An Alternative TAGE-like Conditional Branch Predictor

PIERRE MICHAUD, Inria, University of Rennes, CNRS, IRISA, France

TAGE is one of the most accurate conditional branch predictors known today. However, TAGE does not exploit its input information perfectly, as it is possible to obtain significant prediction accuracy improvements by complementing TAGE with a statistical corrector using the same input information. This article proposes an alternative TAGE-like predictor making statistical correction practically superfluous.

CCS Concepts: • **Computer systems organization** → **Superscalar architectures**;

Additional Key Words and Phrases: Branch prediction, PPM, TAGE, Bayesian confidence estimation

ACM Reference format:

Pierre Michaud. 2018. An Alternative TAGE-like Conditional Branch Predictor. *ACM Trans. Archit. Code Optim.* 15, 3, Article 30 (August 2018), 23 pages.
<https://doi.org/10.1145/3226098>

1 INTRODUCTION

The branch predictor is the keystone of modern superscalar microarchitectures. Reducing the number of branch mispredictions is a relatively simple way to increase performance and simultaneously decrease energy consumption.

Research in branch prediction has been a decades-long effort, largely focused on predicting the direction of conditional branches, as this is generally the greatest source of mispredictions. Since 2006, research in branch prediction has progressed only slowly. A plausible reason is that modern branch predictors might be close to a prediction accuracy limit. However, this is not certain either, as the only mathematically proven limit is zero misprediction, which is still far from being attained.

Today, a highly accurate conditional branch predictor based on the published state of the art is either TAGE-like [69], perceptron-based [29, 35], or a combination of those [30, 66]. Both types of predictors have distinct advantages: TAGE exploits the limited predictor storage very efficiently, whereas perceptron-based predictors can easily combine different sorts of input information.

This article focuses on TAGE-like predictors. Except for the first branch prediction championship, which was held in 2004 before TAGE was introduced [2], all four subsequent championships were won by predictors based on TAGE [3–6]. It is not known with certainty which commercial processors implement a TAGE-like predictor, as this information is rarely disclosed. Admittedly, a PPM-like predictor is implemented in the IBM zEC12 [1], and a TAGE predictor is implemented in the Phytium Mars [82] and in the IBM POWER9 [26].

Although TAGE is a very accurate predictor, it does not exploit its input information perfectly, as significant prediction accuracy improvements are obtained by complementing TAGE with a

Authors' address: P. Michaud, Inria, University of Rennes, CNRS, IRISA, Campus de Beaulieu, Rennes F-35000, France; email: pierre.michaud@inria.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

1544-3566/2018/08-ART30 \$15.00

<https://doi.org/10.1145/3226098>

statistical corrector using the same input information [63]. The statistical corrector, even small, makes the whole predictor more complex. More importantly perhaps, this imperfection of TAGE revealed by the statistical corrector is a hint that our understanding of TAGE-like predictors is incomplete.

This article proposes an alternative TAGE-like predictor, called BATAGE, making statistical correction practically superfluous. BATAGE has the same global structure as TAGE but uses a different tagged-entry format and different prediction and update algorithms.

Section 2 provides the historical and technical context for situating the contributions of BATAGE. Section 3 analyzes the *cold-counter* problem, which is the main cause for TAGE needing statistical correction. Section 4 proposes to replace the conventional up/down counter in TAGE with two counters counting separately the taken and not-taken occurrences. A new method is proposed for estimating prediction confidence using Bayesian probabilities. Section 5 describes the BATAGE predictor and its features, including

- a new prediction automaton called *dual-counter*,
- prediction and update algorithms exploiting the dual-counter and Bayesian confidence estimation,
- a new method called Controlled Allocation Throttling (CAT) for reducing the allocation frequency.

Finally, Section 6 provides an experimental evaluation of BATAGE.

2 A SHORT HISTORY OF CONDITIONAL BRANCH PREDICTION

BATAGE is derived from TAGE, which is the result of decades of research in branch prediction. This section provides a short history of conditional branch prediction. Its purpose is to furnish the historical and technical context for situating the contributions of BATAGE.

Figure 1 illustrates the improvements in prediction accuracy over the years. It shows the average MPKI (mispredictions per 1,000 instructions) over the CBP 2016 traces for various branch predictors since 1993. Appendix A provides a detailed legend of Figure 1.

2.1 The 1980s

The *bimodal* predictor, that is, a table of up/down saturating counters accessed with a hash of the branch address, was introduced by James E. Smith in the early 1980s [70].¹

In 1984, Johnny Lee and Alan Jay Smith proposed a branch prediction method consisting of branch history bit vectors stored in the Branch Target Buffer (BTB) entries and used for accessing a prediction table trained offline with representative workloads [40]. They were focusing on understanding the prediction accuracy limits of this method. They did not describe a hardware implementation. Although Lee and Smith mentioned in passing the possibility of having a global branch history (using today's terminology), their study focused on per-branch history bit vectors.

2.2 Two-level Prediction

In a seminal article published in 1991, Yeh and Patt proposed a branch predictor derived from Lee and Smith's method but using online instead of offline training [78]. They introduced the *Pattern History Table* (PHT), a prediction table holding 2-bit up/down counters and accessed with the history bit vector of the branch being predicted. They called this predictor *Two-Level Adaptive*.

¹The 2-bit automaton used in the branch predictor of the S-1 computer built at the Lawrence Livermore Laboratory during the 1970s was not an up/down counter [40]. The S-1 automaton is generally less accurate than the 2-bit up/down counter [49, 51, 78].

Fig. 1. Average number of mispredictions per 1,000 instructions (MPKI) for various conditional branch predictors on the CBP 2016 traces for 8KB, 32KB, and 64KB storage budgets (see Appendix A).

In the 1991 article, Yeh and Patt considered only per-branch (aka local) history bit vectors. The first articles describing a global-history predictor were published in 1992 by Yeh and Patt [79] and by Pan, So, and Rameh [50]. In a global-history predictor, there is a global history register shared by all the branches. Yeh and Patt generalized the label “Two Level,” including global-history predictors as a special case [79, 80]. This first wave of two-level schemes culminated with the *gshare* global-history predictor introduced by McFarling [42]. For several years, *gshare* was the reference conditional branch predictor for its simplicity and high prediction accuracy.²

2.3 Meta-Prediction and De-Aliasing Techniques

In 1993, McFarling introduced *meta-prediction*, a method for combining several different prediction methods, leading to *hybrid* predictors more accurate than *gshare* [42]. As Figure 1 shows, the *bimodal/gshare* hybrid predictor of 1993, and the improved version of 1994 [44], have an MPKI significantly lower than *gshare*. More research on hybrid predictors was published in the following years [9, 10, 18, 19, 41].

During the same period, some articles studied the *aliasing* problem stemming from the limited branch predictor size [53, 74, 81]. There was a wave of studies trying to overcome the aliasing problem [7, 17, 39, 47, 72]. This line of research culminated with the branch predictor of the Alpha EV8 processor [67], derived from Seznec’s 2bc-gskew predictor [68].

The MPKIs of some emblematic “de-aliased” predictors are shown in Figure 1 (1997–1999). These predictors are indeed more accurate than *gshare*. But compared against a well-tuned bimodal/*gshare* hybrid, the accuracy of de-aliased predictors is not much better, if better at all. The discrepancy between Figure 1 and the progresses claimed in the original articles is perhaps due, in part, to the CBP 2016 traces having different characteristics from the benchmarks used at the time. Still, the use of McFarling’s meta-predictor as a de-aliasing technique was certainly overlooked back then. Nevertheless, research on de-aliased predictors improved our understanding of branch prediction. Eventually, the lessons from de-aliased predictors led to significant accuracy gains with the branch predictor of the Alpha EV8 processor (2bc-gskew-EV8 in Figure 1).

2.4 The Perceptron

In 1999–2000, two research teams, independently, started exploring the use of artificial neural networks for branch prediction [35, 77]. In particular, Jiménez and Lin’s perceptron predictor had an important impact, not so much for the demonstrated prediction accuracy as for the completely different, thought-provoking branch prediction algorithm [35, 37]. Remarkably, the perceptron was the first branch predictor that could take advantage of a global history longer than 60 branches. It started the trend toward very long global histories.³

In their 2001 article, Jiménez and Lin emphasized a limitation of the original perceptron: *linear separability*. In the following years, several studies tackled linear separability [31, 32, 56, 58]. In particular, Seznec’s RHSP was the first perceptron-based predictor to demonstrate a prediction accuracy substantially better than that of de-aliased predictors [56, 58]. The hardware complexity of perceptron-based predictors was greatly reduced in 2004 when, inspired from Seznec’s MAC-RHSP [58], Tarjan and Skadron with the Hashed Perceptron [75, 76] and Seznec with the GEHL predictor [57, 59], decoupled the perceptron width from the global history length.

²On the CBP 2016 traces, the *average* MPKI of an 8KB *gshare* is about 35% lower than that of a large bimodal predictor.

³A recent study explains in detail how to implement a long global history [52].

Fig. 2. The PPM-like predictor of CBP 2004 [2, 45].

2.5 TAGE

In 1996, Chen et al. considered adapting Prediction by Partial Matching (PPM), a data compression technique [13], to branch prediction [12]. More specifically, they applied PPM to local-history two-level prediction, assuming a PHT with as many entries as the number of distinct local-history values of all possible lengths up to a maximum length. Extra bits were stored in the PHT to detect cold PHT entries [11]. Another research team confirmed the potential of PPM for branch prediction, but without considering hardware implementation constraints [20]. PPM was also proposed for indirect jump prediction [38].

Conventional BTBs feature tags [25, 40, 71], and in 1997–1998, when researchers started exploring more sophisticated indirect jump predictors, the use of tags was an obvious option [8, 14, 15]. However, the use of tags in conditional branch predictors was counterintuitive, as this meant much fewer 2-bit up/down counters at equal storage. The fact that tags are a cost-effective option for conditional branch prediction was discovered by McFarling with the serial-BLG predictor [43] and by Eden and Mudge with the YAGS predictor [17].

In 2001, I introduced *gtags*, an approximation of PPM consisting of several cascaded YAGS using increasing global history lengths⁴ [46], somewhat similar to Driesen and Hölzle’s multi-cascaded indirect jump predictor [16], but for conditional branches. I observed that prediction accuracy increases with the number of distinct global history lengths. However, at the time, it seemed difficult to implement a cost-effective *gtags* with more than three tables. In particular, I noticed that freshly allocated entries sometimes generated many mispredictions, a problem that I call today the *cold-counter* problem (see Section 3). In 2004, I submitted to the first branch prediction championship a predictor called *PPM-like* [45], which was a 5-table *gtags* but with two key improvements (Figure 2). The first improvement was the introduction of a *u* bit in each tagged entry for estimating the usefulness of that entry, to better manage the limited predictor storage. The second improvement was the introduction of an *m* bit in tagless entries for controlling the initialization of the up/down prediction counter on allocation of a tagged entry, to mitigate the cold counter problem. The

⁴In 2001, I quickly discarded the use of arithmetic progressions for the global history lengths and used a geometric progression instead (8,16,32,64) [46]. The fact that a geometric progression is often close to optimal was discovered and emphasized by Seznec a few years later [59, 69].

PPM-like predictor was the first predictor not derived from the perceptron with a prediction accuracy substantially better than that of de-aliased predictors (Figure 1, 8KB budget).

In 2005, Seznec improved the PPM-like predictor quite substantially, which led to TAGE [69]. Unlike the PPM-like predictor, which uses a global history of branch directions, TAGE uses a path history [48, 81]. The key features distinguishing TAGE from PPM-like are as follows:

- TAGE considers not only the longest matching path (*pred*) but also the second-longest matching one (*altpred*).
- A tiny meta-predictor selects the final prediction between *pred* and *altpred* depending on the state of the *pred* up/down counter. This is for tackling the cold counter problem. On tagged entry allocation, the up/down counter is always initialized according to the branch outcome (TAGE has no *m* bit).
- Each tagged entry contains a *u* counter (generalization of the *u* bit).
- The *pred* entry is considered useful when it prevents the occurrence of a misprediction.
- The *u* counters are periodically decremented *en masse* so that dead entries are not locked forever.
- TAGE allocates a single entry per misprediction.

All four branch prediction championships between 2006 and 2016 were won by André Seznec with a TAGE based predictor. Figure 1 shows two TAGE versions, the original 2006 version, and the latest 2016 version (see Appendix A). Seznec did a few changes over the years (in particular, he improved the management of *u* counters in 2011 [63]), sometimes depending on the storage budget considered. Nevertheless, the fundamentals of TAGE have not changed. The better accuracy of TAGE 2016 over TAGE 2006 is largely due to TAGE 2016 having more banks (and path lengths) and using bank interleaving, where a tagged bank is shared by several path lengths [62, 66].

2.6 Statistical Correction

Gao and Zhou noticed that, in a PPM-like predictor, the longest match is not always the most accurate [22], corroborating and generalizing what Seznec observed for TAGE [69]. The PMPM predictor they submitted to CBP 2006 combined some GEHL features (geometric path lengths, perceptron algorithm, adaptive threshold) with some TAGE features (tags, *u* counters) [22]. However, they did not succeed in outperforming a well-tuned TAGE [3, 61].

In 2006, Seznec noticed that, for huge storage budgets, and using the same input information, GEHL outperforms TAGE [60]. Later, he found that this was due to TAGE behaving suboptimally on hard-to-predict branches, that is, weakly-biased branches with little or no global correlation. He proposed to complement TAGE with a statistical corrector, and found that a perceptron-inspired predictor such as GEHL is a cost-effective statistical corrector for TAGE [63]. Seznec won the last three championships with this kind of predictor, called TAGE-SC [62, 65, 66].

Prediction accuracy has improved a lot since gshare: As shown in Figure 1, the 64KB TAGE-SC-L of CBP 2016 has an *average* MPKI 43% lower than a 64KB gshare.⁵

3 THE COLD COUNTER PROBLEM

The cold counter problem is specific to PPM/TAGE-like predictors. Yet, although some features were introduced in the PPM-like predictor and in TAGE to reduce it, the cold-counter problem was neither named nor emphasized. I briefly mentioned *some* aspects of the problem in 2001 [46],

⁵It should be recalled that the CBP rules do not limit predictor complexity besides a fixed storage budget. The 64KB TAGE-SC-L of CBP 2016 features a TAGE with 30 interleaved tagged banks and a statistical corrector with 20 weight tables and 3 local-history tables [66].

Fig. 3. Illustration of the cold counter problem: misprediction probability for a wide-enough up/down counter on a branch having probability q to be taken, as a function of the number of past occurrences.

and it was only implicit in subsequent articles [45, 69]. The article that introduced statistical correction seems to downplay the importance of the meta-predictor in TAGE (called USE_ALT_ON_NA) but does not explain that it is the statistical corrector that renders the meta-predictor practically superfluous [63]. The importance of the meta-predictor, which is real without the statistical corrector, is stated explicitly only in the source code of TAGE-SC-L in 2016 [6]. It is safe to say that the cold-counter problem has long been underestimated and did not receive the attention it deserves. This section provides a mathematical analysis of the cold-counter problem.

Initially, an up/down counter has no information about the branch (or path) it is trying to predict. After one occurrence of the branch, the up/down counter provides a prediction that is generally better than a random guess. In fact, if the branch is biased very strongly toward one direction, an up/down counter trained with a single occurrence of the branch is close to an ideal predictor for that branch. However, if the branch is weakly biased, it may take many occurrences of the branch for the counter to guess the branch bias with high confidence (assuming the counter is wide enough to record that amount of branch history). The number of occurrences necessary to train a wide-enough counter can be surprisingly large, as can be seen by modeling a single branch as a Bernoulli process, with probability q for the branch to be taken: The probability $M(n)$ to mispredict the branch with a (wide-enough) up/down counter previously trained with n occurrences of the branch is

$$\text{for } n \text{ odd, } M(n) = M(n+1) = \sum_{k=0}^{\frac{n-1}{2}} \binom{n}{k} q^{k+1} (1-q)^{n-k} + \sum_{k=\frac{n+1}{2}}^n \binom{n}{k} q^k (1-q)^{n-k+1},$$

where it is assumed that the counter is initialized in state 0 when the first occurrence is taken, in state -1 when it is not taken.⁶

Figure 3 shows $M(n)$ as a function of n , for $q = 0.1$ and for $q = 0.3$. The misprediction probability decreases as we accumulate knowledge about the branch and, in the limit, approaches the optimal misprediction probability $\min(q, 1-q)$. However, convergence is slow. For prediction accuracy to be within 5% of the optimal accuracy, it takes 8 occurrences of the branch with $q = 0.1$, and 20 occurrences with $q = 0.3$.

⁶The fact that $M(n)$ decreases only on odd values of n can be understood intuitively as follows. The prediction from the counter is a majority vote over the past n occurrences. When n is odd, the counter cannot change its prediction on the $n+1$ th occurrence unless the number of taken and not-take occurrences becomes equal, in which case the prediction is no better than a random guess.

Fig. 4. Tagged entry in TAGE (left) vs. BATAGE (right). The payload of a TAGE entry consists of an up/down counter (typically, 3 bits) and a u counter (1 or 2 bits). The payload of a BATAGE entry consists of two counters n_1 and n_0 (typically, 3 bits each) counting respectively the number of taken and not-taken occurrences.

PPM/TAGE-like predictors suffer from the cold counter problem, because they allocate tagged entries on mispredictions, resetting the up/down counter of newly allocated entries. Worse, if there is no global correlation, the up/down counter tends to be initialized in the wrong direction. In the PPM-like predictor, the cold counter problem was somewhat mitigated by trying to initialize the up/down counter intelligently. TAGE provided a more effective solution to the cold counter problem with the meta-predictor selecting between the longest and second-longest matching paths. Nevertheless, TAGE does not solve the cold counter problem completely. This is why Seznec obtained significant accuracy gains by combining TAGE with a GEHL statistical corrector, as perceptron-based predictors such as GEHL barely suffer from the cold-counter problem.

4 BAYESIAN CONFIDENCE ESTIMATION

While complementing TAGE with a statistical corrector is a possible solution to the cold-counter problem, it increases the branch predictor complexity and the branch prediction delay. In this article, I propose to solve the cold-counter problem without a statistical corrector.

There are two fundamental reasons why the cold-counter problem exists in TAGE:

- (1) TAGE allocates tagged entries on nearly every branch misprediction.
- (2) TAGE cannot estimate precisely the “temperature” of an up/down counter.

This section addresses the second point (the first point is addressed in Section 5.6).

On a tagged entry allocation, the counter is initialized in state 0 if the branch is taken, in state -1 if the branch is not taken. Subsequently, every time this entry is used to make a prediction, the counter is incremented or decremented depending on the branch direction. If we find the counter in state -4 , then we are sure that the counter has been updated at least 3 times after the initialization, so we know that the counter is relatively warm. However, if we find the counter is state -1 , we do not know if the counter has just been initialized or has been updated multiple times. That is, we cannot distinguish between a cold counter and a weakly biased branch.

Therefore, I propose to replace the up/down counter with two counters n_1 and n_0 counting respectively the number of taken and not-taken occurrences. This way, the “temperature” of a tagged entry can be estimated precisely. Incidentally, the introduction of n_1 and n_0 also permits removing the u counter (explanation in Section 5.3). This change of tagged entry format, shown in Figure 4, is the foundation of BATAGE. Upon a tagged entry allocation, n_1 and n_0 are both reset, and either n_1 or n_0 is incremented according to the branch direction. On subsequent updates, n_1 or n_0 increase depending on the branch direction (ignoring for the moment the limited counter width). The prediction is given by the sign of $n_0 - n_1$.

I propose to tackle the cold counter problem by associating prediction confidence levels with values of n_1 and n_0 . This way, when a tagged entry gives a low-confidence prediction, we can choose to ignore it and go to the next hitting tagged entry.

4.1 Estimating Branch Predictability

In 1981, Smith observed that, when an up/down counter is saturated, this is an indication that the branch prediction provided by that counter is likely to be correct [70]. Smith’s method can be

generalized as follows. Let $x \in [-2^c, 2^c - 1]$ be the value of the up/down counter. A fine-grained confidence level can be obtained as $|2x + 1|$: the higher $|2x + 1|$, the more confident the prediction [64]. A coarse-grained confidence level can be obtained by comparing $|2x + 1|$ with a threshold [36].

Defining confidence levels from two counters n_1 and n_0 is not as straightforward. Obviously, there exists many different possibilities. Defining the confidence level as $|n_1 - n_0|$ is definitely *not* what we want, as this would be practically equivalent to having an up/down counter. Instead, we would like to exploit the information that n_1 and n_0 provide and that is not available from an up/down counter.

For example, let us consider two predictors giving different predictions:

- **predictor A:** $n_1 = 2, n_0 = 0$ (predict *taken*)
- **predictor B:** $n_1 = 3, n_0 = 6$ (predict *not-taken*)

Given that sole information, which predictor is the most reliable? If we were using $|n_1 - n_0|$ as confidence level, then we would select prediction B. However, there is nothing obvious here. Maybe we should trust predictor A instead, even though it is colder. In fact, the question is not well defined. To answer it, we must make some assumptions.

I propose to assume that a branch has a constant but unknown probability q to be taken for a given path, with q uniformly distributed in $[0, 1]$. This assumption is somewhat arbitrary. Nevertheless it does not introduce any bias toward *taken* or *not-taken* and it is consistent with the fact that we know nothing about branches but what we learn through n_1 and n_0 . With this assumption, and using a Bayesian reasoning, it is possible to define confidence levels. Let us start from a uniform prior distribution $P_0(q) = 1$, and let us apply Bayes' theorem to obtain a posterior distribution $P(q|n_1, n_0)$:

$$P(q|n_1, n_0) = \frac{P(n_1, n_0|q) \times P_0(q)}{\int_0^1 P(n_1, n_0|t) P_0(t) dt} = \frac{P(n_1, n_0|q)}{\int_0^1 P(n_1, n_0|t) dt}.$$

As $P(n_1, n_0|q) = \binom{n_1+n_0}{n_1} q^{n_1} (1-q)^{n_0}$, we obtain

$$P(q|n_1, n_0) = \frac{q^{n_1} (1-q)^{n_0}}{B(n_1 + 1, n_0 + 1)},$$

where $B(x, y) = \int_0^1 t^{x-1} (1-t)^{y-1} dt$ is the beta function. With the posterior distribution, we can obtain an estimate of the misprediction probability, which we denote $\hat{m}$. If $n_1 < n_0$, then the prediction is *not-taken*, and the misprediction probability is q , hence

$$\hat{m} = \int_0^1 q P(q|n_1, n_0) dq = \frac{B(n_1 + 2, n_0 + 1)}{B(n_1 + 1, n_0 + 1)}$$

If $n_1 > n_0$, then the prediction is *taken*, the misprediction probability is $1 - q$, and we obtain a formula for $\hat{m}$ similar to the one above but with n_1 and n_0 permuted. Using the properties of the beta function, more precisely $B(x + 1, y) = \frac{x}{x+y} B(x, y)$ and $B(x, y) = B(y, x)$, we obtain the following simple formula:

$$\hat{m} = \frac{1 + \min(n_1, n_0)}{2 + n_1 + n_0}. \quad (1)$$

Formula (1) is known as Laplace's rule of succession. Notice that $\hat{m} = 1/2$ when $n_0 = n_1$. Table 1 gives the value of $\hat{m}$ for n_1 and n_0 in $[0, 7]$.

Estimator $\hat{m}$ provides an estimation of branch predictability, with low $\hat{m}$ meaning high confidence in the prediction. Going back to the previous example, prediction A gives $\hat{m}(2, 0) = 0.25$, prediction B gives $\hat{m}(3, 6) \approx 0.36$. Hence prediction A is deemed more reliable than prediction B.

Table 1. Estimator $\hat{m}$ as a Function of n_1 and n_0 (Formula (1))

		n_1							
		0	1	2	3	4	5	6	7
n_0	0	0.50	0.33	0.25	0.20	0.17	0.14	0.12	0.11
	1	0.33	0.50	0.40	0.33	0.29	0.25	0.22	0.20
	2	0.25	0.40	0.50	0.43	0.38	0.33	0.30	0.27
	3	0.20	0.33	0.43	0.50	0.44	0.40	0.36	0.33
	4	0.17	0.29	0.38	0.44	0.50	0.45	0.42	0.38
	5	0.14	0.25	0.33	0.40	0.45	0.50	0.46	0.43
	6	0.12	0.22	0.30	0.36	0.42	0.46	0.50	0.47
	7	0.11	0.20	0.27	0.33	0.38	0.43	0.47	0.50

It must be emphasized that $\hat{m}$ is an *estimate* of the misprediction probability, not the actual misprediction probability, which is unknown to us. Moreover, there is no reason for q to be uniformly distributed in $[0, 1]$. The actual distribution of q is an empirical property depending on the application and on the path length. In particular, $\hat{m}$ ignores global correlations, which tend to push q away from $1/2$. Nevertheless, I observed in my experiments the practical efficacy of estimator $\hat{m}$, which is the basis of the BATAGE (Bayesian TAGE) predictor described in the following sections.

4.2 Defining Confidence Levels

BATAGE does not use the full spectrum of $\hat{m}$ values, but coarsens it into a few confidence levels: high, medium, and low. In this article, medium confidence is synonymous with unknown predictability. Notice that a single occurrence of a branch gives no information about its predictability. Therefore, the case $n_1 + n_0 = 1$ is a situation where predictability is unknown. This corresponds to $\hat{m} = 1/3$, which is the natural threshold to use here. This leads to the following confidence levels:

$$\begin{aligned}
 \hat{m} < 1/3 &\longrightarrow \text{high} \\
 \hat{m} = 1/3 &\longrightarrow \text{medium} \\
 \hat{m} > 1/3 &\longrightarrow \text{low}
 \end{aligned}$$

Table 1 uses a color code for the confidence levels: red for high confidence, orange for medium confidence, and black for low confidence.

A hardware implementation does not need to compute $\hat{m}$ explicitly. It is sufficient to compute a value $\bar{m} \in \{0, 1, 2\}$ directly from n_1 and n_0 :

$$\begin{aligned}
 \bar{m} = 0 &\longrightarrow \text{high} \\
 \bar{m} = 1 &\longrightarrow \text{medium} \\
 \bar{m} = 2 &\longrightarrow \text{low}
 \end{aligned}$$

The 2-bit value $\bar{m}$ is obtained with simple combinational logic as follows. Define

$$\begin{aligned}
 \text{medium} &= (n_1 = (2n_0 + 1)) \vee (n_0 = (2n_1 + 1)) \\
 \text{low} &= (n_1 < (2n_0 + 1)) \wedge (n_0 < (2n_1 + 1)),
 \end{aligned}$$

then $\bar{m}$ is the concatenation of bits *low* and *medium*:

$$\bar{m} = 2 \times \text{low} + \text{medium}.$$

5 BATAGE

Like TAGE, BATAGE does not name one particular predictor configuration but a predictor family based on some common principles. This section describes those principles, focusing on what

Table 2. MPKI of a 32k-entry Bimodal Predictor on the CBP 2016 Traces

counter width	2 bits	3 bits	4 bits	5 bits
predictor size	8KB	12KB	16KB	20KB
average MPKI	10.47	10.07	10.12	10.29

distinguishes BATAGE from TAGE. The parts of the predictor that are not described are essentially similar to TAGE.

5.1 The Dual-Counter

So far, we have assumed that n_1 and n_0 can increase indefinitely. In practice however, the n_1 and n_0 counters in a BATAGE entry are narrow (typically, 3 bits each), for the same reasons that TAGE has narrow prediction counters: Under limited storage, the statistical benefit of using wide counters is not high enough to justify using storage for this, and using wide counters actually hurts prediction accuracy for certain branches not behaving like a Bernoulli process. For instance, Table 2 gives the average MPKI of a large bimodal predictor for different widths of the up/down counter. On the CBP 2016 traces, a large bimodal predictor is optimal with 3-bit up/down counters. 5-bit counters are significantly worse than 3-bit counters because of non-Bernoulli branch behaviors.

So we must deal with the fact that n_1 and n_0 cannot exceed a limit $nmax$ (with 3-bit counters, $nmax = 7$). BATAGE uses the following update method (C language):

```

if (taken) { // branch is taken
    if (n1 < nmax) n1++;
    else if (n0 > 0) n0--;
} else { // branch is not taken
    if (n0 < nmax) n0++;
    else if (n1 > 0) n1--;
}

```

In other words, the counter corresponding to the branch direction is incremented, unless that counter equals $nmax$, in which case the other counter is decremented if not null. I call (n_1, n_0) and the algorithm above the *dual-counter* automaton.

It should be noted that, after having hit the limit $nmax$ once, the dual-counter behaves like an up/down counter with $2 \times nmax + 1$ states. Although formula (1) was established assuming unlimited n_0 and n_1 , BATAGE always uses the confidence levels shown in Table 1.

The dual-counter is used only in tagged entries, not in the tagless table, where this would be a waste of storage. The tagless table of BATAGE contains conventional up/down counters, like TAGE.

5.2 BATAGE Prediction Algorithm

Let $L_1, L_2, \dots, L_G$ denote the path lengths from shortest to longest, where G is the number of predefined path lengths. Similar to TAGE, a BATAGE prediction starts by reading simultaneously the tagless entry and the G tagged entries.

BATAGE computes the 2-bit confidence level $\overline{m}_i$ for all tagged entries $i \in [1, G]$, as described in Section 4.2. This computation is done in parallel with tag comparisons. When there is a tag miss at L_i , we set $\overline{m}_i = 3$.

Fig. 5. Divide-and-conquer implementation of the prediction selection circuit for TAGE (left) and BATAGE (right). The circuit for TAGE takes as input the hit/miss bit, the prediction bit, and the low-confidence bit (needed for the final selection between $pred$ and $altpred$). The circuit for BATAGE takes as input the $\bar{m}$ value (2 bits) and the prediction bit. This example assumes 7 tagged banks. The path lengths increase from left to right. The leftmost input is the bimodal entry.

We also define a confidence level $\bar{m}_0 \in \{0, 1, 2\}$ for the tagless entry, as follows. The tagless entry of BATAGE contains a 3-bit up/down counter $x \in [-4, 3]$. The value $\bar{m}_0$ is set according to the following equivalence:

x	-4	-3	-2	-1	0	1	2	3
n_1	1	1	1	1	2	3	4	5
n_0	5	4	3	2	1	1	1	1
$\bar{m}_0$	0	0	1	2	2	1	0	0

The BATAGE prediction algorithm is (C-like pseudo-code):

```

j = G
for (i=G-1; i>=0; i--) {
  if ( $\bar{m}_i < \bar{m}_j$ ) j=i;
}
// entry j provides the final prediction

```

In other words, the final prediction is provided by the entry with the smallest $\bar{m}$ value, with priority to the longest path in case of equality. The prediction selection circuit of BATAGE is no more complex than that of TAGE. A divide-and-conquer approach is possible, allowing a tree structure, as illustrated in Figure 5. Other implementations are possible. In some implementations, it might be advantageous to use a 3-bit thermometer code for $\bar{m}$ ($0 \rightarrow 000$, $1 \rightarrow 001$, $2 \rightarrow 011$, $3 \rightarrow 111$), which allows to compute the minimum of multiple $\bar{m}$ values with a bitwise AND.

5.3 Decay Mechanism

In TAGE, the u counter is for protecting useful entries against eviction. There are no u counters in BATAGE. Instead, we use the confidence level as an indication of utility: A high confidence prediction is more likely to be useful than a medium or low confidence one. More precisely, the following rule is used in BATAGE:

a tagged entry giving a high confidence prediction cannot be evicted.

Table 3. Updating the Hitting Entries

entry	action	condition	comment
$L_h > L_p$	update with branch direction	always	entries were skipped probably because they are cold
L_p	update with branch direction	$L_p = L_0$ <u>or</u> L_p not high confidence <u>or</u> L_n not high confidence <u>or</u> incorrect prediction from L_n	
	decay	$L_p > L_0$ <u>and</u> L_p high confidence <u>and</u> L_n high confidence <u>and</u> correct prediction from L_n	entry is probably useless
L_n	update with branch direction	$L_p > L_0$ <u>and</u> L_p not high confidence	improves prediction accuracy only slightly

L_0 is the tagless entry. L_p is the path length that provided the BATAGE prediction. Hitting entries are denoted L_h (L_0 is one of them). When $L_p > L_0$, L_n denotes the longest L_h such that $L_h < L_p$.

Yet, a high confidence entry that is no longer used should not be locked forever. For such entry to become evictable, there must be a way to *decay* it. We define the action of decaying an entry as follows (C language):

```

if (n1 > n0) n1--;
if (n0 > n1) n0--;

```

In other words, we decay an entry by decrementing its highest count, n_1 or n_0 . Decay leaves the numerator in formula (1) unchanged but decreases the denominator. Hence decay increases $\hat{m}$. In particular, decaying a high confidence dual-counter multiple times eventually leads to a medium confidence state (see Table 1). In practice, it is unnecessary to decay a dual-counter further than medium confidence.

5.4 Updating the Hitting Entries

Predictor update is done at the in-order retirement pipeline stage [63]. We define $L_1, L_2, \dots, L_G$ as in Section 5.2. Moreover, let L_0 denote the tagless entry. At update, BATAGE needs the following information: which path length L_p provided the BATAGE prediction (Section 5.2), and all the path lengths L_h corresponding to a tag hit (including L_0). Table 3 describes how the hitting entries are updated. When $L_p > L_0$, L_n denotes the longest L_h such that $L_h < L_p$. That is, L_n is the next hitting entry after the entry that provided the BATAGE prediction. Entries $L_h < L_n$ are not updated. Entries $L_h > L_p$ are always updated with the branch direction, as described in Section 5.1. Those entries exist because shorter path lengths could not deliver 100% accuracy, but they were skipped probably because they are cold. By being updated, they can warm up and become useful. The dual-counter in entry L_p is decayed when that entry is deemed useless, that is, when $L_p > L_0$ and the prediction from L_n is correct and L_p and L_n are both high confidence.

5.5 Allocating New Entries

TAGE-like predictors allocate tagged entries on branch mispredictions, because it is hoped that longer paths will provide more accurate predictions. Allocating a tagged entry means setting the tag and initializing the dual-counter according to the branch outcome ($n_0 + n_1 = 1$). Like TAGE,

BATAGE allocates entries only on mispredictions, at retirement. BATAGE allocates at most one entry per misprediction, like the original 2006 TAGE.

Let L_m be the longest hitting path length ($L_m = L_0$ when there are no hits) and $s \geq 1$ a small random number. The allocation algorithm selects the entry corresponding to the shortest path length $L_r \geq L_{m+s}$ such that the entry at L_r is not high confidence ($\bar{m} \neq 0$). If such L_r does not exist, then no allocation is done, and we set $r = G + 1$. Then, independent of allocation's success, the entries $L_i \in [L_{m+s}, L_{r-1}]$ (which are all high confidence) are decayed with probability PDEC, as described in Section 5.3. The best value of PDEC is determined empirically. PDEC=1/4 works well on the CBP 2016 traces.

5.6 Controlled Allocation Throttling (CAT)

While Bayesian confidence estimation helps mitigate the cold-counter problem, it does not solve the root cause, which is the allocation of tagged entries on every branch misprediction.

I observed that, on traces with hard-to-predict branches, prediction accuracy improves when the allocation probability is set to a value less than 1, which I call *allocation throttling*. This observation is valid both for TAGE and BATAGE. The problem is that allocation throttling degrades prediction accuracy severely on some traces. In fact, an allocation probability of one, as in TAGE, is the best fixed allocation probability on average.

The problem is to find the best allocation probability automatically. That is, we are looking for an effective *controlled allocation throttling* (CAT) method. After trying several methods that did not work, I made progress toward a solution when I conjectured that one could tell whether there are too many or too few allocations by looking at the global predictor state, more precisely the distribution of dual-counter states. I noticed that, when entries are allocated too aggressively, many entries are in medium or low confidence state, either because allocated entries are not reused or because these entries are useless and kept in medium/low confidence state by the update rules (Table 3). Conversely, medium/low confidence entries are very few when the allocated entries are quickly reused and converted into high confidence entries. However, my attempts to adjust allocation throttling to obtain a fixed number of medium/low confidence entries did not produce good results. Then, I realized that what matters is not the absolute number of medium/low confidence entries but their proportion with respect to the number of *moderately high* confidence entries. The rest of this section provides a detailed description of the CAT method I found.

The main idea is to compare the amount of entries in a medium/low confidence state, denoted **NHC** (for *not high confidence*), and the amount of entries in a *moderately high confidence* state, denoted **MHC**, where MHC states are the states such that $0.17 < \hat{m} < 1/3$. That is, the MHC states are the high confidence states from which are excluded $n_0 = 0, n_1 \geq 4$ and $n_1 = 0, n_0 \geq 4$ (see Table 1). MHC states are generally transient states. Besides statistical fluctuations due to random branch behavior, allocation and decay are two major causes for the existence of MHC entries.

When NHC entries significantly outnumber MHC entries, this is generally because allocation is easy (little decay) but allocated entries stay in a medium or low confidence state, hence allocation is probably too aggressive. Conversely, when MHC entries outnumber NHC entries, this is an indication that allocation is difficult (much decay) and allocated entries quickly evolve toward a high confidence state, so allocation can be more aggressive.

A possible implementation of CAT would be to monitor the fractions F_n and F_m of NHC and MHC entries respectively. In practice though, we just want to know the ratio F_m/F_n . This ratio can be estimated as follows. On an allocation attempt, the average number of skipped entries is $1/F_n - 1$, neglecting the fact that the number G of tagged banks is finite. The fraction of high confidence entries that are MHC is $F_m/(1 - F_n)$. Hence the average number of skipped MHC entries


```

if (misprediction) {
  r = random in [0, MINAP-1];
  if (r >= ((cat*MINAP)/(CATMAX+1))) {
    mhc = 0;
    s = random in [1, SKIPMAX];
    for (i=m+s; i ≤ G; i++) {
      if ( $\hat{m}[i] < 1/3$ ) { // high confidence, protected
        decay with probability PDEC;
        if ( $\hat{m}[i] > 0.17$ ) mhc++;
      } else { // not high confidence, take the entry
        overwrite tag;
        initialize dual-counter;
        cat = cat + 3 - 4 * mhc; // CATR=3/4
        cat = min(CATMAX, max(0, cat));
        break; // stop here
      }
    }
  }
}

```

Fig. 6. BATAGE allocation algorithm (C-like pseudo-code). L_m is the longest hitting path length, $\hat{m}[i]$ is the $\hat{m}$ estimator for the entry at L_i . The *cat* counter is initially set to zero.

is

$$\left(\frac{1}{F_n} - 1\right) \times \frac{F_m}{1 - F_n} = \frac{F_m}{F_n},$$

which is the quantity we are looking for.

When the average number of skipped MHC entries is below a certain threshold CATR, the allocation probability is decreased if possible. The allocation probability cannot fall below a value $1/\text{MINAP}$. When the average number of skipped MHC entries is greater than CATR, the allocation probability is increased if possible. Typically, the optimal CATR value lies between $1/4$ and $3/4$ and it tends to increase with G .

Figure 6 describes the BATAGE allocation algorithm. The *cat* counter controls the allocation probability. The allocation probability is 1 when *cat* is small, it is $1/\text{MINAP}$ when *cat* is close to CATMAX. The maximum *cat* value, CATMAX, is typically much larger than the total number of tagged entries. The predictor can warm up more quickly if *cat* is reset to zero when an application starts running.

6 EXPERIMENTAL EVALUATION

I evaluated BATAGE with the CBP 2016 simulation infrastructure, which provides a set of 220 traces for tuning predictors and a distinct set of 440 traces for evaluating them [6].

The average MPKIs of TAGE and BATAGE are shown in Figure 1 (2016 and 2017) for 8KB, 32KB and 64KB storage budgets. To obtain a stand-alone 8KB TAGE, I started from the 7KB TAGE embedded in the 8KB TAGE-SC-L that won CBP 2016 [66], and I tuned it without the statistical corrector. In particular, I replaced 3-bit counters in tagged entries with 4-bit counters (this improves the average prediction accuracy) and I scaled the predictor up to 8KB by increasing the number of banks and the number of path lengths. For the 32KB and 64KB TAGE, I proceeded similarly but starting from the 56KB TAGE embedded in the winning 64KB TAGE-SC-L.

The prediction accuracy gain of BATAGE over TAGE is modest on average, but tangible: It is roughly equivalent to the accuracy gain that the 2001 perceptron brought over an e-gskew. The 8KB BATAGE even matches the MPKI of the CBP 2016 winner, despite being simpler (no statistical corrector, no local history, no loop predictor).

Fig. 7. MPKI difference to 7KB TAGE for BATAGE, TAGE-GSC, and BATAGE-GSC (higher is better). The traces are sorted for each curve separately so that the curve is non-decreasing.

Table 4. Average MPKI of TAGE and BATAGE with and without Statistical Correction

predictor	TAGE	TAGE-GSC	BATAGE	BATAGE-GSC
size	7KB	8KB	7KB	8KB
MPKI	4.42	4.27	4.29	4.24

The rest of this section provides a detailed evaluation of BATAGE, focusing on a 7KB storage budget. The 7KB TAGE component embedded in the 8KB TAGE-SC-L of CBP 2016 features two u bits per tagged entry. After replacing the 3-bit up/down counter in the tagged entry of TAGE with a 4-bit counter, the tagged entries of TAGE and BATAGE have the same payload size (6 bits). This way, the 7KB TAGE and BATAGE configurations can be compared with equal tag size and equal number of entries and banks.

6.1 BATAGE Does Not Need Statistical Correction

Figure 7 and Table 4 compare TAGE and BATAGE with and without a global-history statistical corrector (GSC). The GSC is the 1KB statistical corrector implemented in the 8KB TAGE-SC-L of CBP 2016, but with the local-history components disabled so that the effect is pure statistical correction. Figure 7 shows the MPKI difference to the 7KB TAGE (higher is better) for the 440 traces. Table 4 provides the average MPKI (arithmetic mean over the 440 traces).

BATAGE outperforms TAGE significantly on several tens of traces (rightmost part of the red curve). The statistical corrector helps TAGE significantly, as expected [63]. The average MPKI of the 8KB TAGE-GSC is very close to that of the 7KB BATAGE. However, the 8KB BATAGE-GSC is only marginally better than the 7KB BATAGE. In practice, BATAGE does not need statistical correction.

6.2 Controlled Allocation Throttling Is an Essential Feature of BATAGE

Two features of BATAGE contribute to solve the cold-counter problem:

Table 5. BATAGE: Impact of Disabling CAT and Restricting the Prediction Algorithm

CAT	enabled			disabled		
prediction selection	$k = G + 1$	$k = 2$	$k = 1$	$k = G + 1$	$k = 2$	$k = 1$
MPKI	4.29	4.38	4.62	4.43	4.58	4.77

Fig. 8. Impact of disabling controlled allocation throttling on a 7KB BATAGE.

- Bayesian confidence estimation takes into account the “temperature” of tagged entries,
- Controlled allocation throttling decreases the number of allocations hence the number of cold entries.

The impact of the BATAGE prediction algorithm can be evaluated by restricting the selection of the final prediction among the first k hitting entries. The case $k = G + 1$ is the BATAGE prediction algorithm. The case $k = 1$ corresponds to selecting the prediction from the longest hitting path, as in the PPM-like predictor [45]. The case $k = 2$ corresponds to selecting between the longest and second longest hitting paths, as in TAGE (but without a meta-predictor).

The impact of CAT can be evaluated by disabling it, that is, by setting the allocation probability equal to one. When CAT is disabled and k is small, the use of decay in the update policy (Table 3) hurts prediction accuracy, because the safety net provided by the BATAGE prediction algorithm has been removed. In these cases (no CAT, small k), the update policy must be modified so that the entry providing the final prediction is always updated.⁷

Table 5 gives the average MPKI for the 7KB BATAGE and for configurations with CAT disabled and/or a small k . Both disabling CAT and limiting k hurt prediction accuracy. Figure 8 shows, for each trace, the MPKI degradation from disabling CAT. Several tens of traces benefit substantially from CAT (leftmost part of the curve), and a few traces experience a degradation, yet barely noticeable (rightmost part of the curve).

CAT is an essential feature of BATAGE: Without it, BATAGE would be no better than TAGE (compare tables 4 and 5). The management of u counters in TAGE is very effective at preserving

⁷However, the use of decay is important for CAT, as it is one of the mechanisms allowing to detect useless allocations by keeping useless entries in a medium confidence state.

Fig. 9. APKI difference to TAGE (higher means fewer allocations).

useful entries. BATAGE has no u counter and instead relies on dual-counter decay. However, without CAT, dual-counters are decayed too aggressively on certain traces. CAT is effective not only at reducing the number of allocations but also at adjusting decay intensity.

Figure 9 compares the number of allocations per 1000 instructions (APKI) of TAGE and BATAGE, showing also the impact of disabling CAT. Without CAT, BATAGE would do more allocations than TAGE on average. This is because TAGE does not allocate entries when the overall prediction is incorrect but the prediction from the longest matching path is correct.⁸ However, with CAT enabled, the number of allocations is reduced dramatically on many traces (Figure 9, right part of the red curve).

6.3 CAT on TAGE?

Given the benefit brought by CAT on BATAGE, it would be satisfying to find a CAT method for TAGE.⁹ The dashed curves in Figure 10 show the prediction accuracy change when the allocation probability in TAGE is set to a fixed value, here $1/2$ and $1/4$. Allocation throttling with a fixed allocation probability improves prediction accuracy on certain traces but degrades it on other traces. In TAGE as in BATAGE, allocation throttling must be controlled. However, the CAT method described in Section 5.6 relies on specificities of BATAGE. It is not clear how TAGE could imitate BATAGE. This is a question for future research.

Nevertheless, it is possible to evaluate the potential prediction accuracy improvements that a hypothetical CAT method would bring to TAGE. Figure 11 depicts the experiment I did. I simulated a BATAGE in lockstep with TAGE, the *cat* counter of BATAGE being communicated to TAGE continuously so that TAGE can set its allocation probability like BATAGE at every instant. The result of this experiment is shown in Figure 10 as the solid red curve. The rightmost part of the curve shows that there is a potential for significant improvement on several tens of traces, motivating future research on CAT for TAGE.

⁸This feature of TAGE decreases the number of allocations and brings modest MPKI reductions.

⁹In a private discussion I had with André Seznec in July 2017, he told me that he was aware of the potential benefit of allocation throttling but had not found a way to control it.

Fig. 10. Allocation throttling on TAGE. The dashed curves are for a fixed allocation probability. The curve labeled “magic CAT” is for TAGE with a CAT counter provided by BATAGE.

Fig. 11. Doing CAT “magically” on TAGE.

7 CONCLUSION

The dual-counter is central to BATAGE. It permits solving the cold-counter problem by allowing Bayesian confidence estimation and by providing a way to control allocation throttling. My original motivation for getting rid of the u counter of TAGE was that the dual-counter needs more bits than a conventional up/down counter, and keeping the u counter would have made the tagged entry larger. Although dual-counter decay allows to recover some of the benefit of the u counter, the two mechanisms work differently, and perhaps they provide distinct advantages. Maybe future advances on TAGE-like prediction will come from a better understanding of the differences between TAGE and BATAGE. Nevertheless, dual-counter decay seems to be essential for CAT to work on BATAGE. Finding a CAT method for TAGE is a topic for future research.

APPENDIX

A LEGEND FOR FIGURE 1

Figure 1 shows the average MPKI (arithmetic mean over the CBP 2016 traces) of various conditional branch predictors over the years, at 8KB, 32KB and 64KB storage budgets. The predictors are *gshare* [42], the 1993 *bimodal/gshare* hybrid [42], a *bimodal/gshare* where the meta-prediction table is accessed by hashing together the branch address and the global branch history [10, 44], *e-gskew* [47], *bimode* [39], YAGS [17], the original *2bc-gskew* [68], a *2bc-gskew* tuned along the lines of the Alpha EV8 predictor [55, 67], the original global-history *perceptron* [35], PBNP, [31], MAC-RHSP [58], PWL [32], O-GEHL [59], the PPM-like predictor [45], TAGE [69], L-TAGE [61], PMPM [22], FTL [27], ISL-TAGE [62, 63], FTL++ [28], OH-SNAP [33, 73], TAGE-SC-L [65, 66], SSHP [34], BFN [23, 24], MPP [29], and BATAGE (this article).

All the predictors were evaluated on the 440 CBP 2016 traces [6]. Predictors whose name is followed by the mention “CBP” in parentheses participated in a branch prediction championship. For these predictors, I used the original source code written by the authors, downloaded from the

championship's web site. I did not modify the source code except for the minimal modifications necessary to simulate the predictor under the CBP 2016 simulation infrastructure.¹⁰ I could not simulate the winning predictor of CBP 2004 [21], as it uses information not available in the CBP 2016 traces. Also, some CBP predictors are omitted for clarity.

For the 2bc-gskew-EV8, MAC-RHSP and TAGE 2006 predictors, I used the source code downloaded from André Seznec's website [54]. I obtained the TAGE 2016 predictor by disabling the statistical corrector in Seznec's CBP 2016 source code and by tuning TAGE parameters to fit the storage budget considered. I implemented the other predictors from the descriptions provided in the original articles.

I tuned the non-CBP predictors using the 223 training traces of CBP 2016. This tuning was done by varying the explicit parameters (for instance, the path length). I did not modify the original algorithms except for the bimode and YAGS predictors, which I made slightly more accurate with algorithm modifications simple enough to be considered part of the tuning.

All the non-CBP predictors and five CBP predictors (PPM-like, O-GEHL, L-TAGE, ISL-TAGE, BFN) use only a global branch or path history as first history level. The predictors using the perceptron algorithm are shown in red, those not using it in blue (the statistical corrector in ISL-TAGE and TAGE-SC-L is GEHL-like).

REFERENCES

- [1] J. Bonnano, A. Collura, D. Lipetz, U. Mayer, B. Prasky, and A. Saporito. 2013. Two level bulk preload branch prediction. In *Proceedings of the International Symposium on High Performance Computer Architecture (HPCA'13)*.
- [2] CBP. 2004. Championship Branch Prediction (CBP-1). Retrieved from <https://www.jilp.org/cbp/>.
- [3] CBP. 2006. Championship Branch Prediction (CBP-2). Retrieved from <http://hpca23.cse.tamu.edu/taco/camino/cbp2/>.
- [4] CBP. 2011. Championship Branch Prediction (CBP-3). Retrieved from <https://www.jilp.org/jwac-2/>.
- [5] CBP. 2014. Championship Branch Prediction (CBP-4). Retrieved from <https://www.jilp.org/cbp2014/>.
- [6] CBP. 2016. Championship Branch Prediction (CBP-5). Retrieved from <https://www.jilp.org/cbp2016/>.
- [7] P.-Y. Chang, M. Evers, and Y. N. Patt. 1996. Improving branch prediction accuracy by reducing pattern history table interference. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT'96)*.
- [8] P.-Y. Chang, E. Hao, and Y. N. Patt. 1997. Target prediction for indirect jumps. In *International Symposium on Computer Architecture (ISCA'97)*.
- [9] P.-Y. Chang, E. Hao, T.-Y. Yeh, and Y. N. Patt. 1994. Branch classification: A new mechanism for improving branch predictor performance. In *Proceedings of the International Symposium on Microarchitecture (MICRO'94)*.
- [10] P.-Y. Chang, E. Hao, T.-Y. Yeh, and Y. N. Patt. 1995. Alternative implementations of hybrid branch predictors. In *Proceedings of the International Symposium on Microarchitecture (MICRO'95)*.
- [11] I.-C. Chen. 1997. *Enhancing the Instruction Fetching Mechanism Using Data Compression*. Ph.D. Dissertation. University of Michigan.
- [12] I.-C. K. Chen, J. T. Coffey, and T. N. Mudge. 1996. Analysis of branch prediction via data compression. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'96)*.
- [13] J. G. Cleary and I. H. Witten. 1984. Data compression using adaptive coding and partial string matching. *IEEE Trans. Commun.* 32, 4 (April 1984).
- [14] K. Driesen and U. Hözlze. 1998. Accurate indirect branch prediction. In *Proceedings of the International Symposium on Computer Architecture (ISCA'98)*.
- [15] K. Driesen and U. Hözlze. 1998. The cascaded predictor: Economical and adaptive branch target prediction. In *Proceedings of the International Symposium on Microarchitecture (MICRO'98)*.
- [16] K. Driesen and U. Hözlze. 1999. Multi-stage cascaded prediction. In *Proceedings of the International Euro-Par Conference on Parallel Processing*. LNCS 1685.
- [17] A. N. Eden and T. Mudge. 1998. The YAGS branch prediction scheme. In *Proceedings of the International Symposium on Microarchitecture (MICRO'98)*.

¹⁰The CBP 2016 traces are based on the ARM architecture, while previous championships were based on the Intel x86 architecture. Hence, for the old CBP predictors, I right-shifted every branch address by 2 bits.

- [18] M. Evers. 2000. *Improving Branch Prediction by Understanding Branch Behavior*. Ph.D. Dissertation. University of Michigan.
- [19] M. Evers, P.-Y. Chang, and Y. N. Patt. 1996. Using hybrid branch predictors to improve branch prediction accuracy in the presence of context switches. In *Proceedings of the International Symposium on Computer Architecture (ISCA'96)*.
- [20] E. Federovsky, M. Feder, and S. Weiss. 1998. Branch prediction based on universal data compression algorithms. In *Proceedings of the International Symposium on Computer Architecture (ISCA'98)*.
- [21] H. Gao and H. Zhou. 2005. Adaptive information processing: An effective way to improve perceptron predictors. *J. Instruct.-Level Parallel*. 7 (April 2005).
- [22] H. Gao and H. Zhou. 2007. PMPM: Prediction by combining multiple partial matches. *J. Instruct.-Level Parallel*. 9 (May 2007).
- [23] D. Gope and M. H. Lipasti. 2014. Bias-free branch predictor. In *Proceedings of the International Symposium on Microarchitecture (MICRO'14)*.
- [24] D. Gope and M. H. Lipasti. 2014. Bias-free neural predictor. In *Championship Branch Prediction (CBP-4)*. Retrieved from <https://www.jilp.org/cbp2014/>.
- [25] R. N. Ibbett. 1972. The MU5 instruction pipeline. *Comput. J.* 15, 1 (Feb. 1972).
- [26] IBM. 2018. POWER9 Processor User's Manual, version 2.0. Retrieved from <https://openpowerfoundation.org>.
- [27] Y. Ishii. 2007. Fused two-level branch prediction with ahead calculation. *J. Instruct.-Level Parall.* 9 (May 2007).
- [28] Y. Ishii, K. Kuroyanagi, T. Sawada, M. Inaba, and K. Hiraki. 2011. Revisiting local history for improving fused two-level branch prediction. In *Championship Branch Prediction (CBP-3)*. Retrieved from <https://www.jilp.org/jwac-2/>.
- [29] D. Jiménez. 2016. Multiperspective perceptron predictor. In *Championship Branch Prediction (CBP-5)*. Retrieved from <https://www.jilp.org/cbp2016/>.
- [30] D. Jiménez. 2016. Multiperspective perceptron predictor with TAGE. In *Championship Branch Prediction (CBP-5)*. Retrieved from <https://www.jilp.org/cbp2016/>.
- [31] D. A. Jiménez. 2003. Fast path-based neural branch prediction. In *Proceedings of the International Symposium on Microarchitecture (MICRO'03)*.
- [32] D. A. Jiménez. 2005. Piecewise linear branch prediction. In *Proceedings of the International Symposium on Computer Architecture (ISCA'05)*.
- [33] D. A. Jiménez. 2011. OH-SNAP: Optimized hybrid scaled neural analog predictor. In *Championship Branch Prediction (CBP-3)*. Retrieved from <https://www.jilp.org/jwac-2/>.
- [34] D. A. Jiménez. 2014. Strided sampling hashed perceptron predictor. In *Championship Branch Prediction (CBP-4)*. Retrieved from <https://www.jilp.org/cbp2014/>.
- [35] D. A. Jiménez and C. Lin. 2001. Dynamic branch prediction with perceptrons. In *Proceedings of the International Symposium on High Performance Computer Architecture (HPCA'01)*.
- [36] D. A. Jiménez and C. Lin. 2002. *Composite Confidence Estimators for Enhanced Speculation Control*. Technical Report TR-02-14. The University of Texas at Austin, Department of Computer Sciences.
- [37] D. A. Jiménez and C. Lin. 2002. Neural methods for dynamic branch prediction. *ACM Trans. Comput. Syst.* 20, 4 (Nov. 2002).
- [38] J. Kalamatianos and D. R. Kaeli. 1998. Predicting indirect branches via data compression. In *Proceedings of the International Symposium on Microarchitecture (MICRO'98)*.
- [39] C.-C. Lee, I.-C. K. Chen, and T. N. Mudge. 1997. The bi-mode branch predictor. In *Proceedings of the International Symposium on Microarchitecture (MICRO'97)*.
- [40] J. K. F. Lee and A. J. Smith. 1984. Branch prediction strategies and branch target buffer design. *Computer* 17, 1 (Jan. 1984).
- [41] G. H. Loh and D. S. Henry. 2002. Predicting conditional branches with fusion-based hybrid predictors. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT'02)*.
- [42] S. McFarling. 1993. *Combining Branch Predictors*. Technical Report TN-36. DEC WRL.
- [43] S. McFarling. 1998. Branch predictor with serially connected predictor stages for improving branch prediction accuracy. (March 1998). Patent No. 6374349.
- [44] S. McFarling, S. C. Steely Jr., J. Emer, and E. McLellan. 1994. Trainable apparatus for predicting instruction outcomes in pipelined processors. (May 1994). Patent No. 5758142.
- [45] P. Michaud. 2005. A PPM-like, tag-based predictor. *J. Instruct.-Level Parall.* 7 (April 2005).
- [46] P. Michaud and A. Seznec. 2001. *A Comprehensive Study of Dynamic Global History Branch Prediction*. Technical Report RR-4219. INRIA.
- [47] P. Michaud, A. Seznec, and R. Uhlig. 1997. Trading conflict and capacity aliasing in conditional branch predictors. In *Proceedings of the International Symposium on Computer Architecture (ISCA'97)*.
- [48] R. Nair. 1995. Dynamic path-based branch correlation. In *International Symposium on Microarchitecture (MICRO'95)*.
- [49] R. Nair. 1995. Optimal 2-bit branch predictors. *IEEE Trans. Comput.* 44, 5 (May 1995).

- [50] S.-T. Pan, K. So, and J. T. Rameh. 1992. Improving the accuracy of dynamic branch prediction using branch correlation. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'92)*.
- [51] C. G. Ponder and M. C. Shebanow. 1990. An information-theoretic look at branch-prediction. In *Studies in Branch-prediction*, C. G. Ponder (Ed.). UCRL-ID-106077. Lawrence Livermore National Laboratory.
- [52] D. J. Schlais and M. H. Lipasti. 2016. BADGR: A practical GHR implementation for TAGE branch predictors. In *Proceedings of the International Conference on Computer Design (ICCD'16)*.
- [53] S. Sechrest, C.-C. Lee, and T. Mudge. 1996. Correlation and aliasing in dynamic branch predictors. In *Proceedings of the International Symposium on Computer Architecture (ISCA'96)*.
- [54] A. Seznec. (2017). Retrieved from <https://team.inria.fr/pacap/members/andre-seznec/>.
- [55] A. Seznec. 2003. An optimized 2bcgskew branch predictor. Retrieved from <https://team.inria.fr/pacap/members/andre-seznec/>.
- [56] A. Seznec. 2003. *Redundant History Skewed Perceptron Predictors: Pushing Limits on Global History Branch Predictors*. Technical Report PI-1554. IRISA.
- [57] A. Seznec. 2004. The O-GEHL branch predictor. In *Championship Branch Prediction (CBP-1)*.
- [58] A. Seznec. 2004. *Revisiting the Perceptron Predictor*. Technical Report PI-1620. IRISA.
- [59] A. Seznec. 2005. Analysis of the O-geometric history length branch predictor. In *Proceedings of the International Symposium on Computer Architecture (ISCA'05)*.
- [60] A. Seznec. 2007. The idealistic GTL predictor. *J. Instruct.-Level Parallel*. 9 (May 2007).
- [61] A. Seznec. 2007. The L-TAGE branch predictor. *J. Instruct.-Level Parall*. 9 (May 2007).
- [62] A. Seznec. 2011. A 64 Kbytes ISL-TAGE branch predictor. In *Championship Branch Prediction (CBP-3)*. Retrieved from <https://www.jilp.org/jwac-2/>.
- [63] A. Seznec. 2011. A new case for the TAGE branch predictor. In *International Symposium on Microarchitecture (MICRO'11)*.
- [64] A. Seznec. 2011. Storage free confidence estimation for the TAGE branch predictor. In *Proceedings of the International Symposium on High Performance Computer Architecture (HPCA'11)*.
- [65] A. Seznec. 2014. TAGE-SC-L branch predictors. In *Championship Branch Prediction (CBP-4)*. Retrieved from <https://www.jilp.org/cbp2014/>.
- [66] A. Seznec. 2016. TAGE-SC-L branch predictors again. In *Championship Branch Prediction (CBP-5)*. Retrieved from <https://www.jilp.org/cbp2016/>.
- [67] A. Seznec, S. Felix, V. Krishnan, and Y. Sazeides. 2002. Design tradeoffs for the Alpha EV8 conditional branch predictor. In *Proceedings of the International Symposium on Computer Architecture (ISCA'02)*.
- [68] A. Seznec and P. Michaud. 1999. *De-aliased Hybrid Branch Predictors*. Technical Report RR-3618. INRIA.
- [69] A. Seznec and P. Michaud. 2006. A case for (partially) tagged geometric history length branch prediction. *J. Instruct.-Level Parall*. 8 (Feb. 2006).
- [70] J. E. Smith. 1981. A study of branch prediction strategies. In *Proceedings of the International Symposium on Computer Architecture (ISCA'81)*.
- [71] M. K. Smotherman, E. H. Sussenguth, and R. J. Robelen. 2016. The IBM ACS project. *IEEE Ann. Hist. Comput*. 38, 1 (Jan. 2016).
- [72] E. Sprangle, R. S. Chappell, M. Alsup, and Y. N. Patt. 1997. The agree predictor: A mechanism for reducing negative branch history interference. In *Proceedings of the International Symposium on Computer Architecture (ISCA'97)*.
- [73] R. St. Amant, D. A. Jiménez, and D. Burger. 2008. Low-power, high-performance analog neural branch prediction. In *Proceedings of the International Symposium on Microarchitecture (MICRO'08)*.
- [74] A. R. Talcott, M. Nemirovsky, and R. C. Wood. 1995. The influence of branch prediction table interference on branch prediction scheme performance. In *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT'95)*.
- [75] D. Tarjan and K. Skadron. 2004. *Revisiting the Perceptron Predictor Again*. Technical Report CS-2004-28. University of Virginia.
- [76] D. Tarjan and K. Skadron. 2005. Merging path and gshare indexing in perceptron branch prediction. *ACM Trans. Arch. Code Optimiz.* 2, 3 (Sept. 2005).
- [77] L. N. Vintan and M. Iridon. 1999. Towards a high performance neural branch predictor. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN'99)*.
- [78] T.-Y. Yeh and Y. N. Patt. 1991. Two-level adaptive training branch prediction. In *Proceedings of the International Symposium on Microarchitecture (MICRO'91)*.
- [79] T.-Y. Yeh and Y. N. Patt. 1992. Alternative implementations of two-level adaptive branch prediction. In *Proceedings of the International Symposium on Computer Architecture (ISCA'92)*.

- [80] T.-Y. Yeh and Y. N. Patt. 1993. A comparison of dynamic branch predictors that use two levels of branch history. In *Proceedings of the International Symposium on Computer Architecture (ISCA '93)*.
- [81] C. Young, N. Gloy, and M. D. Smith. 1995. A comparative analysis of schemes for correlated branch prediction. In *Proceedings of the International Symposium on Computer Architecture (ISCA '95)*.
- [82] C. Zhang. 2015. Mars: a 64-core ARMv8 processor. Hot Chips. Retrieved from <http://www.hotchips.org/archives/>.

Received March 2018; revised April 2018; accepted May 2018