SimpleScalar toolset [2]. we simulated applications from the SPEC2000cpu integer benchmark suite with reference inputs, MiBench [8] with the large inputs and the MediaBench [14] multimedia benchmark suites with expanded inputs. Some applications (e.g., the `lame` MP3 encoder) are not included because we could not compile them in our Alpha environment due to unsupported libraries. We used 100M instruction simulation points chosen by SimPoint 2.0 [19]. Our applications were compiled on an Alpha 21264 with Compaq `cc` with full optimizations. For our IPC simulations, we used the MASE simulator from SimpleScalar 4.0 [13]. We simulated a four-wide out-of-order processor; the details are listed in Table 1. The processor parameters were chosen to model a machine with a level of aggressiveness similar to a Intel Pentium-III/Pentium-M microarchitecture.

We used CACTI 3.2 [23] to estimate the energy consumption of 90nm implementations of the branch predictors. Since CACTI does not simulate tables with non-power-of-two numbers of entries, we simply rounded-up the sizes of our structures to the next largest power of two. While this introduces some slight overestimation in power consumption, a realistic implementation of a neural branch predictor would likely use SRAMs with a power-of-two

Parameter	Value	Parameter	Value
Machine	4-wide	Integer Units	ALU:2, Mult:2
IFQ Size	8 entry	FP Units	Add:1, Mult:1, Div:1
Scheduler	24 entry	Latencies	Same as Pentium-M [7]
LSQ Size	24 entry	Memory Ports	2
ROB Size	64 entry	ITLB/DTLB	64 entry each
IL1, DL1	16KB/4-way	Branch Penalty	13 cycles
Unified L2	512KB/8-way	DRAM Latency	200 cycles

Table 1. The processor configuration used for our IPC simulations.

	Common Parameters		Modulo-Path	+ Folded Path	
Size	History Length	Rows per SRAM (# Perceptrons)	Path Length	Path Length	f = Path Folding
1KB	24	40	8	11	2
2KB	24	81	8	11	2
4KB	31	128	6	15	2
8KB	32	248	7	11	2

Table 2. Parameters for the baseline PBNP and versions using modulo path-history and folded-modulo-path history.

number of entries. We used CACTI to estimate the energy consumption of both the predictor tables of weights as well as the checkpoint tables. We also extrapolated the energy consumption of the predictors' adders based on the logic model of CACTI's row decoders. For our predictor die-area estimates, we use the register bit equivalent (rbe) methodology proposed by Mulder et al. [18].

We simulated a large number of PBNP configurations to find the best parameter settings. For the modulo-path and folded-modulo-path versions, we maintained the same branch history length and number of entries per table of weights (i.e., same total number of perceptrons) while allowing the predictor pipeline depth and path history length to vary. The final configurations are listed in Table 2. We could have potentially improved the performance of the modulo-path versions by allowing the history length and number of perceptrons to change. However, we decided to keep these parameters the same as the baseline PBNP to directly quantify the impact of our techniques.

4.2. Predictor Accuracy

The usage of modulo-path history potentially compromises the prediction accuracy of the path-based neural predictor (PBNP) due to the reduction in the total amount of unique path information. While the modulo-path versions of the PBNP make for simpler and more practical implementations, a substantial reduction in accuracy and overall performance would simply make both conventional *and* modulo-path versions of the predictor undesirable. Figure 4 shows the average prediction accuracy of the different versions of PBNPs across a range of predictor sizes for SPECint, MediaBench and MiBench, respectively. Overall, the modulo-path modifications only slightly increase the misprediction rates of the predictors, with a greater sensitivity at the smallest hardware budget. For future high-performance processors however, the predictor sizes are more likely to be toward 8KB or larger [1], making the sensitivity at the smaller sizes less of a problem. This is a very positive result as it means that we can employ the simpler modulo-path versions of the neural predictor without crippling performance.

4.3. Predictor Accuracy

Given that our modified versions of the PBNP do not affect prediction accuracy by much, we expect that the overall performance will also be similar to that of the original PBNP. Figure 5 shows the geometric mean IPC rates across the SPECint, MediaBench and MiBench applications, respectively, for different predictor sizes. Overall, the IPC of our simplified PBNPs matches the performance of the original predictor very closely. At an 8KB budget on the SPECint applications, the IPC degradation is only 0.8% for the ModPath predictor, and 1.3% for the Folded version. Across the range, the difference in overall performance is within the noise of the simulator. For the MediaBench applications, we observe 1.5% and 0.5% IPC degradations for the ModPath and Folded-ModPath 8KB predictors, respectively. For MiBench, the performance penalties are 0.1% and 0.0% for the same predictors.

The choice of target applications affects which predictor organization is most appropriate. For example, ModPath performs better for SPECint, while the Folded ModPath is better for MediaBench. On the other hand, the MiBench applications are fairly insensitive to the choice of the neural predictor implementation, which means we can reap energy and area benefits (described in the next section) without *any* performance impact.

It is important to keep in perspective that even though the modulo-path history versions of the PBNP cause a slight performance drop, this is relative to a processor that uses a conventional PBNP. Without our proposed modifications, a processor would not even be able to use the PBNP in the first place. Our contribution is a new design for the PBNP that makes it much more practical while delivering nearly the same benefit as the more complex version. The alternative is a processor with much less performance due to a less sophisticated non-neural prediction algorithm [12].

4.4. Energy Impact

We anticipate that the modulo-path PBNPs will consume less energy per access for two reasons. The first is that the folded organization reduces the total number of SRAM arrays which reduces the per-table overhead such as the row decoders. However, packing the same number of bits into

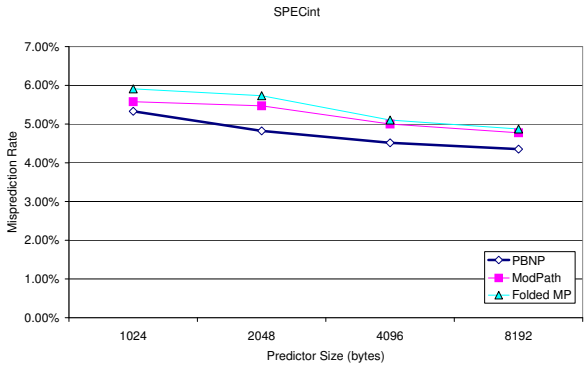

(a)

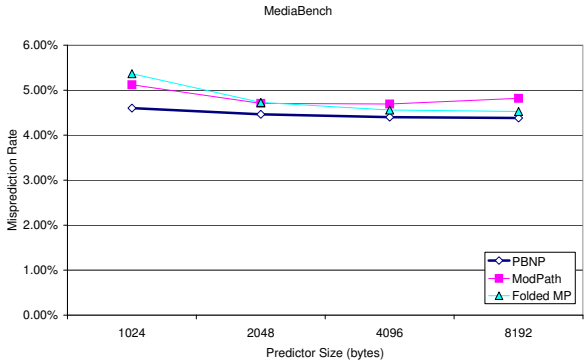

(b)

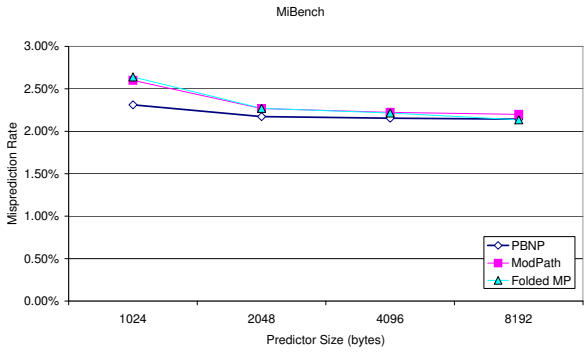

(c)

Figure 4. Arithmetic mean misprediction rates across the (a) SPECint, (b) MediaBench and (c) MiBench applications for the original path-based neural predictor (PBNP) as well as versions using modulo- and folded-modulo-path history.

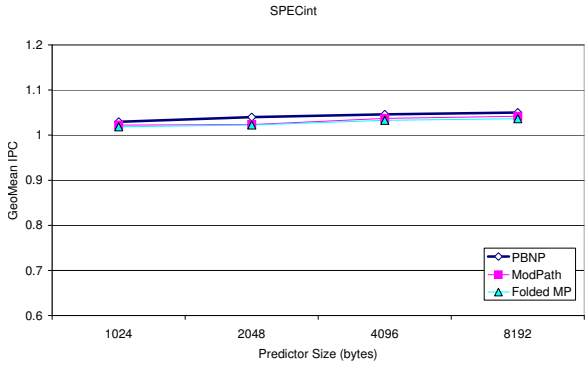

(a)

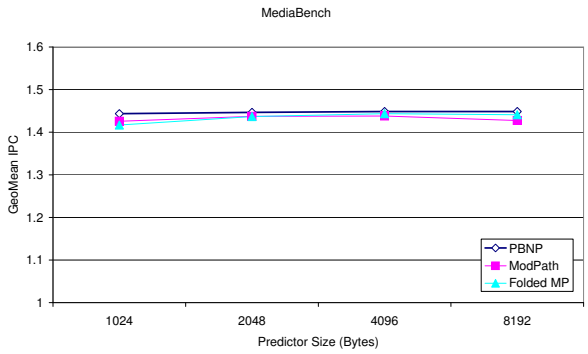

(b)

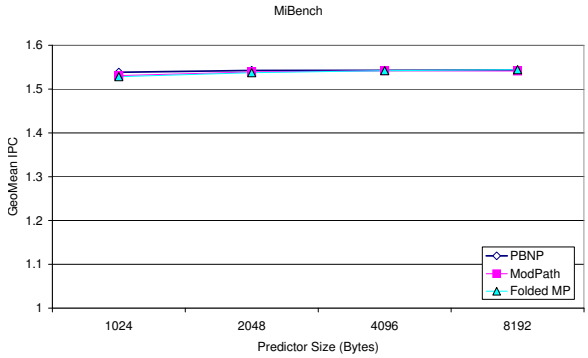

(c)

Figure 5. Geometric mean IPC rates for the (a) SPECint, (b) MediaBench and (c) MiBench applications for the original path-based neural predictor (PBNP) as well as versions using modulo- and folded-modulo-path history.

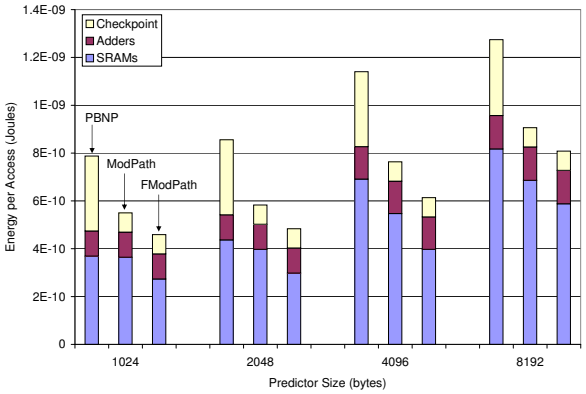

Figure 6. Energy per access including the lookups in the SRAMs for the predictor weights, the adders for computing the predictor output, and one checkpoint access.

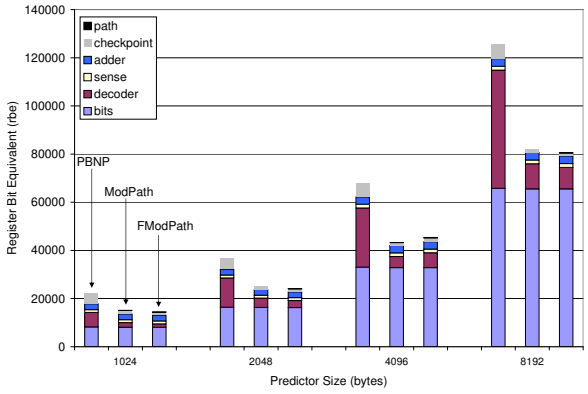

Figure 7. Area requirements of the different versions of PBNP in register bit equivalents (rbe).

a smaller number of tables tends to increase the wordline lengths, which could increase power. The second reason for an energy reduction is in the decrease of the checkpointing overhead. By reducing the total number of SRAM tables, we reduce the length of the predictor pipeline, thereby requiring less state to be checkpointed for each branch. Figure 6 shows the overall energy consumption per predictor access, which includes the predictor portion (SRAMs and adders) as well as the checkpointing overhead. The adders' energy consumption does not vary because we maintain the same history length between the conventional PBNP and the modulo-path versions. The results clearly show a substantial reduction in the predictor energy consumption, ranging from 30% and 42% for 1KB modulo-path and folded-modulo-path versions, respectively, to 29% and 37% for the 8KB configurations.

4.5. Area

To estimate the die area of each predictor configuration, we use the process independent register bit equivalent (rbe) metric proposed by Mulder et al. [18]. The rbe methodology provides for a way to estimate the area overhead of the decoder/driver logic, sense amps, and other related circuitry. Figure 7 shows the estimated areas of the predictor configurations in rbe's. The majority of the area reduction comes from a reduction in the total number of SRAM tables which reduces the overhead of duplicated decode logic (the total number of predictor bits remains constant). This estimate may be slightly generous in that additional wordline repeaters/drivers may need to be inserted to avoid a substantial increase in the overall predictor access time. Any additional drivers would add to the area overhead; however, the drivers should still take up less area than a full decode tree. The checkpoint overhead[1] (includes bitcells, decoders, sense amps) also contributes a small but non-negligible amount of area. Depending on the overall hardware budget for the PBNP, the modulo-path history can reduce area requirements by 18-37% without any substantial impact on performance.

Note that in our earlier discussions and comparisons, we kept the overall hardware budget the same between the original and optimized PBNP configurations. However, the hardware budget as measured by bits of storage is only a proxy for the actual die area required for the predictor. In a practical setting, our 8KB optimized predictors would require substantially less area to implement (up to 37%). Instead of reducing the die footprint of the predictor, we could instead reclaim the area to add more perceptron entries to the tables of weights. This would help to relieve capacity conflicts in the predictor structures, but could still maintain the same hardware cost in die area as the original PBNP even though it contains a larger total number of bits of state.

5. Conclusions

Despite the high accuracy of the neural-based branch predictors, none have yet been implemented in any commercial processors. We believe that the primary obstacles to the adoption of neural predictors is in the complexity of the previously proposed schemes. The modulo-path history predictors proposed in this paper provide substantial reductions in the hardware complexity as measured by the number of tables, the predictor pipeline depth, and the checkpointing overhead, while simultaneously reducing predictor energy consumption and die-area requirements. The reduction in power and area can potentially be used to reduce the cost of the processor, or they could also be traded to im-

[1] We assume that at most 1 out of 4 instructions will be a branch, and therefore for a ROB of 64 entries, we use 16 checkpoints.

plement larger predictors than was previously possible with a conventional neural predictor organization. We believe that research aimed at providing practical implementations of sophisticated predictors is critical to successfully transferring this technology to industrial implementations.

Acknowledgments

Daniel Jiménez is supported by grants from NSF (CCR-0311091 and CCF-0545898). Gabriel Loh is supported by funding and equipment from Intel Corporation.

References

[1] The 1st JILP Championship Branch Prediction Competition (CBP-1). `http://www.jilp.org/cbp`.

[2] T. Austin, E. Larson, and D. Ernst. SimpleScalar: An Infrastructure for Computer System Modeling. *IEEE Micro Magazine*, pages 59–67, February 2002.

[3] M. Co, D. A. B. Weikle, and K. Skadron. A Break-Even Formulation for Evaluating Branch Predictor Energy Efficiency. In *Proceedings of the Workshop on Complexity-Effective Design*, Madison, WI, USA, June 2005.

[4] V. Desmet, H. Vandierendonck, and K. D. Bosschere. A 2bcgskew Predictor Fused by a Redundant History Skewed Perceptron Predictor. In *Proceedings of the 1st Championship Branch Prediction Competition*, pages 1–4, Portland, OR, USA, December 2004.

[5] H. Gao and H. Zhou. Adaptive Information Processing: An Effective Way to Improve Perceptron Predictors. In *Proceedings of the 1st Championship Branch Prediction Competition*, pages 1–4, Portland, OR, USA, December 2004.

[6] K. Ghose and M. B. Kamble. Reducing Power in Superscalar Processor Caches Using Subbanking, Multiple Line Buffers and Bit-Line Segmentation. In *Proceedings of the International Symposium on Low Power Electronics and Design*, pages 70–75, San Diego, CA, USA, August 1999.

[7] S. Gochman, R. Ronen, I. Anati, A. Berkovitz, T. Kurts, A. Naveh, A. Saeed, Z. Sperber, and R. C. Valentine. The Intel Pentium M Processor: Microarchitecture and Performance. *Intel Technology Journal*, 7(2), May 2003.

[8] M. R. Guthaus, J. S. Ringenberg, D. Ernst, T. M. Austin, T. Mudge, and R. B. Brown. MiBench: A Free, Commercially Representative Embedded Benchmark Suite. In *Proceedings of the 4th Workshop on Workload Characterization*, pages 83–94, Austin, TX, USA, December 2001.

[9] D. A. Jiménez. Fast Path-Based Neural Branch Prediction. In *Proceedings of the 36th International Symposium on Microarchitecture*, pages 243–252, San Diego, CA, USA, December 2003.

[10] D. A. Jiménez. Piecewise Linear Branch Prediction. In *Proceedings of the 32nd International Symposium on Computer Architecture*, pages 382–393, Madison, WI, USA, June 2005.

[11] D. A. Jiménez, S. W. Keckler, and C. Lin. The Impact of Delay on the Design of Branch Predictors. In *Proceedings of the 33rd International Symposium on Microarchitecture*, pages 4–13, Monterey, CA, USA, December 2000.

[12] D. A. Jiménez and C. Lin. Neural Methods for Dynamic Branch Prediction. *ACM Transactions on Computer Systems*, 20(4):369–397, November 2002.

[13] E. Larson, S. Chatterjee, and T. Austin. MASE: A Novel Infrastructure for Detailed Microarchitectural Modeling. In *Proceedings of the 2001 International Symposium on Performance Analysis of Systems and Software*, pages 1–9, Tucson, AZ, USA, November 2001.

[14] C. Lee, M. Potkonjak, and W. H. Mangione-Smith. MediaBench: A Tool for Evaluating and Synthesizing Multimedia and Communication Systems. In *Proceedings of the 30th International Symposium on Microarchitecture*, pages 330–335, Research Triangle Park, NC, USA, December 1997.

[15] G. H. Loh. The Frankenpredictor. In *Proceedings of the 1st Championship Branch Prediction Competition*, pages 1–4, Portland, OR, USA, December 2004.

[16] S. McFarling. Combining Branch Predictors. TN 36, Compaq Computer Corporation Western Research Laboratory, June 1993.

[17] P. Michaud. A PPM-like, Tag-Based Predictor. *Journal of Instruction Level Parallelism*, 7:1–10, 2005.

[18] J. M. Mulder, N. T. Quach, and M. J. Flynn. An Area Model for On-Chip Memories and its Application. *IEEE Journal of Solid-State Circuits*, 26(2):98–106, February 1991.

[19] E. Perelman, G. Hamerly, and B. Calder. Picking Statistically Valid and Early Simulation Points. In *Proceedings of the 2003 International Conference on Parallel Architectures and Compilation Techniques*, pages 244–255, New Orleans, LA, USA, September 2004.

[20] A. Seznec. Analysis of the O-GEometric History Length Branch Predictor. In *Proceedings of the 32nd International Symposium on Computer Architecture*, Madison, WI, USA, June 2005.

[21] A. Seznec, S. Felix, V. Krishnan, and Y. Sazeides. Design Tradeoffs for the Alpha EV8 Conditional Branch Predictor. In *Proceedings of the 29th International Symposium on Computer Architecture*, Anchorage, AK, USA, May 2002.

[22] A. Seznec and A. Fraboulet. Effective Ahead Pipelining of Instruction Block Address Generation. In *Proceedings of the 30th International Symposium on Computer Architecture*, San Diego, CA, USA, May 2003.

[23] P. Shivakumar and N. P. Jouppi. CACTI 3.0: An Integrated Timing, Power, and Area Model. TR 2001/2, Compaq Computer Corporation Western Research Laboratory, August 2001.

[24] J. Stark, M. Evers, and Y. N. Patt. Variable Length Path Branch Prediction. *ACM SIGPLAN Notices*, 33(11):170–179, 1998.

[25] D. Tarjan and K. Skadron. Merging Path and Gshare Indexing in Perceptron Branch Prediction. *ACM Transactions on Architecture and Code Optimization*, 2(3):280–300, September 2005.

[26] T.-Y. Yeh and Y. N. Patt. Two-Level Adaptive Branch Prediction. In *Proceedings of the 24th International Symposium on Microarchitecture*, pages 51–61, Albuqueque, NM, USA, November 1991.