

Micro BTB: A High Performance and Storage Efficient Last-Level Branch Target Buffer for Servers

Vishal Gupta*

School of Computer and Communication Sciences, EPFL,
Lausanne, Switzerland
vishal.gupta@epfl.ch

Biswabandan Panda

Indian Institute of Technology Bombay, Mumbai, India
biswa@cse.iitb.ac.in

ABSTRACT

High-performance branch target buffers (BTBs) and the L1I cache are key to high-performance front-end. Modern branch predictors are highly accurate, but with an increase in code footprint in modern-day server workloads, BTB and L1I misses are still frequent. Recent industry trend shows usage of large BTBs (100s of KB per core) that provide performance closer to the ideal BTB along with a decoupled front-end that provides efficient fetch-directed L1I instruction prefetching. On the other hand, techniques proposed by academia, like BTB prefetching and using retire order stream for learning, fail to provide significant performance with modern-day processor cores that are deeper and wider.

We solve the problem fundamentally by increasing the storage density of the last-level BTB. We observe that not all branch instructions require a full branch target address. Instead, we can store the branch target as a branch offset, relative to the branch instruction. Using branch offset enables the BTB to store multiple branches per entry. We reduce the BTB storage in half, but we observe that it increases skewness in the BTB. We revisit the need for skewed indexing and propose a skewed indexed and compressed last-level BTB design called MicroBTB (MBTB) that stores multiple branches per BTB entry. We evaluate MBTB on 100 industry-provided server workloads. A 4K-entry MBTB provides 17.61% performance improvement compared to an 8K-entry baseline BTB design with a storage savings of 47.5KB per core.

CCS CONCEPTS

• **Computer systems organization** → **Superscalar architectures, pipeline computing.**

KEYWORDS

Superscalar cores, Branch Target Buffer, Performance

ACM Reference Format:

Vishal Gupta and Biswabandan Panda. 2022. Micro BTB: A High Performance and Storage Efficient Last-Level Branch Target Buffer for Servers. In *19th ACM International Conference on Computing Frontiers (CF'22)*, May 17–19, 2022, Torino, Italy

*The work was done while the author was an MS student at IIT Kanpur.

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.

CF'22, May 17–19, 2022, Torino, Italy

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9338-6/22/05...\$15.00

<https://doi.org/10.1145/3528416.3530224>




Figure 1: Performance improvements of state-of-the-art BTB designs scaled to 8K entry, normalized to the baseline system with 8K entry BTB (Table 1) for 100 industry provided server workloads. Area overheads (in terms of storage requirements per core) are normalized to 8K entry baseline BTB. Our proposal MBTB with 4K entries, outperforms all the competitive BTB designs, with storage savings. Refer section 2 for the details about these designs. Refer Table 1 for the baseline system parameters.

17–19, 2022, Torino, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3528416.3530224>

1 INTRODUCTION

Branch predictor and branch target buffer (BTB) are the two key front-end structures that play a significant role in providing high performance. Branch predictors predicts the direction (taken or not-taken) of a conditional branch. BTB provides information about whether an instruction is a branch instruction or not, and the target of the branch instruction. To prevent the processor from going on the wrong execution path, both the structures need to be accurate. Modern branch predictors like hashed perceptron[1] and TAGE[2] are highly accurate. However, with the increase in code footprint in server workloads, BTB and L1I cache see frequent misses. The code footprint ranges in multi-megabytes because of the deepening software stack with frequent updates/patches.

State-of-the-art techniques for reducing front-end bottlenecks like Shotgun[3] and SN4L+Dis+BTB[4] employ L1I prefetcher to prefetch cache blocks in the L1I. These techniques decode the prefetched cache blocks and pre-fill the branches into the BTB. With an increase in code footprint, L1I misses are high, and the L1I prefetchers in these techniques are not able to prefetch the instruction blocks in a timely manner. L1I prefetcher in Shotgun learns the control flow of a program using the retire order instruction stream. With modern-day high-performing processors becoming deeper and wider with each generation, the fetch stage of the pipeline is now much further ahead in the instruction stream compared to the retire stage of the pipeline. We show in Section 5 that the increase

in the number of in-flight instructions causes lateness in Shotgun, resulting in lower performance.

Recent industry trend shows that modern processors employ a decoupled front-end [5–8] with a multi-level BTB design. A decoupled front-end decouples the branch prediction unit from the instruction fetch and allows the branch predictor to generate the address of future instructions. Modern processors like Arm Neo-verse N1[5], AMD Zen2 [6], Samsung Exynos M3 [7] and IBM z15 [8] have high BTB capacities of 6K, 7K, 16K and 144K entries, respectively. With technology scaling slowing down [9], the number of transistors available because of reduction in transistor size is slowing down. BTB size cannot increase with the increase in code footprint without reducing the size of other structures, which creates a need to reduce the size of the BTB while providing similar or better performance.

Previous work[10] for reducing the storage overhead of BTB uses a Page Number (PN) cache to store the page numbers and use a smaller index to the PN-cache instead of a page number, for storing branch tag and target. The PN-cache needs to be accessed when accessing the BTB and also to get the branch target if it is a BTB hit. The PN-cache adds extra latency to the critical path of accessing the BTB.

Opportunity: Figure 1 shows average performance improvement with various state-of-the-art BTB designs [3, 4, 11, 12] normalized to a baseline system[13] with two-level BTB, a 128-entry L1 BTB and an 8192-entry L2 BTB. We do not show the performance and area overheads for Phantom BTB[14, 15] and Confluence [16] for 8K-entry BTB as it demands significant storage. Our baseline front-end is a decoupled one with a 24-entry (192 instructions) fetch target queue (FTQ) and a fetch directed instruction prefetcher (FDIP)[17]. We use a 57-bit virtual address (keeping in mind the recent trend of a five-level page table[18]) with a 57-bit instruction pointer (IP). We use 100 industry provided server workloads, which are available through the 1st championship value prediction (CVP-1) [19]. We extend ChampSim [20] with detailed extensions to the front-end and the memory hierarchy. In the baseline system (Table 1) for 100 server workloads, the average BTB and L1I MPKI are 8.6 and 54.94, respectively.

Figure 1 shows the performance improvement with the ideal BTB, ideal L1I, and the ideal front-end. For the ideal BTB, the first instance of a branch instruction is a miss, and the rest of the instances are hits. For the ideal L1I, L1I misses are converted to hits only when there is some space left in the miss status handling registers (MSHRs). MSHRs are allocated on an L1I miss and deallocated on an L1I fill. This method of calculating ideal numbers for the L1I takes into the account the bandwidth constraint between L1I and L2. The ideal front-end has both ideal BTB and ideal L1I.

A recent work [21] has shown that with a decoupled front-end, an FDIP prefetcher, and the ideal BTB, the ideal L1I does not provide a significant performance improvement. We also observe the same in Figure 1 where an ideal front-end provides 3% more improvement compared to the ideal BTB. Figure 1 also shows that state-of-the-art BTB designs like Shotgun and SN4L+Dis+BTB provide marginal performance improvement or require high storage like Skewed BTB[12]. Please note that Figure 1 shows the performance of state-of-the-art BTB designs with a decoupled front-end. To the best of

our knowledge, this is the first work that evaluates the impact of a decoupled front-end on state-of-the-art BTB designs.

Our Contributions: We observe that not all branches require a full target, and we can encode the branch target in fewer bits. FDIP-X[11] uses four different BTBs to store branches with different encoding types. We propose a single compressed L2BTB design called MicroBTB (MBTB), storing one or multiple branches per BTB entry. MBTB uses skewed indexing to reduce the conflict misses arising due to halving the storage.

In summary, we make the following key contributions:-

- We show that the BTB misses are costly compared to the L1I misses with a decoupled front-end (Section 3).
- We propose a skewed and compressed L2BTB design called MicroBTB (MBTB) to mitigate these costly BTB misses (Section 4).
- We show that a 4K-entry MBTB provides an average performance improvement of 17.61% compared to an 8K-entry baseline BTB, resulting in 51% storage savings (Section 5).

2 RELATED WORK

In this section, we provide an overview of some of the state-of-the-art BTB designs.

Phantom BTB[15]: Phantom BTB proposes a hierarchical BTB design where at the first level, it uses a conventional BTB. However, it exploits temporal correlation among control flow jumps at the second level, packs multiple first-level BTB misses as temporal groups, and stores them at the last-level cache (LLC) using virtualization. A second-level BTB with 4K temporal groups can demand up to 256KB of LLC space shared among all the cores. Unfortunately, with this design, the performance does not improve significantly as the level-1 BTB incurs frequent misses, and getting a response from the level-2 BTB is usually delayed because it is limited by the LLC access latency.

Air BTB[16]: Air BTB is a block-based BTB design in which the BTB is in sync with the L1I cache. It stores a branch bitmap to identify branch instructions inside the cache block. Prefetch or demand blocks which are filled into the L1I cache are predecoded, and the branches are stored in the Air BTB.

Shotgun BTB[3]: Shotgun BTB proposes a new BTB design segregated by the type of control flow jump. It uses three kinds of BTB: (i) C-BTB for conditional branches, (ii) U-BTB for unconditional branches, and (iii) RIB for return. Shotgun insight is that the global control flow of a program is determined by unconditional branches, whereas conditional branches determine the local control flow. Most of the BTB is dedicated to unconditional branches where it stores the spatial footprint and uses that for prefetching the next instruction blocks. Shotgun pre-decodes conditional branches from these prefetched blocks, which helps achieve a high hit rate for conditional branches despite its small size.

SN4L+Dis+BTB [4]: SN4L+Dis+BTB uses baseline BTB design but uses an additional BTB prefetch buffer to hold predecoded branches. It proposes a next-line and discontinuity-based prefetcher for L1I prefetching and performs a Shotgun style BTB prefetching by pre-decoding the prefetched blocks.

FDIP-X [11]: FDIP-X uses four BTBs with different branch target offsets, which is the distance between the branch instruction and its




Figure 2: (a) Starvation cycles per kilo instructions for baseline, ideal BTB, ideal L1I, and the ideal front-end (b) Breakdown of branch resolution latency with front-ends having two different queue sizes for the front-end structures: FTQ and Decode queue. Fetch stage deals with the stalls when the instruction is in the fetch stage of the pipeline, fetching the instruction for L1I cache. Fetch queue deals with the stalls when an instruction is in the FTQ.

target. The insight that drives FDIP-X is that branch offset lengths are not distributed equally: conditional branches have shorter offsets than unconditional branches, enabling FDIP-X to store a higher number of branches in the same storage budget.

Skewed BTB[12]: Skewed BTB design uses different set indexing function for each way to increase the utilization of BTB entries. Utilization increases because a branch instruction that can go to a single set in the baseline design can go to multiple sets with the skewed BTB design.

3 MOTIVATION

In this section, we first elaborate on why BTB misses are costly compared to the L1I misses. We then show that the BTB miss cost increases with an increase in queue size of the front-end structures like FTQ and decode queue. Finally, we discuss about how to improve the storage density of the BTB to reduce these costly BTB misses.

BTB misses are costly compared to L1I misses: Figure 1 shows that the ideal front-end provides only 3% performance improvement on top of the ideal BTB. To understand this performance trend, we use the metric *starvation cycles per kilo instructions (SCKI)*. This metric provides the number of cycles for which the decode stage is empty, and no instruction is sent to the ROB. Higher SCKI indicates that the front-end is a bottleneck as it under-utilizes the bandwidth between decode stage and ROB. Figure 2(a) shows that with an ideal BTB, the SCKI reduces to 110 from 289, but with an ideal front-end, the SCKI reduces to 71.

On top of ideal BTB, ideal L1I does not decrease the SCKI, significantly, because with a decoupled front-end, on an L1I miss, the branch predictor continues to predict future instructions. The request for cache block containing these instructions is sent to the L1I, effectively hiding the latency for future instructions as multiple L1I requests are going in parallel. Recent work[21] also shows that L1I prefetchers like EIP[22] and FNL+MMA[23] do not provide significant performance improvement with a decoupled front-end and FDIP prefetcher.

On a BTB miss, the fetch pipeline is either stalled if the branch predictor predicts the direction of the branch is taken with high confidence or the pipeline goes on the wrong path until the branch instruction executes. For a direct branch/call instruction, the branch gets resolved in the decode stage, whereas for indirect branch/call

instruction, the branch gets resolved in the execute stage. Once the branch instruction is resolved, the front-end is flushed if it was going on the wrong path, and the fetch stage continues with the correct branch target.

After resolving the branch, the front-end is re-steered to the predicted taken path. There are multiple cycles for which the decode stage is empty, and nothing is added to the ROB, while the instructions gets fetched from the L1I cache. The back-end of the processor does not have any instructions to execute, making the BTB miss costly in terms of performance. This motivates us to design a BTB that can reduce the BTB misses and help in increasing the overall performance.

Modern processors use a decoupled front-end with an increase in queue size for the front-end structures like FTQ to support high-performing FDIP prefetcher. We next discuss the impact of this increase in queue size on branch resolution latency.

Impact of increasing queue size: Figure 2(b) shows the breakdown of branch resolution latency per branch type for a front-end with a large FTQ and a small FTQ. Large FTQ has 24 entries, where each entry can hold up to eight instructions from the same cache line, storing 192 instructions [13]. Small FTQ has 18 entries, where each entry can store one instruction. This showcases the impact of different queue sizes on the branch resolution latency, which is the time it takes for a branch instruction to resolve a BTB miss. A branch instruction is first fetched from the L1I cache and then sent to the decode queue for decoding. For direct branch/call instruction, the branch gets resolved in the decode stage, whereas for indirect branch/call instruction, the branch gets resolved in the execute stage.

On a BTB miss, for a return instruction, the front-end gets to know the type of instruction in the decode stage, and the fetch stage is re-steered to access the return address stack (RAS). Since the RAS accurately captures the return instruction target, the branch gets resolved after the decode stage. If the target stored in RAS is wrong, then the branch gets resolved in the execute stage.

Figure 2(b) shows that increasing the FTQ and decode queue size increases the branch resolution latency. Note that the increase in latency may not be on the critical path since the older instructions that precede the BTB missing branch can still represent useful work. The breakdown shows that the time spent in the fetch and decode queue increases as the size of the queue increases. The decode stage




Figure 3: Frequency of occurrence of branch instructions whose target is within a specified offset. Offset refers to the distance between branch IP and its target.

Variant	56 bits Tag Array		32 bits Offset Array	
0	T1: 28 bits	Branch Target: Upper 25 bits	Branch Target: Lower 32 bits	
1	T1: 28 bits	T2: 28 bits	O1: 16 bits	O2: 16 bits

(a) Variants of L2BTB entry

IP	37 bits	A2 (10 bits)	A1 (10 bits)	IP	29 bits	A3 (28 bits)
Set 0	= A2 ⊕ A1		Set 1	= σ ¹ (A2) ⊕ A1		Tag Variant 0 = A3
Set 2	= σ ² (A2) ⊕ A1		Set 3	= σ ³ (A2) ⊕ A1		Tag Variant 1 = A3

Where, ⊕ is XOR and σⁿ is n-bit right circular shift

(b) Hash function for set indexing.

(c) Hash function for tag.

Figure 4: (a) Variants of MBTB entry. T1 and T2 store the hash of the branch instruction and O1 and O2 store the branch offset (b) Skewed function that generates four indices for four banks (c) Tag calculation for different MBTB entry variants.

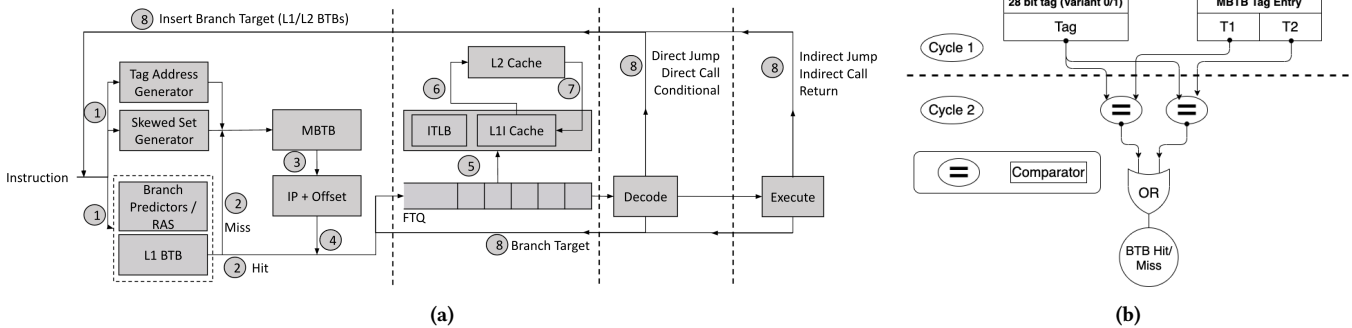


Figure 5: (a) Flow of a branch instruction with a decoupled front-end and the MBTB (b) Implementation of the MBTB tag comparison.

decodes a fixed number of instructions per cycle, but if the queue size increases, then the time spent by an instruction in the queue increases. This adds extra latency in resolving the branch, which forces the front-end to be on the wrong path for higher number of cycles. With a processor similar to our baseline system having a 24-entry FTQ (192 instructions) and a 60-instruction decode queue, there are many instructions in the front-end that need to be flushed once the mis-predicted branch is resolved, and the fetch stage moves to the correct target. For indirect branches, the penalty is more because the increase in queue size increases the number of instructions in the ROB. There is a delay in the execution of branch instruction because multiple instructions competes for the shared resource in the execute stage.

Need for a storage efficient BTB design: BTB misses are costly compared to L1I misses, and this cost is increasing with the increase in queue size with modern processors. With the slowdown

of Moore’s law, increasing the BTB size is not practical, and this motivates to design a BTB with increased storage density.

Scope for BTB Compression: We alleviate this increase in BTB cost by increasing the storage density of the last-level BTB that requires storing multiple branches per BTB entry. Figure 3 shows the branch offset bits required to encode the branch target of a branch instruction relative to the IP of the branch instruction. We find that storing a 57-bit branch target is not required for many branch instructions as observed in the previous work [11]. We store a branch IP and its offset in a BTB entry instead of the branch target. On a BTB hit, the offset is added to the branch IP to get the branch target. Since the bits required to store the offset are less than the branch target, a single BTB entry can store multiple branches.

Figure 3 also shows the branch target offset required per branch

type. Conditional and direct jump instructions have shorter offsets as these branches define the local control flow of the program, whereas call/return instructions have longer offsets as these branches define the global control flow of the program [3].

4 MICRO BTB: DESIGN AND IMPLEMENTATION

Using our observation that multiple branches can be stored using offsets, we propose a skewed and compressed L2 BTB called MicroBTB (MBTB). MBTB stores a hash of the branch instruction as tag and branch target as offset. Our baseline front-end is a decoupled one with FTQ that stores the predictions from the branch prediction unit until these instructions get consumed by the L1I. L1BTB design is the same as in the baseline system. Halving the storage increase conflict misses, and we use skewed indexing to reduce these conflict misses with the MBTB. We use a return address stack (RAS) that is a part of the baseline front-end to get the branch target of a return branch instruction. We use ITTAGE [2] for indirect branch prediction.

Compressed MBTB entry: Figure 4(a) shows two different types of variants which an MBTB entry can store. We find that increasing the number of variants increases the complexity and decreases performance gains. Section 5 shows the effect of more variants with MBTB. Variant 0 is uncompressed, and it stores a single branch instruction and its target. Variant 1 stores two branch instructions per entry depending on the size of offsets bit required to encode the distance between branch instruction and its target. Each offset field's most significant bit stores the direction of the branch *i.e.* whether it is a forward or a backward branch. Branch instruction that requires more than 15 bits for encoding the offset uses the uncompressed variant 0. Branch instructions that require less than 15 bits for encoding the offset uses the compressed variant.

Skewed set and Tag address generator: To minimize BTB conflict misses by spreading out the BTB accesses across all the sets, we generate four different MBTB set numbers using a skewing function [24, 25]. The skewing function generates different set numbers mapped to four different banks for addresses that would have otherwise conflicted using the baseline indexing function. Figure 4(b) describes the skewing function used. Figure 4(c) shows that MBTB uses lower 28 bits of the branch IP for both variant 0 and variant 1.

4.1 Design

Figure 5(a) shows the flow of a branch instruction through a decoupled front-end with MBTB.

BTB access: The address of the branch instruction goes through L1 BTB and branch predictors (①). Simultaneously, the address goes through the skewed set generator and the tag address generator (①). If there is a hit in L1 BTB (②), then the branch target from L1 BTB goes to the FTQ and we drop the current request to MBTB. If we get a miss in L1 BTB (②), then we access the MBTB. The skewed set generator provides four different set numbers for each bank. We check the variant and compare the tag present in each entry based on the variant, to check if it is a hit or not. Simultaneously, we extract the offset from MBTB entry. If there is a MBTB hit (③), then we add/subtract the offset to the branch IP based on

the most significant bit in the offset field and send it to the FTQ (④). If there is a MBTB miss, then the branch instruction goes to the subsequent stages of the processor pipeline to get the branch target.

Instruction fetch: Instructions from the FTQ send requests to L1I (⑤). Since, L1I is virtually indexed, physically tagged (VIPT), ITLB is accessed to get the physical address and then L1I is accessed. If the request hits in L1I, the instruction is sent back to FTQ, else if it misses, it goes to the lower level of cache hierarchy (⑥). Instruction cache block is fetched from the lower levels of cache hierarchy and filled in L1I (⑦) and the instruction is sent to the FTQ.

BTB insertions and updates: For direct jump/call and conditional types of branches, we get the target in the decode stage, and we insert the target into the L1 and L2 BTBs (⑧). For indirect jump/call and return types of branches, we get the target after the branch is executed in the execute stage and then we update the L1 and L2 BTBs (⑧). Note that the FTQ also gets updated with the target. BTB updates during the wrong-path execution affect performance. However, with the MBTB, BTB MPKI is extremely low (more details in Section 5) and the wrong-path BTB updates do not affect the overall performance improvement significantly.

4.2 Implementation

MBTB is a 4096-entry skewed L2BTB with four banks, each of 1024 sets that are direct-mapped, implemented as four SRAM tables. To get the effect of 4-way skewed associative mapping, we generate four BTB set numbers scattered across four banks for a given branch IP. On an MBTB insertion, if we do not find a BTB entry to allocate, then we choose a victim out of those four BTB entries through a random replacement policy. We find marginal utility with more sophisticated replacement policies as the BTB MPKI drops significantly due to compression. Each entry in MBTB has a 56-bit tag field, a 32-bit offset field, two bits to determine the branch type, and one bit to determine the variant, totaling 91 bits.

Storage savings: MBTB has 4096 entries where each entry is of 91 bits. The storage overhead for MBTB is **45.5KB**. The baseline system uses an 8192-entry L2BTB where each entry has a 32-bit tag field, a 57-bit target field, 2-bit LRU, and two bits to determine the branch type totaling 93 bits. The storage overhead for baseline L2BTB is **93KB**. MBTB provides a storage savings of **47.5KB** and an average performance improvement of **17.61%** compared to an 8K entry baseline BTB. Section 5 provides detailed performance improvement.

MBTB Tag Comparison: The tag address and the skewed set generator for the MBTB operate on the *same clock cycle* when the L1 BTB is accessed. On the first cycle, the skewed set generator generates four different set numbers mapped to four different banks from a 57-bit IP, and concurrently the tag address generator generates a tag for different variants. If there is an L1 BTB miss, then we check the MBTB.

Figure 5(b) shows the steps for MBTB tag comparison. The tag field of the MBTB is of 56 bits, consisting of two 28-bit chunks. On cycle 1, the tags T1 and T2 are read from the MBTB entry selected by the skewed set generator. On cycle 2, the 28-bit comparator compares the tag bits from the tag address generator to the bits stored in the MBTB entry. The OR gate combines the result, and

Table 1: Parameters of the baseline system.

Core	4 GHz, out of order, hashed perceptron branch predictor for conditional branches [1] that merges [26], [27], and [1] with branch prediction accuracy of 99.6%. ITTAGE [2] for indirect branches, 24-entry FTQ, 60-entry decode queue, OoO width: 6, ROB size: 352, RAS: 32 entries
BTB	L1BTB: 128 entry (64 sets, 2 way, LRU), 1 cycle L2BTB: 8192 entry (2048 sets, 4 way, LRU), 2 cycle
Caches	L1I: 32KB (8-way), LRU, 16 MSHRs, FDIP prefetcher[17], 4 cycles L1D: 48KB (12-way), LRU, 16 MSHRs, IPCP prefetcher[28], 5 cycles L2: 512KB (8-way), SRRIP[29], 32 MSHRs, IPCP prefetcher[28], 10 cycles L3: 2MB, 16 way, DRRIP[29], 64 MSHRs, 20 cycles
DRAM	4 GB, 1 channel, 6400 MT/sec

Table 2: BTB designs with their storage requirements.

Front-ends	BTB	L1I prefetcher	Total storage
Baseline	2048 sets and 4 ways, 32-bit tag, 57-bit target IP, two LRU bits, two branch type bits	FDIP prefetcher	93 KB (8K entry)
Shotgun	Basic block based BTBs: 4096 entry UBTB (4-way, 114 bits per entry), 2048 entry CBTB (4-way, 97 bits per entry), and 2048 entry RIB (4-way, 40 bits per entry)	BTB directed prefetcher	91.25 KB (8K entry)
SN4L+DIS+ BTB	Baseline BTB (93KB for 8K entry)	As described in the paper [4], 7.6KB	100.6 KB (8K entry)
Skewed BTB	2048 sets and 4 way, 32-bit tag, 57-bit target IP, two branch type bits	FDIP prefetcher	91 KB (8K entry)
FDIPX BTB	6144 entry 8-bit offset BTB (4 ways, 29 bits per entry), 6144 entry 13-bit offset BTB (4 way, 34 bits per entry), 6144 entry 23-bit offset (4 way, 44 bits per entry) and 896 entry 57-bit offset BTB (4 way, 77 bits per entry)	FDIP prefetcher	66.92 KB
MBTB	1024 sets and 4 way, 56-bit tag field, 32-bit offset field, two branch type bits and one variant bits	FDIP prefetcher	45.5 KB (4K entry) 91 KB (8K entry)

if the output is one, then it is a BTB hit, or else it is a BTB miss. All the offsets are extracted concurrently with the tag comparison process. After the tag comparison, the offset is added to the branch IP to get the branch target. The branch target is added to the FTQ. **Storing Indirect Branches in MBTB:** Direct and conditional types of branches have a single target, for which the variant is fixed, unlike indirect branches. Indirect branches can have multiple offsets depending on the distance between branch IP and the target IP. An indirect branch stored in MBTB in variant-1 format can require variant-0 if the offset between the current IP and the branch target is more than 15 bits. In this case, we invalidate the entry with variant-1, and we select a victim amongst the four entries provided by the skewed set generator using a random replacement policy. We evict the victim entry and then insert the indirect branch with variant-0.

Storing Return Branches in MBTB: Return type of branch instructions do not need to store their branch target in the BTB as RAS is used to get the target. We still need to track the return instruction in the BTB because the pipeline only knows about the type of instruction in the decode stage. At the fetch stage, BTB helps in identifying branch instruction and its type. We store the return instructions with variant-1 as it takes the least storage.

5 EVALUATION

5.1 Evaluation Methodology

We evaluate BTB designs based on the following metrics: (i) performance in terms of IPC improvement, (ii) BTB miss coverage, and

(iii) SCKI. We first evaluate the MBTB's effectiveness in eliminating BTB misses. We then discuss the significance of skewed indexing in MBTB. Finally, we show the sensitivity studies on the number of variants, MBTB size, MBTB ways, and MBTB access latency.

Workloads and infrastructure: We use 100 server workloads released publicly with CVP-1 [19] (co-located with ISCA '18) which shows the highest sensitivity to BTB and L1I. A subset of these workloads were used in 1st instruction prefetching championship [31] co-located with ISCA 2020. These workloads contain instruction traces of both user and kernel level. We use an in-house extension of ChampSim [20] simulator. We extend ChampSim with detailed *de-coupled* front-end [17], back-end, memory hierarchy, and a detailed memory system. Our extension supports BTBs, RAS, FTQ, VIPT L1 caches, and a detailed virtual memory system (five-level page table walker (PTW) and four memory management unit (MMU) caches to back the PTW). The three-level cache hierarchy also stores the page translations. Table 1 shows the parameters used, which are similar to the recent Intel Sunny Cove [13][21]. For each trace, we first warmup for 50M instructions and then report performance for the next 50M instructions.

BTB designs: Based on the performance improvement and storage overhead, we use the following recent works for our evaluation: (i) Shotgun [3] (ii) SN4L+Dis+BTB [4], (iii) Skewed BTB [12] and (iv) FDIP-X [11]. Table 2 provides the details about these techniques. We use PACTI [32] to measure BTB access latency for different BTB organizations.

5.2 BTB MPKI and Performance Improvement

Figure 6(a) shows the BTB misses per kilo instructions (MPKI) of different state-of-the-art BTB designs. The change in BTB MPKI directly correlates with the SCKI for different BTB designs, as shown in Figure 6(b). Figure 6(c) shows the performance improvement with state-of-the-art BTB designs and a 4K and 8K entry MBTB. A 4K-entry MBTB compared to an 8K-entry L2BTB baseline provides an average performance improvement of 17.61%, whereas an 8K-entry MBTB provides an average performance improvement of 16.91%. Ideal BTB provides an average performance improvement of 20.38%.

MBTB of 4K-entry decreases the BTB MPKI from 8.6 to 1.35. The decrease in BTB MPKI results in SCKI reduction from 289 to 138, which is closer to SCKI of the ideal BTB of 110. The decrease in BTB MPKI stalls the decode stage for fewer cycles, resulting in the reduction in the SCKI. MBTB of 4K-entry is able to capture, on average, 6266 branches in the L2BTB. This increase in the number of branches is what helps MBTB to reduce the BTB MPKI significantly. *An 8K-entry MBTB reduces the MPKI from 1.35 to 1.24 compared to a 4K-entry MBTB, but it still provides lower performance because of an increase in access time of the MBTB. A 4K-entry MBTB can be accessed in two cycles, whereas an 8K-entry MBTB takes three cycles to access.* An 8K-entry MBTB will be more effective with higher code footprint workloads or when workloads are running in parallel on a single core *i.e.* simultaneous multi-threading (SMT).

Shotgun BTB of 8K entry increases MPKI from 8.6 to 9.29 compared to a baseline system. Shotgun has a large BTB for unconditional branches and a small BTB for conditional branches. It uses BTB prefetching to prefetch conditional branches and uses the retire order stream to learn what to prefetch. Retire order stream




Figure 6: State-of-the-art BTB designs and MBTB evaluated on the following metrics :- (a) L2BTB Misses per Kilo Instructions (Lower the better) (b) Stall Cycles per Kilo Instructions (Lower the better) (c) Performance Improvement (Higher the better). 8K-entry MBTB provides lower performance compared to 4K-entry MBTB due to increase in access latency. 4K-entry MBTB outperforms all other competitive BTB organizations.




Figure 7: (a) Number of in-flight instruction with a system similar to Intel Sunny Cove[13] and Broadwell[30] micro-architecture (b) Performance improvement of MBTB with only skewed indexing, only compression and both skewed indexing and compression compared to 8K-entry L2BTB baseline.




Figure 8: Variant sensitivity of a 4K entry MBTB.

provides better learning as it is free from the wrong execution path taken by the processor due to inaccurate branch prediction or a BTB miss. However, with modern-day high-performing processors becoming deeper and wider with each generation, these schemes fail to improve performance. For example, Intel’s Broadwell[30] has a 192-entry Reorder Buffer (ROB) with an instruction throughput of four instructions per cycle, which increases to 352-entry ROB and six instructions per cycle with Intel’s Sunny Cove processor[33]. The increase in ROB size increases the number of in-flight instructions. Figure 7(a) shows the number of in-flight instructions in the pipeline for a particular phase of a server trace for a system similar to Broadwell and Sunny Cove processors. The number of in-flight instructions is taken every ten cycles during the simulation phase. The in-flight instruction includes the instruction in the fetch, decode and execute stage of the pipeline.

In the baseline system with 24 entry (192 instruction) FTQ, 60-entry decode queue, and 352-entry ROB, there can be a maximum

of 604 instructions in the pipeline if all the queues are full. Figure 7(a) shows that on average, in the baseline system, there are 200 instructions in-flight for this trace, with a maximum of 557. Shotgun[3] uses a configuration similar to the Broadwell processor in which the average number of in-flight instruction is 45, with a maximum reaching around 100. This 4x increase in the number of instructions in the pipeline leads to an increase in BTB MPKI with Shotgun. As the number of instructions increases, the fetch stage is much further in the program execution order than the retire stage, and since, Shotgun learns based on the retire order stream, it gets delayed. Shotgun is not able to issue timely prefetch requests for BTB prefetching, which results in higher BTB MPKI.

SN4L+Dis+BTB decreases the BTB MPKI slightly from 8.6 to 8.4. This is because it uses BTB pre-decoding to decode branches from the instruction cache blocks brought by the SN4L+Dis prefetcher. Due to large code footprint server workloads, the baseline system has an L1I MPKI of 54. SN4L+Dis prefetcher reduces the L1I MPKI to 46. Since the L1I prefetcher cannot reduce the L1I MPKI significantly, BTB pre-decoding dependent on the L1I prefetcher cannot decode the branches in a timely manner.

5.3 Importance of Skewed Indexing

Figure 7(b) shows the performance improvement of a 4K and 8K entry MBTB with skewed indexing, compression, and both skewed indexing and compression. For a 4K-entry MBTB with just skewed

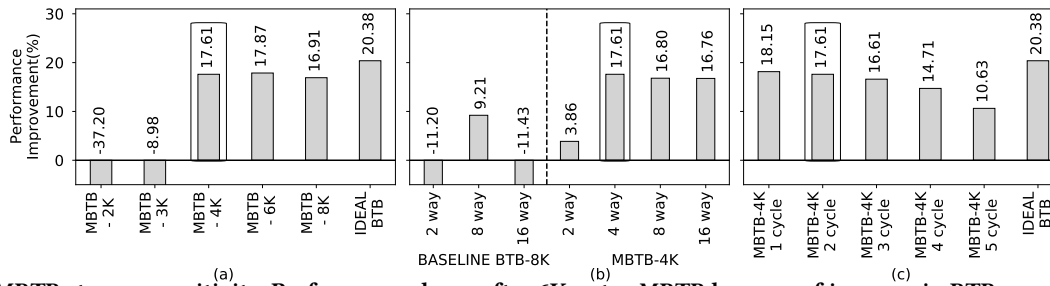


Figure 9: (a) MBTB storage sensitivity. Performance drops after 6K-entry MBTB because of increase in BTB access time (b) Way sensitivity with 8K-entry Baseline BTB and 4K-entry MBTB (c) Access latency sensitivity with 4K-entry MBTB.

indexing, the performance decreases by 24.88%, whereas with compression, the performance decreases by 3.42%. When skewed indexing and compression are used, MBTB of 4K entry improves performance by 17.61% compared to 8K-entry baseline BTB.

The increase in performance is mainly due to the increase in the number of branches stored in the BTB. Only skewed indexing is able to capture 4090 branches, whereas only compression captures 5496 branches. With compression and skewed indexing, MBTB is able to capture 6266 branches, which results in improved performance. Note that the baseline BTB captures less than 2800 branches. The skewed indexing allows for better utilization of BTB entries, and it spreads the branches to multiple sets. This improves the compression scheme as many branches can map to entries which would not have been possible without skewed indexing.

With an 8K-entry BTB size, we see an opposite trend, and only skewed indexing performs better compared to using compression and with both skewed indexing and compression. This is because of the additional cycle latency added due to tag comparison when using compression. The increase in latency decreases the performance slightly from 17.08% with only skewed indexing to 16.91% with skewed indexing and compression.

5.4 Sensitivity study

Sensitivity to the number of variants: Figure 8(a) shows the performance improvement of a 4K-entry MBTB with increasing the number of variants from two to four. The third variant stores four branch instructions per entry. MBTB stores branch instructions in the third variant if the offset requires less than eight bits. The fourth variant can store eight branch instructions per entry. MBTB stores branch instructions in the fourth variant if the offset requires less than four bits.

Figure 8(a) shows that with increasing the number of variants, the performance decreases. This is due to an increase in the number of evictions with the increase in the number of variants as shown in Figure 8(b). In an MBTB with four variants, a branch that requires variant-1 can evict an entry with variant-3, resulting in evictions of multiple branches. Insertion of these evicted branches on reuse will evict more branches, resulting in higher BTB MPKI.

Storage sensitivity: Figure 9(a) shows the MBTB performance with different storage size from 2K to 8K entries. We fix the number of ways to four and increase the number of sets. 4K-entry MBTB provides a sweet spot in terms of storage savings and performance improvement. Adding 2K-entry on top of a 4K-entry MBTB improves performance by only 0.26%.

Sensitivity to the number of ways: Figure 9(b) shows the performance improvement of 8K-entry baseline BTB and 4K-entry MBTB from 2 to 16 ways normalized to an 8K-entry baseline BTB with four ways. We observe an increase in performance when doubling the ways from four to eight in baseline BTB, but a decrease in performance with 16 ways due to an increase in access latency. 4K-entry MBTB with four ways outperforms baseline BTB with different number of ways.

With a 2-way MBTB, the performance is less compared to a 4-way MBTB. This is because the number of branches captured reduces from 6266 branches with a 4-way MBTB to 5549 branches with a 2-way MBTB. With a 2-way MBTB, the skewed indexing can map branches to two sets instead of four sets with a 4-way MBTB. The reduction in the spread of branches causes a decrease in the number of branch instructions captured, resulting in lower performance.

Sensitivity to L2BTB access latency: Figure 9(c) shows the performance improvement of an 4K-entry MBTB when varying the access latency from one to five cycles. With a 4-cycle MBTB of 4K entry, the performance still improves by 14.71% compared to an 8K-entry baseline BTB. This shows that as long as the BTB is accurate in providing branch targets, a slight increase in BTB latency does not impact much on the performance.

6 SUMMARY

In this paper, we make a case for a compressed and skewed indexed L2BTB design called MicroBTB(MBTB) to mitigate BTB bottlenecks and provide storage savings. Compression helps in increasing the storage density, whereas skewed indexing helps in providing uniform BTB access across BTB sets. Contrary to the industry trend of large BTBs, we make a case for relatively small yet effective L2BTB.

Averaged across 100 industry-provided server workloads, a 4K-entry MBTB provides 17.61% performance improvement compared to an 8K-entry baseline BTB providing a storage savings of 47.5KB. Overall, MBTB is a lightweight and high-performance BTB design that can scale to future workloads with a higher branch footprint and can be easily adopted by the industry.

7 ACKNOWLEDGEMENTS

The authors would like to thank Mainak Chaudhuri, Shankar Balachandran, Anant Nori, and Niranjan Soundararajan for their valuable feedback on the earlier draft. This work is supported by the SRC grant SRC-2922.001.

REFERENCES

- [1] D. Tarjan and K. Skadron, "Merging path and gshare indexing in perceptron branch prediction," *ACM Trans. Archit. Code Optim.*, vol. 2, p. 280–300, sep 2005.
- [2] A. Seznec and P. Michaud, "A case for (partially) tagged geometric history length branch prediction," *J. Instr. Level Parallelism*, vol. 8, 2006.
- [3] R. Kumar, B. Grot, and V. Nagarajan, *Blasting through the Front-End Bottleneck with Shotgun*, p. 30–42. New York, NY, USA: Association for Computing Machinery, 2018.
- [4] A. Ansari, P. Lotfi-Kamran, and H. Sarbazi-Azad, "Divide and conquer frontend bottleneck," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 65–78, 2020.
- [5] Pellegrini *et al.*, "The arm neoverse n1 platform: Building blocks for the next-gen cloud-to-edge infrastructure soc," *Micro'20*.
- [6] Suggs *et al.*, "The amd "zen 2" processor," *IEEE Micro*, 2020.
- [7] B. Grayson, J. Rupley, G. Z. Zuraski, E. Quinnell, D. A. Jiménez, T. Nakra, P. Kitchin, R. Hensley, E. Brekelbaum, V. Sinha, and A. Ghiya, "Evolution of the samsung exynos cpu microarchitecture," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 40–51, 2020.
- [8] N. Adiga, J. Bonanno, A. Collura, M. Heizmann, B. R. Prasky, and A. Saporito, "The IBM z15 high frequency mainframe branch predictor industrial product," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 27–39, 2020.
- [9] Brooks, "What's the future of technology scaling?," 2018.
- [10] A. Seznec, "Don't use the page number, but a pointer to it," in *Proceedings of the 23rd Annual International Symposium on Computer Architecture, ISCA '96*, (New York, NY, USA), p. 104–113, Association for Computing Machinery, 1996.
- [11] Asheim *et al.*, "Fetch-directed instruction prefetching revisited," 2020.
- [12] A. Seznec, "A case for two-way skewed-associative caches," in *Proceedings of the 20th Annual International Symposium on Computer Architecture, ISCA '93*, (New York, NY, USA), p. 169–178, Association for Computing Machinery, 1993.
- [13] Intel Sunny Cove. https://en.wikichip.org/wiki/intel/microarchitectures/sunny_cove.
- [14] C. Kaynak, B. Grot, and B. Falsafi, "Shift: Shared history instruction fetch for lean-core server processors," in *2013 46th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 272–283, 2013.
- [15] I. Burcea and A. Moshovos, "Phantom-btb: A virtualized branch target buffer design," vol. 37, p. 313–324, mar 2009.
- [16] C. Kaynak, B. Grot, and B. Falsafi, "Confluence: Unified instruction supply for scale-out servers," in *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 166–177, 2015.
- [17] G. Reinman, B. Calder, and T. Austin, "Fetch directed instruction prefetching," in *Proceedings of the 32nd Annual ACM/IEEE International Symposium on Microarchitecture, MICRO 32*, (USA), p. 16–27, IEEE Computer Society, 1999.
- [18] Intel 5-level Paging and 5-level EPT. https://software.intel.com/sites/default/files/managed/2b/80/5-level_paging_white_paper.pdf.
- [19] Championship Value Prediction. <https://www.microarch.org/cvp1/>.
- [20] ChampSim. <https://github.com/ChampSim/ChampSim>.
- [21] Y. Ishii, J. Lee, K. Nathella, and D. Sunwoo, "Re-establishing fetch-directed instruction prefetching: An industry perspective," in *2021 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pp. 172–182, 2021.
- [22] A. Ros and A. Jimborean, "A cost-effective entangling prefetcher for instructions," in *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pp. 99–111, 2021.
- [23] Seznec, "The fnl+mml instruction cache prefetcher," in *IPC-1 @ ISCA'20*.
- [24] F. Bodin and A. Seznec, "Skewed associativity enhances performance predictability," in *Proceedings of the 22nd Annual International Symposium on Computer Architecture, ISCA '95*, (New York, NY, USA), p. 265–274, Association for Computing Machinery, 1995.
- [25] F. Bodin and A. Seznec, "Skewed associativity improves program performance and enhances predictability," *IEEE Transactions on Computers*, vol. 46, no. 5, pp. 530–544, 1997.
- [26] D. Jimenez, "Fast path-based neural branch prediction," in *Proceedings. 36th Annual IEEE/ACM International Symposium on Microarchitecture, 2003. MICRO-36.*, pp. 243–252, 2003.
- [27] S. Mcfarling, "Combining branch predictors," tech. rep., 1993.
- [28] S. Pakalapati and B. Panda, "Bouquet of instruction pointers: Instruction pointer classifier-based spatial hardware prefetching," in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 118–131, 2020.
- [29] A. Jaleel, K. B. Theobald, S. C. Steely, and J. Emer, "High performance cache replacement using re-reference interval prediction (rrip)," in *Proceedings of the 37th Annual International Symposium on Computer Architecture, ISCA '10*, (New York, NY, USA), p. 60–71, Association for Computing Machinery, 2010.
- [30] Broadwell Micro-architecture. <https://en.wikichip.org/wiki/intel/microarchitectures/broadwell>.
- [31] First Instruction Prefetching Championship. <https://research.ece.ncsu.edu/ipc/>.
- [32] Pcacti tool, Online. Available: <https://sportlab.usc.edu/downloads/>.
- [33] Intel Sunny Cove micro-architecture. <https://www.anandtech.com/show/14514/examining-intels-ice-lake-microarchitecture-and-sunny-cove/3>.