



Evaluating Persistent Memory Range Indexes: Part Two

Yuliang He
Simon Fraser University
georgeh@sfu.ca

Duo Lu
Simon Fraser University
luduol@sfu.ca

Kaisong Huang
Simon Fraser University
kha85@sfu.ca

Tianzheng Wang
Simon Fraser University
tzwang@sfu.ca

ABSTRACT

Scalable persistent memory (PM) has opened up new opportunities for building indexes that operate and persist data directly on the memory bus, potentially enabling instant recovery, low latency and high throughput. When real PM hardware (Intel Optane Persistent Memory) first became available, previous work evaluated PM indexes proposed in the pre-Optane era. Since then, newer indexes based on real PM have appeared, but it is unclear how they compare to each other and to previous proposals, and what further challenges remain.

This paper addresses these issues by analyzing and experimentally evaluating state-of-the-art PM range indexes built for real PM. We find that newer designs inherited past techniques with new improvements, but do not necessarily outperform pre-Optane era proposals. Moreover, PM indexes are often very competitive with or even outperform indexes tailored for DRAM, highlighting the potential of using a unified design for both PM and DRAM. Functionality-wise, these indexes still lack good support for variable-length keys and handling NUMA effect. Based on our findings, we distill new design principles and highlight future directions.

PVLDB Reference Format:

Yuliang He, Duo Lu, Kaisong Huang, and Tianzheng Wang. Evaluating Persistent Memory Range Indexes: Part Two. PVLDB, 15(11): 2477 - 2490, 2022.

doi:10.14778/3551793.3551808

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/sfu-dis/pibench-ep2>.

1 INTRODUCTION

The recently commercialized persistent memory (PM) devices, represented by Intel Optane Persistent Memory (Optane PMem) [16], deliver persistence, high capacity, lower cost and fast speed on the memory bus. There have been many PM-inspired (re)designs in data-intensive systems [3, 4, 9, 45, 55, 63, 76]. In particular, much exciting progress has been made on devising single-level persistent OLTP indexes that directly operate and store data on PM without involving the storage stack, even before real PM devices became available [2, 11, 12, 27, 40, 56, 64, 68, 71, 73]. As we have shown in the past [44], although these pre-Optane proposals do not perform as expected (e.g., much slower) on real PMem devices due to inaccurate assumptions and emulation, several building blocks (unsorted

nodes, fingerprinting, hybrid DRAM-PM structures) have proved useful for devising new indexes on PM.

The availability of real PM devices has further enabled a new breed of PM-based indexes [13, 34, 46, 48, 49, 77], which are tailor-made for Intel Optane PMem. Although they appear/claim to perform better than pre-Optane proposals on real PM, it remains unclear (1) *how they compare against each other*, as these new indexes appeared roughly concurrently and/or were proposed by different research communities (e.g., VLDB/SIGMOD vs. SOSP/OSDI), (2) *how they are different from (or similar to) the pre-Optane proposals* (i.e., what “legacy” have pre-Optane proposals left for the new breed?), and (3) *what further challenges and opportunities remain in this area*. The goal of this paper is to answer these questions, which is the key to pushing actual adoption of these new indexes and PM-based systems in general. We do so by conducting a comprehensive evaluation of five state-of-the-art PM indexes, including DPTree [77], μ Tree [13], LB⁺-Tree [46], ROART [49] and PACTree [34]. We also include as a reference FPTree [56], a representative and arguably the best-performing design from the pre-Optane era [44].

This work can be seen as a sequel to the “first episode” of our past work which benchmarked and distilled the aforementioned useful building blocks from pre-Optane era PM indexes [44], but is not a mere repeat of what was done before with newer indexes. On the one hand, we follow the similar methodology and focus on representative range indexes. On the other hand, in this paper we give a snapshot of the latest state of this area and highlight new findings, observations and perspectives that were often omitted in the past, especially on variable-length key support, NUMA-awareness and new potential impact of existing PM indexes on future work. We summarize our findings below. ❶ The key optimization target remains to be reducing read and (more often) write operations on PM to preserve PM bandwidth (and thus achieving higher performance). ❷ As we predicted in past work, several pre-Optane techniques are effective on real PM hardware and widely adopted by new proposals, including fingerprinting [56], unsorted nodes and leveraging DRAM. ❸ Although modern PM-enabled servers are often multi-socket, most state-of-the-art PM indexes are still only optimized for single-socket machines and consciously avoided NUMA effect in their experimental evaluations; handling NUMA effect remains an unsolved problem. ❹ PM programming infrastructure (e.g., PM allocators and runtime) is still far from ideal in many cases and requires further improvements. ❺ Finally and perhaps most profoundly, for the first time, we observe that techniques employed by PM indexes can also be useful in devising high-performance volatile in-memory indexes. Surprisingly, when running on DRAM without the extra fences and cacheline flushes, under certain workloads a PM index can match or even outperform well-tuned indexes specifically designed for DRAM, such as HOT [6] and Masstree [50]. This highlights the potential of unifying persistent and volatile indexing to simplify the design and implementation of future systems.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 15, No. 11 ISSN 2150-8097.
doi:10.14778/3551793.3551808

The remaining sections expand on more details and insights. We give the necessary background on PM in Section 2, followed by a review of pre-Optane PM range indexes in Section 3. Sections 4–5 then survey and analyze state-of-the-art PM range indexes. Sections 6–7 present our experimental results and observations. We give an outlook of future indexes in Section 8. Section 9 covers related work and Section 10 concludes this paper. We have open-sourced our code and results at <https://github.com/sfu-dis/pibench-ep2>.

2 PERSISTENT MEMORY

We briefly review the properties of PM (in particular Intel Optane PMem¹) and its implications on software; readers already familiar with the literature may skim and fast forward to Section 3.

Collectively, “PM” refers to a class of devices that offer byte-addressability (like DRAM), high capacity and persistence (like SSDs) on the memory bus. They can be built using various materials and techniques, such as PCM [70], STT-RAM [24] and memristor [62]. However, only the 3D XPoint [17] based Optane PMem so far has delivered high capacity promised by the PM vision. Thus, we target Optane PMem in this paper. Optane PMem’s performance generally falls between DRAM and SSDs.² The first generation (100 series) exhibits ~300ns latency and ~5–40GB/s bandwidth depending on the access pattern, and with sequential/read accesses being faster than random/write accesses. The latest 200 series further gives ~30% higher bandwidth [31].

Although PM offers persistence, it is still behind multiple levels of CPU caches. Software normally goes through CPU caches to access data stored on PM using load and store instructions, and explicitly issues cacheline flush instructions (e.g., CLWB and CLFLUSHOPT) and fences to ensure data is correctly persisted [30]. However, PM-capable platforms use asynchronous DRAM refresh (ADR) to guarantee data flushed from the CPU caches will first land on the CPU’s write buffer, which is power failure protected. As a result, a PM write is considered complete once the data is forced to the ADR domain, without necessarily having arrived at physical PM media. More recent platforms further feature enhanced ADR (eADR) which also protects the CPU cache, effectively providing durable CPU caches and potentially sparing the need of cacheline flush instructions (although fences are still needed for correct ordering). eADR is only available for the more recent 200-series Optane PMem and Skylake platforms [30]. Our experiments are based on the 100-series PMem and Cascade Lake platform without eADR (detailed setups in Section 6). However, we do not expect our conclusions to change based on recent results obtained from evaluating the impact of eADR [22]. We leave it as future work to document eADR’s detailed behavior and impact on PM range indexes.

Optane PMem can operate under the Memory, App Direct, or Dual modes [30]. The Memory mode uses the system’s DRAM as a hardware-controlled cache and presents bigger but volatile memory. The App Direct mode enables persistence, thus allowing building persistent indexes. The Dual mode combines both by allowing part of PM to be configured for the Memory or App Direct mode. Same as other work, we focus on the App Direct mode as it gives software the flexibility to use DRAM and PM as needed.

¹Also known as Optane DC Persistent Memory Module (DCPMM) [28].

²Exceptions apply under certain workloads and hardware configurations [26].

3 PREVIOUSLY ON PM RANGE INDEXES

Indexing had received much attention even before real PM was available [2, 11, 12, 27, 40, 56, 64, 68, 73]. This section reviews pre-Optane designs to set the stage for our new evaluations.

3.1 Early Assumptions in the Pre-Optane Era

Due to the lack of real devices, early proposals had to “guesstimate” the properties of PM based on prototypes and simulations [39, 70, 75].³ Common assumptions included (1) limited write endurance, (2) 3–5× higher latency but similar bandwidth to DRAM’s [58, 67], (3) writes slower than reads, (4) 8-byte atomic PM writes [15], and (5) volatile CPU caches and reorderings by the CPU. Out of these, the assumption on bandwidth turned out to be inaccurate: Optane PMem’s lower-than-DRAM bandwidth is a major factor that limits performance [44]. The speed gap between sequential and random accesses was also largely left out. Endurance so far has not been a major issue for Optane PMem which warrants virtually unlimited endurance during the usual replacement cycle of 3–5 years [28, 51].

Nevertheless, these assumptions rightfully suggested that classic in-memory indexes would not guarantee correctness (customized recovery protocols are necessary), nor perform well on PM. Concurrency control must be carefully considered given PM’s higher latency. These led to the development of numerous PM indexes even before actual PM hardware was available. The key is reducing PM read/write operations and avoiding unnecessary cacheline flushes and fences to both improve performance and reduce wear.

3.2 Presumed Designs and Building Blocks

To reduce PM accesses, several building blocks have been proposed around re-designing the tree architecture, node structure and concurrency control. Now we discuss them in the context of previously evaluated PM indexes (wBTree [12], BzTree [2], FPTree [56] and NV-Tree [73]). Table 1 lists their main design choices. The first three dimensions (architecture, node structure and concurrency) received the most attention in the past; in this work, we also consider variable-length keys, PM management and NUMA-awareness.

Tree Architecture: PM-Only → DRAM-PM Hybrid. Most pre-Optane proposals adapt in-memory B+-trees. Some (e.g., wBTree and BzTree) place the entire tree on PM to allow instant recovery. But doing so leads to much slower traversal speed compared to DRAM indexes. A key contribution by FPTree and NV-Tree was to leverage the fact that B+-tree’s inner nodes only guide search traffic and are reconstructible using leaf nodes. This allows improved traversal speed by loosening the consistency requirements for inner nodes, by omitting flushes/fences [73] and/or placing all the inner nodes in DRAM [56]. Upon restart, the inner nodes can be rebuilt using B+-tree bulk loading algorithms. The downside is recovery time can scale with data size, sacrificing instant recovery.

Node Structure: Sorted → Unsorted. Traditional B+-trees keep keys sorted for fast binary search. The drawback of inheriting this design on PM is insertions may shift keys, causing excessive PM reads and writes, while having the risk of incomplete updates upon hardware failure. Moreover, binary search becomes less (or not at all) beneficial with small nodes commonly used by in-memory

³Various materials can be used to build PM, yet they may perform differently. This partially forced past work to make general assumptions for potentially wider applicability.

Table 1: Design choices of pre-Optane PM index proposals (wBTree, BzTree, FPtree and NV-Tree) [44]. Common building blocks such as using DRAM to store reconstructible data (e.g., inner nodes), using unsorted nodes and fingerprinting have proved to be useful on real Optane PMem. The impact of variable-length keys, PM allocator and NUMA effect were largely unconsidered.

	Architecture	Node Structure	Concurrency	Var. Keys	PM Allocator	NUMA-Aware
wBTree [12]	B+-tree; PM-only	Unsorted; indirection	Single-threaded	Pointer	Emulation/PMDK	No
BzTree [2]	B+-tree; PM-only	Partially unsorted leaf	Lock-free with Persistent MwCAS [68]	Inlined	Emulation/PMDK	No
FPtree [56]	B+-tree; DRAM (inner) + PM (leaf)	Unsorted leaf; fingerprints	HTM (inner) + locking (leaf)	Pointer	Customized/PMDK	No
NV-Tree [73]	B+-tree; PM-only or optionally DRAM+PM	Unsorted leaf; inconsistent inner nodes	Locking	Pointer	Emulation/PMDK	No

B+-trees. A popular solution is to keep nodes unsorted and use a linear scan to retrieve the target key. Most of the surveyed pre-Optane indexes adopted this technique. To mitigate the impact of linear search, BzTree periodically consolidates nodes to become sorted. FPtree proposed fingerprinting which maintains one-byte hashes of the keys in the node and a lookup starts by checking the fingerprints. Only the keys with matched fingerprints will be further examined. This greatly reduces PM accesses, especially for negative search where the key does not exist. Another approach is to keep an indirection array [12] that stores sorted index positions of keys, allowing binary search using the indirection array.

Concurrency Control: Pessimistic \rightarrow Optimistic. PM indexes often prefer lightweight concurrency control over pessimistic lock coupling [59]. FPtree uses separate strategies for inner and leaf nodes: For the former it leverages hardware transactional memory (HTM) to reduce traversal costs, and for the latter it uses traditional locking because inserting into PM-resident leaf nodes may involve cacheline flushes, which in turn will abort HTM transactions. BzTree uses lock-free multi-word compare-and-swap (PMwCAS) [68] that can atomically modify multiple 8-byte words. Without pessimistic locking, FPtree scales to high core count but exhibits high tail latency and low throughput under contention, due to HTM’s inherent limitations [44]. FPtree delegates the detailed HTM-locking interactions to Intel TBB which uses a fast-path that uses HTM and a slow-path that serves as a fallback when HTM abort rate becomes higher than a threshold (10 by default).⁴

3.3 Functionality and PM Management

Now we turn to the remaining three dimensions under consideration in Table 1: support for variable-length keys, NUMA awareness and PM programming infrastructure support.

Most pre-Optane proposals [12, 56, 73] focused on handling fixed-length, 8-byte keys, so did past evaluation efforts. Variable-length keys are usually supported using pointers to keys stored in the (persistent) heap. Some proposals differentiate pointers from inlined values by designating a special “type” bit in the 8-byte key area, effectively limiting the maximum length of keys to be 63 bits. Others, e.g., FPtree, require compile-time customization by specifying whether keys are 8-byte integers or pointers to support full 64-bit

keys. Pointers are used in case both types are required. As a result, accessing shorter keys ≤ 8 bytes is also prone to cache misses caused by pointer chasing. Out of the four indexes, only BzTree provides inlined support for variable-length keys with slotted pages.

None of the surveyed pre-Optane indexes handles NUMA effect. Most of them also do not have a well thought-out design for managing PM. FPtree uses PMDK [29], the current de facto standard PM library, to avoid issues such as PM leaks. For better performance, FPtree has to use customized slabs (large chunks allocated from PMDK allocator) to amortize PM allocation cost. Still, our previous work [44] has identified PM allocation in all these indexes as a main bottleneck that should be avoided by future designs.

4 STATE-OF-THE-ART PM RANGE INDEXES

We survey five representatives of recent PM indexes optimized for Optane PMem: LB⁺-Tree [46], DPTree [77], μ Tree [13], ROART [49] and PACTree [34]. They can be categorized as B+-tree based, trie based and hybrid which makes use of both B+-trees and tries.⁵

4.1 B+-Tree based: LB⁺-Tree and μ Tree

We survey two representative B+-tree based PM range indexes, LB⁺-Tree [46] and μ Tree [13], which represent designs that mainly optimize for high throughput and low tail latency, respectively.

LB⁺-Tree. As evaluated by previous work [33, 44, 72], Optane PMem uses 256-byte granularity internally: accesses smaller than 256 bytes will still incur 256-byte of traffic to/from the physical media. It then becomes important to (1) coordinate PM accesses in 256-byte granularity and (2) reduce unnecessary PM accesses (to save PM bandwidth) [33, 72]. To reach these goals, as shown in Tables 1 and 2, LB⁺-Tree starts with several useful techniques proposed by FPtree: fingerprinting, DRAM+PM architecture and optimistic concurrency with HTM and locking. Compared to FPtree, it uses hand-rolled RTM transactions instead of TBB. We highlight the impact of this design decision later in Section 6.

On top of these existing techniques, LB⁺-Tree proposes new ones to further optimize PM accesses. First, node sizes are set to align with and be multiples of 256 bytes for better CPU cache and PM utilization. Second, to minimize PM writes during inserts, LB⁺-Tree

⁴Details at https://github.com/oneapi-src/oneTBB/blob/v2021.5.0/src/tbb/rtm_mutex.cpp#L33. Our evaluation (Section 6) sets the threshold to 256 for better performance.

⁵In addition to traditional indexes, learned indexes [36] are also being adapted for PM [47]. However, they are still in very early stage. We thus leave it as future work to evaluate them to avoid pre-mature conclusions.

proposes entry moving to bound the number of cacheline writes per insert. As Figure 1(a) shows, a 256-byte leaf node is divided into four 64-byte cachelines. Upon insert, LB⁺-Tree first attempts to insert the record into Line 0 if an empty slot is available in it. Since the header is also in Line 0, only one cacheline write to PM is needed to persist both the record and header. Otherwise, LB⁺-Tree moves data from Line 0 to another line where the new entry is inserted. This proactively spares empty slots in the first cacheline, which will reduce flush operations incurred by future inserts.

Write-ahead logging (WAL) is widely used for crash recovery [29], but incurs additional PM writes. LB⁺-Tree disposes of WAL with logless node splits by storing two sibling pointers in each leaf node and uses an alt bit in node header to indicate the valid pointer. Upon split, a new leaf node is first allocated and tracked by the unused pointer, followed by a redistribution of entries to the last two cachelines of the new node. Then the alt bit and bitmap in the original node are updated using an 8-byte atomic PM write to ensure the tree is always in a consistent state even across failures.

μTree. Different from most PM indexes, μTree mainly optimizes for tail latency caused by PM’s high latency. It places B+-tree inner nodes for tail in DRAM, but redesigns leaf nodes to use both DRAM and PM. As shown in Figure 1(b), leaf nodes in μTree consist of two layers: an array layer and a list layer. The former consists of traditional B+-tree nodes which store pointers to list layer nodes; the latter is a PM-resident singly-linked list where each node stores a key-value pair. μTree uses optimistic locking [43] for the DRAM-resident B+-tree and manages the PM-resident linked list in a lock-free manner. To insert a record, the thread first traverses the in-DRAM B+-tree optimistically without holding any locks, and then inserts a node representing the key-value pair in the linked list using the compare-and-swap (CAS) instruction [20, 23]. It then acquires the lock in the corresponding leaf node in the B+-tree to insert the key, which may trigger splits that will in turn acquire locks from the bottom up in DRAM. This way, B+-tree structural modification operations (SMOs) and actual key-value inserts/removals (which incur PM flushes in the list layer) are decoupled, and multiple threads inserting into the same leaf node could proceed in parallel in the list layer. This could allow more concurrency and removes flushes from the critical path, thus reducing tail latency.

4.2 Trie based: ROART and PACTree

Newer designs have also explored tries on PM. They are mostly PM-only (instead of DRAM-PM hybrids) and based on ART [42]. We survey two representative proposals, ROART [49] and PACTree [34].

ROART. Based on ART, ROART optimizes for range scans, which are known to be a weak point of tries [49]. Figure 1(c) shows the overall architecture of ROART. It is purely on PM but selectively persists metadata entries and reconstructs the inconsistent ones upon recovery (similar to NV-Tree’s selective consistency [73]). To optimize scan, ROART compacts small sub-trees (< 64 entries) into leaf arrays that store pointers to records. This reduces pointer chasing overhead during scan and makes the index shallower. The downside is splits become expensive, as all the keys in a leaf array are divided into subsets based on the first differentiating byte, followed by a node allocation for each subset. This requires more PM writes and fences, and puts higher pressure on the PM allocator.

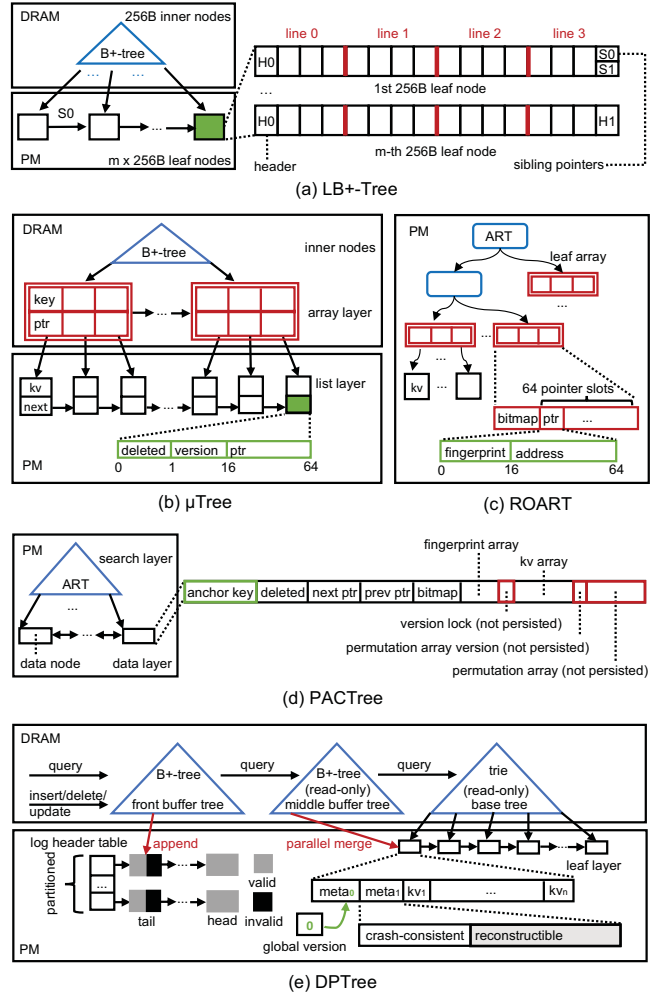


Figure 1: Architecture of five state-of-the-art PM indexes.

To this end, ROART reduces the number of fences by relaxing the order of split steps and using a depth field to detect and resolve inconsistency. It also proposes delayed check memory management (DCMM) to cope with the high PM allocation demand. DCMM performs fast allocations using thread-local pools but delays garbage collection by traversing the entire index at a later time, potentially increasing PM usage. Since ROART is fully in PM, it supports true instant recovery, and natively supports variable-length keys without extra pointer chasing using its trie-based design.

PACTree. PACTree is a PM-only two-layer persistent trie with B+-tree styled leaf nodes. As shown in Figure 1(d), PACTree consists of a search layer and a data layer. The search layer is a durable trie based on concurrent ART that uses read-optimized write exclusion (ROWEX) [43]. The data layer is a doubly-linked list of B+-tree leaf nodes, each of which contains 64 key-value pairs and an anchor key to indicate the smallest key in the node. PACTree stores fingerprints and indirection arrays in leaf nodes to facilitate search and scan, but they are not persisted to reduce PM writes. Upon split, the target leaf node in the data layer is first locked, and then a log entry is

Table 2: Main design choices of state-of-the-art PM range indexes (LB⁺-Tree, DPTree, μ Tree, ROART and PACTree). They base on Intel Optane PMem and inherit designs from pre-Optane proposals, by using DRAM (sometimes more aggressively), lightweight concurrency control and unsorted nodes. Some also advocate customized PM allocators to reduce PM accesses. Variable-length keys and NUMA-awareness are still less considered.

	Architecture	Node Structure	Concurrency	Var. Keys	PM Allocator	NUMA-Aware
LB ⁺ -Tree [46]	B+-tree; DRAM (inner) + PM (leaf)	Unsorted leaf; fingerprints; extra metadata	HTM (inner) + locking (leaf)	Pointer	Customized/PMDK	No
μ Tree [13]	B+-tree; DRAM (B+-tree) + PM (linked list)	Sorted	Locking (array layer) + lock-free (list layer)	Pointer	PMDK	No
DPTree [77]	Hybrid; DRAM (B+-tree, trie inner) + PM (trie leaf)	Unsorted leaf; fingerprints; indirection; extra metadata	Optimistic lock [43] + async. updates	Pointer	PMDK	No
ROART [49]	Trie; PM-only	B+-tree like unsorted leaf; fingerprints	ROWEX [43]	Inlined	Customized/PMDK	No
PACTree [34]	Trie; PM-only or optionally DRAM+PM	Unsorted leaf; fingerprints; indirection	Optimistic lock [43] + async. update	Inlined	Customized/PMDK	Yes

written to a per-thread SMO log in PM. The thread then splits the leaf node and commits without modifying inner nodes in the search layer. A background thread will then finish the remaining SMO in the search layer. This allows worker threads to commit early right after modifying leaf nodes. However, it also creates inconsistencies between the search and data layers. Thus, query threads may need to perform a “last-mile” search after arriving at the data layer to find the correct leaf node using anchor keys. PACTree is the only index that mitigates NUMA effect. It uses separate pools for the search layer, data layer and logs in each NUMA node. It also advocates the use of snooping-based CPU coherence protocols (instead of the default directory-based protocol on most platforms) to avoid poor performance when PM accesses cross NUMA boundaries.

4.3 Hybrid: DPTree

As shown in Figure 1(e), DPTree is a PM-DRAM hybrid index that combines up to two B+-trees (a *front* and a *middle* buffer tree) in DRAM, a trie (*base* tree) that places inner nodes in DRAM and leaf nodes in PM. To search for a key, DPTree first visits the front buffer tree. If the target key is not found, the middle buffer tree (if exists) will be further searched. If the key does not exist in the buffer trees, the base tree will have to be searched. To reduce unnecessary traversals, DPTree maintains a bloom filter per buffer tree. For range queries, DPTree has to search and merge results from all the trees. When the size ratio between the front buffer tree and base tree reaches a pre-defined threshold, DPTree creates a new front buffer tree and turns the previous front buffer tree into a middle buffer tree. Tree merge operations are triggered when the size ratio between the front buffer tree and the base tree reaches a pre-defined threshold, and are performed by background threads. Using the version number and extra set of metadata, DPTree ensures changes are invisible to concurrent queries when a merge is in progress. After merging, the middle buffer tree is destroyed and the inner nodes of the base tree (ART) are rebuilt. Then the global version bit is flipped to expose changes to incoming requests. DPTree also uses selective metadata persistence with reconstructible metadata (e.g., record count and fingerprints) in DRAM.

5 ANALYZING THE STATE-OF-THE-ART

With the high-level designs laid out, now we analyze the new PM indexes in detail and distill common building blocks which can be useful for future PM indexes. We analyze them from the six dimensions in Table 2, followed by empirical evaluation in Section 6.

5.1 Index Architecture

New PM indexes often inherit the PM+DRAM architecture with new optimizations, and consider tries and hybrid structures.

(More Extensive) Use of DRAM. New PM indexes based on B+-tree and hybrid structures (LB⁺-Tree, DPTree and μ Tree) continue to use DRAM to store part of the index. We also observe more aggressive use of DRAM, e.g., DPTree and μ Tree place entire tree structures in DRAM to get more performance gains. The tradeoffs are (1) longer recovery time, (2) more complex programming and (3) higher DRAM consumption which we quantify in Section 6.

Beyond B+-Trees and Monolithic Indexes. New PM indexes also adapt tries (PACTree and ROART), but they default to pure PM designs, potentially leading to suboptimal performance but preserving instant recovery. Notably, PACTree stores keys in both inner and leaf nodes, but it is possible to place inner nodes in DRAM to improve performance. We expect to see more B+-tree based PM indexes utilize DRAM for faster access/write speed, as PM servers will still feature DRAM in the foreseeable future.⁶ As a major departure from just adapting one type of data structure, DPTree combines B+-trees and tries. When it comes to node structure, new indexes are also introducing new designs that no longer use pure trie or B+-tree nodes, which we highlight next.

5.2 Node Structure

New PM indexes base their designs on Optane PMem with node alignment of 256 bytes to reduce unnecessary PM accesses. Several pre-Optane designs—fingerprinting, unsorted (leaf) nodes and selective consistency for metadata—continue to be used by new PM indexes. But they are further optimized with new techniques.

⁶DRAM must be present for PMem to work, even if the software does not need it [26].

Fingerprinting on Steroids. As PM accesses are slow, fingerprinting becomes the most popular approach used by four out of the five new PM indexes. Meanwhile, new techniques are introduced to better store and use fingerprints. LB⁺-Tree uses SIMD instructions to compare up to 64 one-byte fingerprints in one instruction. ROART embeds a two-byte fingerprint inside pointers to key-value pairs, minimizing pointer chasing at the leaf level.

Extra and Selectively Persisted Metadata. LB⁺-Tree and DP-Tree both use an extra set of metadata per leaf node to avoid logging (thus reducing PM writes). For LB⁺-Tree, this allows it to achieve logless split. DP-Tree uses the extra metadata to track PM allocations and hide incomplete changes during tree merge operations.

Not all metadata entries have to be persisted when modified. For instance, version locks in PACTree leaf nodes are only meaningful at runtime; fingerprints and indirection arrays can be rebuilt during recovery (DP-Tree and PACTree) or on demand at runtime by query threads (ROART). Keeping them volatile can significantly reduce PM writes and improve performance, at the cost of slower recovery.

Hybrid Leaf Nodes. All the surveyed PM indexes that adopt trie (DP-Tree, ROART and PACTree) use B+-tree like leaf nodes where a node stores multiple records to reduce insert overhead, as the leaf node is only split when it is full. This design also reduces pressure on the PM allocator, as trie-based indexes typically incur more frequent allocations with varying node sizes compared to B+-trees. Scan performance is also improved with less pointer chasing. Different from other proposals, μ Tree introduces a linked list layer in PM, making its “leaf node” logical. Although this design decouples SMOs and data movement to potentially enable more parallelism, it adds more overhead to scans due to more pointer chasing.

5.3 Concurrency Control

All the surveyed new PM indexes use optimistic concurrency control. They optimize traversals using lock-free read or HTM for inner nodes; locks are only acquired at the leaf level and/or as needed in inner nodes to reduce PM writes. Further, they tend to use background threads for SMOs (e.g., PACTree and DP-Tree). The benefit of offloading SMOs to the background is potentially lower latency for index operations. However, it can be tricky to determine the appropriate number of background threads. Also, with a given CPU budget (e.g., in the cloud), the machine may not have enough resources to spare for the background threads, which may then fall behind the foreground threads and affect the overall progress.

5.4 Functionality and PM Management

As Table 2 lists, among the new PM indexes, only trie-based ROART and PACTree natively support variable-length keys; the others follow pre-Optane proposals to use pointers to keys. With real hardware and libraries like PMDK, all the indexes have taken into account PM management issues (e.g., avoiding persistent leaks and optimizing allocation performance). However, many need a customized allocator for performance reasons. Finally, only PACTree is designed to mitigate NUMA effect. Other indexes would have to use general-purpose approaches [66] that can be applied on any PM index, but they come with limitations (e.g., focus on certain workloads). Such facts indicate that in terms of functionality, new PM indexes have been mainly sticking with the status quo.

6 EVALUATING THE STATE-OF-THE-ART

Now we empirically evaluate new PM indexes and compare them with FPTree, the best-performing pre-Optane PM range index.

6.1 Experimental Setup

We run experiments on a 40-core (80-hyperthread) server equipped with two Intel Xeon Gold 6242R CPUs clocked at 3.10 GHz with 36MB of cache. The server is fully populated with 12×32GB DRAM DIMMs (384GB in total) and 12×128GB Optane PMem DIMMs (1.5TB in total) for maximum bandwidth. Both DRAM and PMem run at 2666MT/s.⁷ The server runs Arch Linux with kernel 5.14.9. All the code is compiled using GCC 11.1 with all the optimizations. Unless otherwise specified, we use PMDK [29]/jemalloc [19] for PM/DRAM allocations.

Benchmarking Framework. We use PiBench [44], a unified framework for benchmarking PM indexes, to stress test the indexes. PiBench generates and issues synthetic workloads of given distributions that consist of user-specified operations (lookup/insert/update/delete/scan). It requires each index implement a set of common interfaces, by extending an abstract C++ class. We create a wrapper for each index that uses PiBench’s interfaces to invoke the index’s internal operations. The wrapper is compiled as a shared library and loaded into PiBench’s address space at runtime.

Metrics and Workloads. We measure throughput (operations per second) and latency at various thread counts. Using PiBench, we collect statistics such as cache misses and bandwidth to aid analysis. We test both fixed-length (8-byte) integer and variable-length keys, following the setups used by previous work [44]. For point queries we use keys chosen randomly under uniform or skewed distributions; for range scans, we uniform randomly choose a start key K and scan 100 records following K . We prefill each index with 100 million key-value pairs, after which we start to run individual and mixes of index operations. The results across runs vary little and so we report the average of three 10-second runs.

6.2 Index Implementations and Parameters

For all the indexes (except FPTree which we had to implement), we use the original authors’ code obtained from their public repositories. We use parameters that lead to the best performance (described below) and make necessary changes to each index for correctness, functionality and fairness.

LB⁺-Tree. (1) The original insert and delete functions do not guarantee persistence; we followed the authors’ suggestions to fix them.⁸ (2) We implemented range scan (similar to FPTree) and update with locking in leaf nodes. (3) We found the PMDK allocator can provide sufficient performance after tuning PMDK parameters and allocating 256-byte aligned leaf nodes, so we also use PMDK for LB⁺-Tree. The inner/leaf nodes contain 15/14 entries (256-byte).

μ Tree. Instead of using chunk-based allocation, the original code in fact uses PMDK’s POBJ_ZALLOC. We changed it to use `pmemobj_allloc` for much better performance. We were not able to verify the correctness of multi-threaded inserts (with keys missing after successful inserts), so we only include μ Tree in single-threaded experiments. Inner/leaf node sizes are both set to 29 entries.

⁷DRAM has to be clocked down from 3200MT/s to 2666MT/s for PMem to work [28].

⁸Details at <https://github.com/schencoding/lbtree/issues/2>.

ROART. The open-source code of ROART supports both PMDK and its customized DCMM allocator [49]; we present numbers under both allocators. Leaf node size is set to 64 entries.

PACTree. We use PACTree’s own NUMA-aware PM allocation. For fair comparison, we pin the background and worker threads to the same CPU cores so that all the indexes use the same amount of CPU cores. Node size is set to 64 entries.

DPTree. The original code misses PM allocator support, so we ported it to use PMDK. We follow the original paper to use an equal number of worker and merge threads. For fair comparison, we also pin the merge and worker threads to the same cores (similar to PACTree’s setup). Inner and leaf nodes contain 31 entries for buffer trees; the base tree uses 256-entry leaf nodes.

FPTree. Since the original implementation is proprietary, we implemented FPTree by strictly following the paper.⁹ We have verified that our implementation performs similarly to the original author’s binary does. We follow past work’s recommendations to set inner/leaf node sizes to 128/64 entries [44, 56].

6.3 Single-threaded Performance

We begin with single-threaded experiments to avoid concurrency complicating our analysis. We run lookup/insert/update/scan operations under the uniform random distribution and report throughput.

Point Queries. LB⁺-Tree, DPtree and μ Tree perform similarly for lookups in Figure 2(a). They are up to $\sim 2\times$ faster than FPTree, ROART and PACTree, which also perform similarly. Overall, DPtree performs the best under single-threaded lookups, largely because of its extensive use of DRAM: if the search key is in the DRAM-resident buffer tree, the entire query can finish without ever accessing PM. If the base tree needs to be visited, only the search in leaf node will incur PM access, which is mitigated by binary search. The other two faster indexes (LB⁺-Tree and μ Tree) also benefit from placing inner nodes and the array leaf layer in DRAM. LB⁺-Tree also extensively uses SIMD instructions and prefetching, which as shown in Figure 3(a) drastically reduces cache misses and our factor analysis showed that prefetching alone improves performance by $\sim 10\%$. For inserts, B+-tree based FPTree/LB⁺-Tree and hybrid DPtree are up to $\sim 2\times$ faster than trie-based ROART and PACTree in Figure 2(b). μ Tree’s use of linked lists in the leaf level cancels out some of B+-tree’s advantages due to high cache miss rates in Figures 3(b–c). The performance of PM allocators is critical for ROART as for each update it needs to allocate a new PM block, and using its own DCMM can double the throughput for updates in Figure 2(c). Compared to lookups, updates incur additional PM writes to update the payload, but will not trigger SMOs compared to inserts. As expected, the update performance of all indexes falls between their lookup and insert performance, but follows the trend of insert performance more closely because (1) PM exhibits higher write latency, and (2) similar to inserts, leaf-level locks can only be released after the new value is flushed, adding delays (although being a constant amount of overhead under a single thread).

Range Scans. As Figure 2(d) shows, range scan performance depends largely on the cost of scanning within and across leaf nodes, i.e., whether the nodes are sorted and they are big or small. Although DPtree needs to search multiple trees and combine results, it is




Figure 2: Single-threaded throughput (uniform distribution). Overall, DPtree and LB⁺-Tree perform the best. FPTree can be very competitive to (or even better than) newer indexes.

still the fastest. DPtree’s leaf nodes use indirection arrays, so the result can be returned directly without sorting, as oppose to trees that use unsorted nodes, e.g., FPTree. LB⁺-Tree inherited a lot from FPTree, but is $\sim 30\%$ slower than FPTree for scans, because its leaf nodes are smaller (14 entries). To scan for the same number of records, compared to FPTree which uses 64-entry nodes, more leaf nodes have to be visited by LB⁺-Tree, causing more cache misses in Figure 3(d). This result highlights the tradeoff between hardware consciousness and optimization goals: using small nodes allows LB⁺-Tree to perform well in point queries, but can penalize scans.

All the tested indexes except ROART first copy the scan results to an array and then optionally sort them before they are returned to the user. ROART, however, first returns an array of leaf pointers without copying. Moreover, to support variable-length keys, it stores in leaf nodes pointers to keys. This mandates the sorting pass to dereference pointers to keys, causing extra cache misses. We note that as an optimization, if key size is known, one may change ROART to copy keys first, which in our tests can double the performance at the cost of supporting variable-length keys. Finally, μ Tree exhibits low scan performance because every record is stored in a linked list node, traversing them results in many cache misses.

Summary. The most effective technique to achieve high performance remains leveraging DRAM, which newer PM indexes adopt more aggressively, by putting more components or even complete trees in DRAM. This is at the cost of higher memory consumption, more complex recovery protocol and higher cost of ownership. Importantly, FPTree remains very competitive. PACTree and ROART are only marginally faster than FPTree for lookups. In Figure 2(b), PACTree, ROART and μ Tree are even slower than FPTree. Only DPtree and PACTree perform better than FPTree for scans.

6.4 Multi-threaded Experiments

Now we measure index throughput with different thread counts under uniform and skewed distributions. We start with one socket and expand to NUMA with two sockets in Section 6.7.

Individual Operations. Figures 4(a–d) show the throughput of lookup/insert/update/scan under the uniform distribution; the shaded areas (over 20 threads) indicate numbers obtained when hyperthreads are also used. All indexes scale well for pure lookups in Figure 4(a), with DPtree and LB⁺-Tree achieving higher raw throughput than others. This result aligns with that obtained in Section 6.3 under a single thread. DPtree uses optimistic lock coupling for its buffer/base trees, and readers can traverse without incurring PM accesses. Its bloom filter also helps avoid unnecessary

⁹Code available at <https://github.com/sfu-dis/fptree>.




Figure 3: Last-level cache misses per operation under a single (a-d) and 20 threads (e-h).

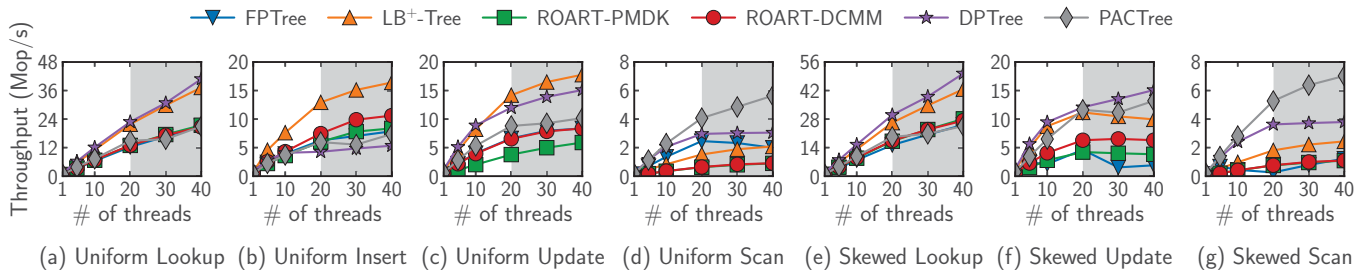


Figure 4: Throughput under uniform (a-d) and skewed (e-g, self-similar with 80% accesses on 20% of keys) distributions.

lookups in the buffer trees. LB⁺-Tree uses HTM which without write operations exhibits little/no aborts.

For inserts, although LB⁺-Tree does not perform the best under a single thread, it scales the best under multiple threads, by being $1.55 \times / 1.96 \times / 3.07 \times / 2.23 \times$ faster than ROART-DCMM/ROART-PMDK/DPTree/PACTree. Although LB⁺-Tree inherits many designs from FPTree, its logless split, new node layout and inner node locks to avoid re-traversals during split further make it $2.09 \times$ faster than FPTree. DPTree’s performance stops scaling beyond 10 threads, mainly due to its 7-phase merge: each time a merge occurs, records in the middle buffer tree will be moved into the base tree, causing the base tree’s inner nodes to be rebuilt. This in turn incurs high garbage collection costs. ROART-DCMM achieves $1.27 \times$ higher performance using its customized PM allocator compared to using PMDK. Not leveraging DRAM also contributes to its lower performance compared to others that do leverage DRAM.

For updates, LB⁺-Tree outperforms others in Figure 4(c), thanks to its fast traversal and node layout design. DPTree is also very competitive, as updates are served in-place without triggering merge operations. Similar to the single-threaded results, for all indexes, updates behave more similarly to inserts (than lookups) as leaf-level locks must be retained until the new value is flushed.

As shown in Figure 4(d), under multiple threads the relative merits of different indexes on range scan are similar to the single-threaded results. The only exception and best performing index is PACTree. It scales to 40 threads, thanks to the combination of (1) its leaf node design that inlines key-value pairs and leverages PM’s fast sequential read, (2) indirection that avoids sorting, and (3) optimistic concurrency that incurs no PM writes for reads. The other trie-based ROART performs the worst although it specifically optimizes for scan since it does not use DRAM and incurs more cache misses (Section 6.3); under high core counts, cache misses are further exacerbated in Figure 3(h). For B+-tree variants, LB⁺-Tree performs much worse for scans due to its use of small nodes (more cache

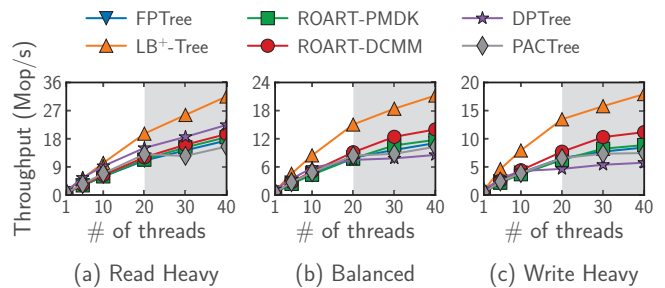


Figure 5: Throughput of mixed workloads (lookups + inserts) under uniform distribution.

misses). In contrast to the single-threaded results, FPTree performs poorly: using larger leaf reduces cache misses, but increases lock contention on leaf nodes. DPTree takes no locks for scans (OLC), but needs to visit multiple indexes and merge results, which contributes to its lower performance than PACTree.

Skewed Accesses. Under the self-similar distribution where 80% of accesses are focused on 20% of all the keys [21], lookups as shown in Figure 4(e) exhibit a similar trend to uniform distribution but with higher raw throughput because of better CPU cache utilization (the working set is smaller). For updates in Figure 4(f), DPTree remains scalable as updates are all performed in the DRAM buffer tree. FPTree shows unstable and low throughput for updates due to frequent HTM aborts and more PM writes caused by SMOs during update: updates to the same record are appended without deduplication, so the node can become full and get split during an update. No index scales under update workloads with hyper-threading. Under contention, locking takes over to become the main bottleneck in FPTree, despite the working set is smaller.

Like FPTree, LB⁺-Tree also uses locking for leaf nodes, but scales under scan because of their different ways of using HTM: FPTree

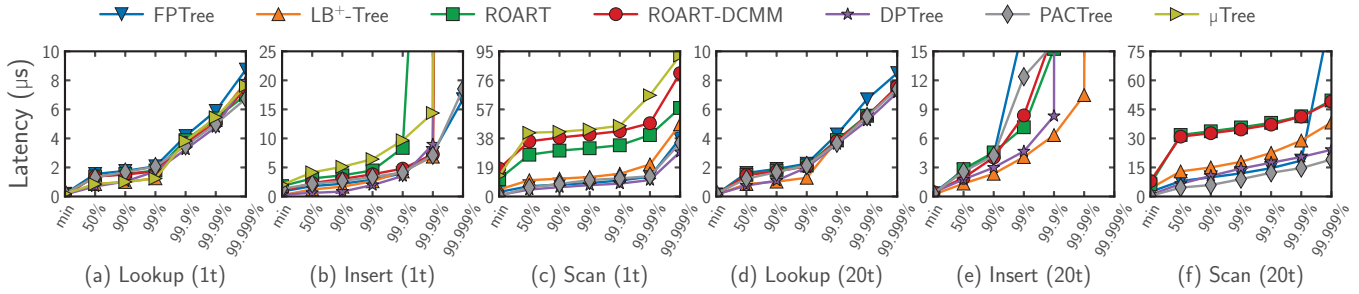


Figure 6: Tail latency of PM indexes under uniform distribution and a single thread (a–c) and 20 threads (d–f).

delegates HTM and locking to TBB (Section 3.2) which has a global fallback path after a pre-defined number (256) of aborts of HTM transactions, whereas LB⁺-Tree directly uses HTM instructions (e.g., `xbegin/xend`). Although scan is read-only, the first leaf node lock (in PM) is acquired inside the HTM transaction at the end of traversal. This incurs much contention under skewed workloads and triggers HTM aborts, leading FPTree to use the slow fallback path protected by a global mutex. DPTree’s optimistic locks allow high scan throughput under skewed accesses as no locks are acquired. PACTree scales with the best performance under skewed accesses for scans, although the gain diminishes with more hyperthreads. ROART scales slightly worse in skewed update: for each update it creates a new leaf and replaces the original leaf pointer inside the leaf array, which becomes more expensive under contention.

Mixed Workloads. We test mixed workloads with different read/insert ratios: read heavy (90% lookups + 10% inserts), balanced (50% lookups + 50% inserts) and write heavy (10% lookups + 90% inserts). As Figure 5 shows, LB⁺-Tree exhibits the best performance and scalability. DPTree scales worse with more inserts as tree merge becomes a major overhead. Finally, FPTree again remains very competitive with ROART, PACTree and DPTree.

6.5 Tail Latency

Like previous work, to balance overhead and accuracy [44] we, sample 10% of all the operations during each run under uniform distribution to rule out the impact of CPU caches. Figure 6 shows the tail latency at varying percentiles under one thread (a–c) and 20 threads (d–f). As expected, we observe no obvious differences between one and 20 threads for lookups.

For inserts, lookups and scans in Figures 6(a–c), μTree shows consistently higher latency (and skyrockets at 99.99% for inserts), although its key design goal is to reduce tail latency (Section 4.1). We observe the reason is in its use of PM-resident linked lists with per-record nodes. To handle an insert, μTree uses ~2900 cycles to allocate a list node and ~2200 cycles to complete a CAS and cacheline flush to insert the allocated node into the linked list. In contrast, LB⁺-Tree only needs around 80 cycles to insert a key into a leaf node. Moreover, the use of linked lists in the leaf layer causes many cache misses during scans: as shown in Figure 3(d), μTree exhibits the highest cache miss ratio. ROART-DCMM exhibits lower latency than ROART-PMDK for inserts, thanks to the better DCMM allocator: in the worst case, a split in ROART could allocate 63 leaf arrays and one inner node. ROART has relatively higher latency

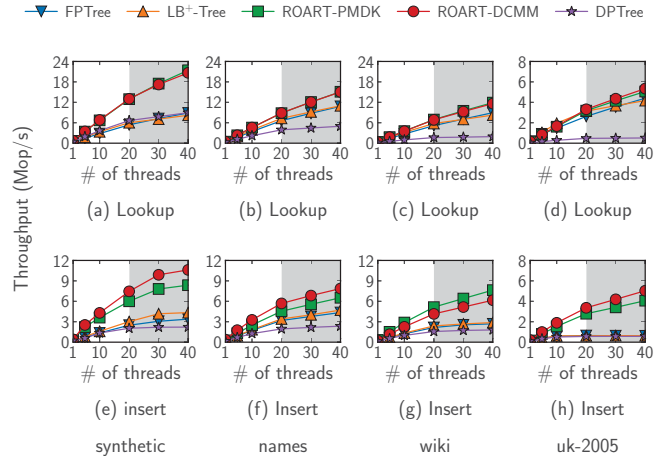


Figure 7: Throughput under variable-length keys using synthetic and real-world datasets [61].

for scans in Figures 6(c) and 6(f), due to pointer chasing at the leaf level to dereference pointers to keys. The other trie-based PACTree directly stores keys in leaf nodes, hence exhibiting lower latency. PACTree also shows relatively low latency for all operations in most cases (except inserts under 20 threads) because SMOs are offloaded to background threads. Under 20 threads, the background threads start to fall behind, requiring worker threads to traverse extra nodes to reach the correct leaf node, thus increasing latency.

6.6 Support for Variable-Length Keys

Now evaluate variable-length key support. We first run the same experiment as Section 6.4 with 8-byte keys, but force the indexes to use their variable-length key support. This may limit the depth of trie-based indexes (thus giving them advantages), but allows us to reason about the efficiency of the variable-length key support by comparing with the fixed-length key experiments; we use three real-world datasets later. For most indexes, including FPTree, LB⁺-Tree and DPTree this means allocating the key in the heap and storing a pointer to the key in the index. ROART uses its own native support to store keys in index nodes themselves. All indexes are plotted except PACTree.¹⁰ As shown in Figures 7(a) and 7(e),

¹⁰PACTree assumes null-terminated strings. This is incompatible with PiBench which may generate keys with `\0`, which will be wrongly treated as short keys.

ROART performs the best in all cases. This is exactly opposite to the case using fixed-length keys (cf. Figure 4). In particular, cache misses caused by pointer chasing (to access keys) dominate the performance of FPTree, LB⁺-Tree and DPTree.

We further test the indexes with three representative real-world datasets [61]: Reddit usernames (names) [52], Wikipedia (wiki) [69] and URLs (uk-2005) [38]. They respectively consist of string keys of up to 25, 256 and 2029 bytes; for space limitation we omit more details about the datasets which can be found elsewhere [61]. As shown in Figures 7(b-d) and 7(f-h), the gaps between ROART and FPTree/LB⁺-Tree shrink, although ROART remains advantageous. With longer keys, trie-based indexes will build more inner nodes to form deeper traversal paths, while the depth of B+-tree variants is not affected as they store pointers to keys. Among all the indexes, DPTree performs much worse with longer keys. We found a main reason is that most lookups (94% according to our profiling results) needed to traverse the base tree and search leaf nodes. Since DPTree maintains multiple indexes, it also requires more complex traversal logic that incurs extensive key comparisons (on average 34 memcmp calls vs. 22 in FPTree), further lowering its performance.

Overall, these results highlight the need to enhance variable-length key support in future PM indexes, for example by combining the best of tries and B+-trees without tradeoffs.

6.7 Impact of NUMA Effect

Now we extend our experiments to use both NUMA nodes on the server. Except for PACTree, we allocate PM and DRAM from the first socket. All the threads are pinned, and we first use all the 40 physical cores across two sockets, before using hyperthreading beyond 40 threads. This allows us to stress the indexes with inter-socket traffic and contrast their behavior with and without NUMA effect. For PACTree, we include two variants which respectively enable and disable its NUMA-aware per-node PM pools.

As shown in Figure 8, NUMA effect has major impact on all indexes' throughput, and no index scales well beyond one socket for all operations. Although not specifically designed for NUMA, LB⁺-Tree achieves the best scalability for lookups, and the highest throughput for inserts and updates (with a dropping trend beyond one socket). We attribute the reason to its frugal use of cacheline flushes and careful node layout designs. Both reduce PM accesses and cross-socket traffic. For lookups, FPTree also does not collapse, whereas the performance of other indexes fluctuate and/or drop beyond 20 threads. This implies HTM is robust to NUMA effect for read-only workloads, thanks to its lightweight conflict detection mechanism that piggybacks on the coherence protocol. Moreover, HTM can use the extra last-level cache in the second NUMA node to track reads [8, 57]. Other approaches (OLC, locking, ROWEX) are unable to take good advantage of the coherence protocol like HTM, contributing to more severe NUMA effect.

PACTree is the only PM index that takes NUMA effect into account by (1) using separate PM pools for each NUMA node and (2) leveraging snooping coherence protocol. With separate PM pools (PACTree-NUMA), PACTree maintains performance beyond one socket without collapsing, but still does not scale as expected. The main culprit is the directory-based coherence protocol that incurs additional PM accesses to update the directory. Thus, PACTree




Figure 8: Impact of NUMA effect for PM indexes. No index scales well under all operations due to additional PM accesses by the directory-based CPU coherence protocol.

Table 3: PM and DRAM consumption (GB) after loading 100 million records with 8-byte keys and 8-byte values.

	FPTree	LB ⁺ -Tree	ROART	DPTree	PACTree	μ Tree
DRAM	0.14	0.34	0.14	1.14	0.18	2.63
PM	2.69	2.54	18.92	4.79	3.2	3.2

advocates using snooping protocols for PM, which broadcasts coherence traffic across all cores, instead of using a directory to record cacheline status, hence does not incur additional PM accesses. However, most platforms default to a directory-based protocol because snooping may not scale to high core counts. Since our server does not allow changing coherence protocols, we were unable to verify the performance of PM indexes using snooping protocols; we leave it as future work. As we have noted in Section 4.2, requiring a certain coherence protocol may inflict issues with other applications and limit the applicability of the index.

6.8 PM and DRAM Space Consumption

New PM indexes are using DRAM more extensively to achieve high performance. The downsides are (1) more complex recovery protocols and (2) higher DRAM space consumption (hence higher cost of ownership). Table 3 lists the amount of DRAM and PM used by each index after loading 100 million records of 8-byte key and values (1.6GB). DPTree and μ Tree practically store complete trees in DRAM, resulting in up to $\sim 18\times$ higher DRAM consumption when compared with FPTree and LB⁺-Tree. PACTree uses a similar amount of PM to μ Tree's and is among the most frugal in using DRAM (second to FPTree) because of its packed design.

Surprisingly, ROART uses 18.92GB to index 1.6GB of data, while others need 2.5–4.8GB. The reason is it requires cacheline-aligned nodes and overprovisions leaf arrays: for 100 million records, it allocates space for around one billion records (17 million leaf arrays) occupying 10.89GB of PM as each leaf array is 640-byte. However, most space (reserved for leaf pointers) are unused. Such overprovisioning is due to ROART's split mechanism. In our experiment, keys are uniform randomly generated, so the 64 records in a full leaf array are usually very different despite they share a common prefix. Then a split operation could allocate 63 leaf arrays in the worst case, leaving each new leaf array with only few records, leading to an average occupancy of $\sim 9\%$ and a waste of over 9GB of PM.




Figure 9: Bandwidth consumption per operation under 20 threads with 100 million records of 8-byte keys and values.

6.9 Bandwidth Utilization and Requirements

Previous evaluation [44] has shown that PM bandwidth is scarce. Since a database system also uses various other components, it is desirable to keep the bandwidth consumption low for indexes. This has been the main focus of newer indexes. We observe that the peak usage across all indexes and operations does not reach the limit ($\sim 10\text{GB/s}/\sim 40\text{GB/s}$ for random write/sequential read). This shows the effectiveness of the bandwidth saving techniques proposed by the surveyed indexes. Due to space limitation, we omit the details on total bandwidth utilization and focus on the bandwidth used per operation, which is more indicative on how frugal (or not) an index uses PM bandwidth. Figure 9 shows the results with 20 threads when running lookups, inserts and scans under 8-byte keys and 8-byte values under the uniform distribution. Overall, LB⁺-Tree exhibits the lowest number of bytes per operation, thanks to its node layout and logless design. ROART exhibits the highest PM reads as it overprovisions node space. In contrast, PACTree, which is also trie-based, has a similar bandwidth requirement to LB⁺-Tree’s, because of its layout designed to reduce unnecessary PM accesses.

7 OBSERVATIONS AND INSIGHTS

In this section, we summarize the major findings based on our experimental results and analysis of the surveyed PM indexes.

1. The rule of thumb remains reducing PM accesses, which was first set in the pre-Optane era. Almost all the design choices (both pre-Optane and new ones) center around this goal, due to PM’s lower bandwidth and higher latency. If the properties of future PM hardware changes, the principles may be revisited.

2. Some building blocks from the pre-Optane era continue to work well and are further optimized by new PM indexes. Most indexes use DRAM to accelerate traversal; some (e.g., DP Tree) even place entire trees in DRAM. However, tries typically cannot take the full advantage of DRAM to store reconstructible data. Fingerprints are also widely used and enhanced by placing them in the spare bits of pointers and accessing them using SIMD instructions.

3. Using extra metadata and selective persistence of metadata can further accelerate performance. The main reason is these approaches can avoid using WAL, which may incur additional PM writes and complicate code logic.

4. Newly Proposed \neq Better. Pre-Optane FPTree is still very competitive and sometimes can even outperform newer indexes. Although μ Tree optimizes for tail latency, it exhibits the highest

latency in many cases. Such results call for careful benchmarking and comprehensive evaluations.

5. All the new PM indexes are tailor-made for one product (Intel Optane PMem), which can be a double-edged sword. While this can deliver high performance, as exemplified by LB⁺-Tree which performs the best in most cases, it could pose challenges when the PM hardware landscape becomes more diverse.

6. Support for NUMA-awareness, efficient PM management and variable-length keys remains inadequate. There have been initial attempts (e.g., using pointers for variable-length keys), but they are usually ad hoc or partial solutions with practical limitations (e.g., requiring a specific coherence protocol).

7. There is no clear “winner” index architecture, but the choice may affect how (efficiently) functionality can be supported. For example, B+-tree (trie) variants perform well for fixed-length (variable-length) keys. But it remains to be explored whether it is easier to add efficient variable-length key support in LB⁺-Tree or to optimize PACTree to match LB⁺-Tree’s performance.

8. Linked lists with small nodes are a bad fit for PM indexes, and cache misses in general should be minimized or hidden. Accessing and scanning through a linked list of individual records incur many cache misses which can dominate the performance and lead to high latency (e.g., in μ Tree), canceling out the positive effects brought by other optimizations.

9. HTM can perform well under NUMA for read-only workloads, but is challenging to handle contention and debug. In particular, the programmability issue is further complicated with other system-level infrastructure: we observed extremely high abort rates under `glibc` version 2.33 which does not use the right instructions that can work with HTM in `memcpy`.¹¹ The bug was fixed in `glibc` 2.34 [54] which is used in our experiments. It is noticeable that the best performing LB⁺-Tree is based on HTM; it therefore remains to be seen in future work whether other approaches can overcome these issues while maintaining high performance.

8 ON THE NEXT PM AND DRAM INDEXES

We give the outlook of the PM indexing space and describe promising future directions for PM and DRAM indexing.

8.1 Future PM Indexes

We identify three promising areas of future work for PM indexes.

1. Efficient Support for Full Functionality. As we have discussed previously, variable-length keys and NUMA-awareness remain open problems for future PM indexes. Importantly, it is desirable to maintain the high performance obtained by existing designs while better supporting full functionality.

2. Wider Applicability/Less Tailor-Made. There are various ways to realize PM, by using new materials (e.g., memristor [62], STT-RAM [24] and Intel 3D XPoint which PMem is based upon) or NVDIMMs which combine flash and DRAM [1, 65]. However, most (if not all) indexes aiming for real PM are tailor-made for Intel Optane PMem; yet certain properties like 256-byte alignment may even change across generations of the same product, and designs based on them may not work well on NVDIMMs, diminishing their applicability. Although some of the hardware efforts are in their

¹¹Details at https://sourceware.org/bugzilla/show_bug.cgi?id=28033.

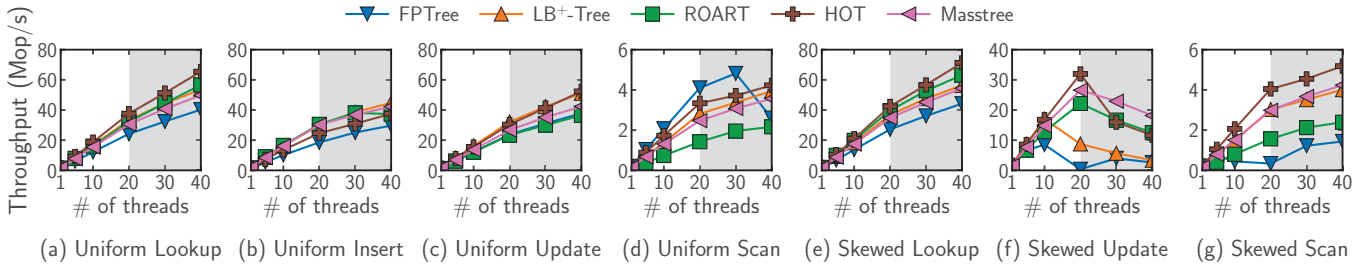


Figure 10: Throughput of PM and DRAM indexes running on DRAM. Without the extra fences, cacheline flushes and PM management code, PM-tailored indexes at least match the performance of HOT and Masstree. LB⁺-Tree performs even better than HOT and Masstree on certain insert workloads (b), and FPTree tops scan performance under uniform distribution (d).

early stage, we argue it is important to consider applicability of future designs on different PM devices.

3. Real-World Adoption and Cost-Effectiveness. Although there have been numerous PM index proposals, we are yet to see major adoption in real systems and applications. Part of the reason is the low cost effectiveness of PM-based servers as identified by other work [26], especially when compared to modern SSDs which can deliver high bandwidth and microsecond-level latency. Therefore, on the hardware side, we hope future work to lower the per GB cost of PM servers. On the software side, PM indexes and data structures in general should focus more on cost/performance.

8.2 Unifying PM and DRAM Indexing

In a similar vein to the point on wider applicability, we observe techniques proposed for PM indexes can also be effective for DRAM. We conduct preliminary experiments to compare the surveyed state-of-the-art PM indexes (with the extra cacheline flushes and fences removed) and two representative DRAM-optimized volatile indexes (HOT [6] and Masstree [50]). Figure 10 shows the throughput obtained by running the same workload as Section 6.4, but purely on DRAM. Under both uniform and skewed distributions, PM indexes perform competitively with DRAM-optimized indexes. In certain cases, PM indexes perform even better than HOT and Masstree which are specifically optimized for DRAM, e.g., LB⁺-Tree for inserts in Figure 10(b) and FPTree (despite being a pre-Optane proposal) in Figure 10(d) for scans. Although it remains to be seen how the optimizations for PM and DRAM indexes compare, and how PM techniques may be used by DRAM indexes (and vice versa), our results indicate it is promising to devise future indexes that would work on both volatile and non-volatile memory. This could greatly simplify implementation and widen the applicability of techniques proposed by both types of indexes.

9 RELATED WORK

Our work is most related to performance studies for PM devices and data structures, PM indexes and PM management issues.

Performance Studies. Early work [33, 72] characterized the performance of Optane PMem, exposing a set of properties different from what were previously assumed by emulations. Gugnani et. al [22] exposed more properties of Optane PMem under various scenarios, e.g., eADR and NUMA, along with case studies. Beyond range indexes, Hu et. al [25] evaluated PM-optimized hash indexes

on Optane PMem. Koutsoukos et. al [35] analyzed the performance of PM-enhanced database engines under TPC-C and TPC-H, and came up with guidelines of tuning the system for best performance.

PM Indexes. In addition to adapting specific indexes, general-purpose approaches such as RECIPE [41], NAP [66] and TIPS [37] present principled methods for converting DRAM indexes into PM indexes. It is interesting future work to evaluate these approaches. Some early efforts have adapted learned indexes [36] for PM. Chen et. al [10] observe that the bigger nodes used by learned indexes can cause excessive PM accesses. APEX [47] transforms the DRAM-based updatable ALEX [18] with concurrency and instant recovery on PM. Hash tables are also being re-designed for PM. CCEH [53] is a failure-atomic variant of extendible hashing that reduces directory management overhead. Dash [48] integrates a set of useful techniques to adapt extendible and linear hashing for PM; the key insight is that both PM reads and writes should be minimized. Clevel [14] is a lock-free version of level hashing [78] that performs asynchronous resizing in the background.

PM Libraries. PMDK [29] is the de facto standard, but may not be the optimal solution: ROART, DPTree and PACTree all devise their own approaches. Designing better PM libraries remains an open area, as seen by many recent alternatives [5, 7, 32, 60, 74].

10 CONCLUSION

We conducted a comprehensive evaluation of representative PM range indexes proposed based on real Intel Optane PMem. These new indexes inherited many useful designs from pre-Optane PM indexes and proposed new building blocks that can be useful for building future PM indexes. Based on our evaluation, we gave a list of observations, insights and future directions. We found the new indexes do not necessarily outperform the pre-Optane proposals, and efficient designs for variable-length keys, PM management and NUMA awareness are still lacking. We also discovered for the first time that a PM range index can match or even outperform DRAM-optimized indexes, highlighting the potential of unifying PM and DRAM indexing to save design and implementation efforts.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and associate editor for their constructive feedback. This work is partially supported by an NSERC Discovery Grant, Canada Foundation for Innovation John R. Evans Leaders Fund and the B.C. Knowledge Development Fund.

REFERENCES

- [1] AgigaTech. 2022. AGIGARAM NVDIMM-N. Retrieved May 15, 2022 from <http://agigatech.com/products/agigaram-nvdimm-n/>.
- [2] Joy Arulraj, Justin J. Levandoski, Umar Farooq Minhas, and Per-Åke Larson. 2018. BzTree: A High-Performance Latch-free Range Index for Non-Volatile Memory. *PVLDB* 11, 5 (2018), 553–565.
- [3] Joy Arulraj, Andrew Pavlo, and Subramanya R. Dulloor. 2015. Let’s Talk About Storage & Recovery Methods for Non-Volatile Memory Database Systems. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD ’15)*, 707–722.
- [4] Lawrence Benson, Hendrik Makait, and Tilmann Rabl. 2021. Viper: An Efficient Hybrid PMem-DRAM Key-Value Store. *Proc. VLDB Endow.* 14, 9 (may 2021), 1544–1556.
- [5] Kumud Bhandari, Dhruva R. Chakrabarti, and Hans-J. Boehm. 2016. Makalu: Fast Recoverable Allocation of Non-Volatile Memory. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, 677–694.
- [6] Robert Binna, Eva Zangerle, Martin Pichl, Günther Specht, and Viktor Leis. 2018. HOT: A Height Optimized Trie Index for Main-Memory Database Systems. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD ’18)*, 521–534.
- [7] Wentao Cai, Haosen Wen, H. Alan Beadle, Chris Kjellqvist, Mohammad Hedayati, and Michael L. Scott. 2020. Understanding and Optimizing Persistent Memory Allocation. In *Proceedings of the 2020 ACM SIGPLAN International Symposium on Memory Management (ISMM 2020)*, 60–73.
- [8] Zixian Cai, Stephen M. Blackburn, and Michael D. Bond. 2021. Understanding and Utilizing Hardware Transactional Memory Capacity. In *Proceedings of the 2021 ACM SIGPLAN International Symposium on Memory Management (ISMM 2021)*, 1–14.
- [9] Cheng Chen, Jun Yang, Mian Lu, Taize Wang, Zhao Zheng, Yuqiang Chen, Wenyuan Dai, Bingsheng He, Weng-Fai Wong, Guoan Wu, Yuping Zhao, and Andy Rudoff. 2021. Optimizing In-Memory Database Engine for AI-Powered on-Line Decision Augmentation Using Persistent Memory. *PVLDB* 14, 5 (2021), 799–812.
- [10] Leying Chen and Shimin Chen. 2021. How Does Updatable Learned Index Perform on Non-Volatile Main Memory?. In *2021 IEEE 37th International Conference on Data Engineering Workshops (ICDEW)*, 66–71.
- [11] Shimin Chen, Phillip B. Gibbons, and Suman Nath. 2011. Rethinking Database Algorithms for Phase Change Memory. In *5th Biennial Conference on Innovative Data Systems Research, CIDR 2011, Asilomar, CA, USA, January 9-12, 2011, Online Proceedings*.
- [12] Shimin Chen and Qin Jin. 2015. Persistent B+-Trees in Non-Volatile Main Memory. *PVLDB* 8, 7 (2015), 786–797.
- [13] Youmin Chen, Youyou Lu, Kedong Fang, Qing Wang, and Jiwu Shu. 2020. uTree: A Persistent B+-Tree with Low Tail Latency. *PVLDB* 13, 11 (2020), 2634–2648.
- [14] Zhangyu Chen, Yu Hua, Bo Ding, and Pengfei Zuo. 2020. Lock-free Concurrent Level Hashing for Persistent Memory. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, 799–812.
- [15] Jeremy Condit, Edmund N. Nightingale, Christopher Frost, Engin Ipek, Benjamin Lee, Doug Burger, and Derrick Coetzee. 2009. Better I/O through Byte-Addressable, Persistent Memory. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles (SOSP ’09)*, 133–146.
- [16] Intel Corporation. 2021. Intel Optane Persistent Memory (PMem). <https://www.intel.ca/content/www/ca/en/architecture-and-technology/optane-dc-persistent-memory.html>
- [17] Rob Crooke and Mark Durcan. 2015. A Revolutionary Breakthrough in Memory Technology. *3D XPoint Launch Keynote (2015)*.
- [18] Jialin Ding, Umar Farooq Minhas, Jia Yu, Chi Wang, Jaeyoung Do, Yinan Li, Hantian Zhang, Badrish Chandramouli, Johannes Gehrke, Donald Kossmann, David Lomet, and Tim Kraska. 2020. ALEX: An Updatable Adaptive Learned Index. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD ’20)*, 969–984.
- [19] Jason Evans. 2006. A Scalable Concurrent malloc (3) Implementation for FreeBSD. In *Proceedings of the BSDCan Conference*.
- [20] Keir Fraser. 2004. *Practical lock-freedom*. Technical Report UCAM-CL-TR-579. University of Cambridge, Computer Laboratory. <https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-579.pdf>
- [21] Jim Gray, Prakash Sundareshan, Susanne Englert, Ken Baclawski, and Peter J. Weinberger. 1994. Quickly Generating Billion-Record Synthetic Databases. In *Proceedings of the 1994 ACM SIGMOD International Conference on Management of Data (SIGMOD ’94)*, 243–252.
- [22] Shashank Gugrani, Arjun Kashyap, and Xiaoyi Lu. 2020. Understanding the Idiosyncrasies of Real Persistent Memory. *PVLDB* 14, 4 (2020), 626–639.
- [23] Timothy L. Harris. 2001. A Pragmatic Implementation of Non-Blocking Linked-Lists. In *Proceedings of the 15th International Conference on Distributed Computing (DISC ’01)*. Springer-Verlag, Berlin, Heidelberg, 300–314.
- [24] M. Hosomi, H. Yamagishi, T. Yamamoto, K. Bessho, Y. Higo, K. Yamane, H. Yamada, M. Shoji, H. Hachino, C. Fukumoto, H. Nagao, and H. Kano. 2005. A novel nonvolatile memory with spin torque transfer magnetization switching: spin-ram. *IEEE International Electron Devices Meeting (IEDM)* (2005), 459–462.
- [25] Daokun Hu, Zhiwen Chen, Jianbing Wu, Jianhua Sun, and Hao Chen. 2021. Persistent Memory Hash Indexes: An Experimental Evaluation. *PVLDB* 14 (2021), 785–798.
- [26] Kaisong Huang, Darien Imai, Tianzheng Wang, and Dong Xie. 2022. SSDs Striking Back: The Storage Jungle and Its Implications on Persistent Indexes. In *12th Annual Conference on Innovative Data Systems Research, CIDR 2022, Chaminade, CA, USA, January 9-12, 2022, Online Proceedings*.
- [27] Deukyeon Hwang, Wook-Hee Kim, Youjip Won, and Beomseok Nam. 2018. Endurable transient inconsistency in byte-addressable persistent B+-tree. In *16th USENIX Conference on File and Storage Technologies (FAST 18)*, 187–200.
- [28] Intel. 2021. Brief: Intel® Optane™ Persistent Memory – The Challenge of Keeping Up with Data. <https://www.intel.ca/content/www/ca/en/products/docs/memory-storage/optane-persistent-memory/optane-dc-persistent-memory-brief.html>
- [29] Intel. 2021. Persistent Memory Development Kit. (2021). <http://pmem.io/pmdk/>.
- [30] Intel Corporation. 2021. Intel 64 and IA-32 Architectures Software Developer’s Manual. (2021). <https://software.intel.com/content/www/us/en/develop/articles/intel-sdm.html>
- [31] Intel Corporation. 2021. Optane DCPMM 200 Series Product Brief. Retrieved August 17, 2021 from <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/optane-persistent-memory-200-series-brief.pdf>.
- [32] Keita Iwabuchi, Karim Youssef, Kaushik Velusamy, Maya Gokhale, and Roger Pearce. 2022. Metall: A persistent memory allocator for data-centric analytics. *Parallel Comput.* 111 (2022).
- [33] Joseph Izraelevitz, Jian Yang, Lu Zhang, Juno Kim, Xiao Liu, Amir Saman Memaripour, Yun Joon Soh, Zixuan Wang, Yi Xu, Subramanya R. Dulloor, Jishen Zhao, and Steven Swanson. 2019. Basic Performance Measurements of the Intel Optane DC Persistent Memory Module. arXiv:1903.05714 [cs.DC]
- [34] Wook-Hee Kim, R. Madhava Krishnan, Xinwei Fu, Sanidhya Kashyap, and Changwoo Min. 2021. PACTree: A High Performance Persistent Range Index Using PAC Guidelines. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles*, 424–439.
- [35] Dimitrios Koutsoukos, Raghav Bhartiya, Ana Klimovic, and Gustavo Alonso. 2021. How to use Persistent Memory in your Database. *CoRR* abs/2112.00425 (2021). arXiv:2112.00425 <https://arxiv.org/abs/2112.00425>
- [36] Tim Kraska, Alex Beutel, Ed H. Chi, Jeffrey Dean, and Neoklis Polyzotis. 2018. The Case for Learned Index Structures. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD ’18)*, 489–504.
- [37] R. Madhava Krishnan, Wook-Hee Kim, Xinwei Fu, Sumit Kumar Monga, Hee Won Lee, Minsung Jang, Ajit Mathew, and Changwoo Min. 2021. TIPS: Making Volatile Index Structures Persistent with DRAM-NVMM Tiering. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, 773–787.
- [38] Laboratory for Web Algorithms. 2022. UK Domain from 2005. <http://data.law.di.unimi.it/webdata/uk-2005/uk-2005.urls.gz>
- [39] Benjamin C. Lee, Engin Ipek, Onur Mutlu, and Doug Burger. 2009. Architecting Phase Change Memory as a Scalable Dram Alternative. In *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA ’09)*, 2–13.
- [40] Se Kwon Lee, K. Hyun Lim, Hyunsub Song, Beomseok Nam, and Sam H. Noh. 2017. WORT: Write Optimal Radix Tree for Persistent Memory Storage Systems. In *15th USENIX Conference on File and Storage Technologies (FAST 17)*, 257–270.
- [41] Se Kwon Lee, Jayashree Mohan, Sanidhya Kashyap, Taesoo Kim, and Vijay Chidambaram. 2019. RECIPE: Converting Concurrent DRAM Indexes to Persistent-Memory Indexes. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP ’19)*, 462–477.
- [42] Viktor Leis, Alfons Kemper, and Thomas Neumann. 2013. The Adaptive Radix Tree: ARTful Indexing for Main-Memory Databases. In *Proceedings of the 2013 IEEE International Conference on Data Engineering (ICDE ’13)*, 38–49.
- [43] Viktor Leis, Florian Scheibner, Alfons Kemper, and Thomas Neumann. 2016. The ART of Practical Synchronization. In *Proceedings of the 12th International Workshop on Data Management on New Hardware (DaMoN ’16)*, Article 3, 8 pages.
- [44] Lucas Lersch, Xiangpeng Hao, Ismail Oukid, Tianzheng Wang, and Thomas Willhalm. 2019. Evaluating Persistent Memory Range Indexes. *PVLDB* 13, 4 (2019), 574–587.
- [45] Gang Liu, Leying Chen, and Shimin Chen. 2021. Zen: A High-Throughput Log-Free OLTP Engine for Non-Volatile Main Memory. *PVLDB* 14, 5 (2021), 835–848.
- [46] Jihang Liu, Shimin Chen, and Lujun Wang. 2020. LB+Trees: Optimizing Persistent Index Performance on 3DXPoint Memory. *PVLDB* 13, 7 (2020), 1078–1090.
- [47] Baotong Lu, Jialin Ding, Eric Lo, Umar Farooq Minhas, and Tianzheng Wang. 2021. APEX: A High-Performance Learned Index on Persistent Memory. *PVLDB* 15, 3 (2021), 597–610.
- [48] Baotong Lu, Xiangpeng Hao, Tianzheng Wang, and Eric Lo. 2020. Dash: Scalable Hashing on Persistent Memory. *PVLDB* 13, 8 (2020), 1147–1161.
- [49] Shaonan Ma, Kang Chen, Shimin Chen, Mengxing Liu, Jianglang Zhu, Hongbo Kang, and Yongwei Wu. 2021. ROART: Range-query Optimized Persistent ART. In *19th USENIX Conference on File and Storage Technologies (FAST 21)*, 1–16.

- [50] Yandong Mao, Eddie Kohler, and Robert Tappan Morris. 2012. Cache craftiness for fast multicore key-value storage. In *Proceedings of the 7th ACM european conference on Computer Systems*. 183–196.
- [51] Chris Mellor. 2019. Is Optane DIMM endurance good enough? Quick answer... Yes, Intel has delivered. <https://blocksandfiles.com/2019/04/04/enduring-optane-dimm-question-is-its-endurance-good-enough-yes-intel-has-delivered/>
- [52] Colin Morris. 2017. Reddit Usernames. <https://www.kaggle.com/datasets/colinmorris/reddit-usernames>
- [53] Moohyeon Nam, Hokeun Cha, Young ri Choi, Sam H. Noh, and Beomseok Nam. 2019. Write-Optimized Dynamic Hashing for Persistent Memory. In *17th USENIX Conference on File and Storage Technologies (FAST 19)*. 31–44.
- [54] Carlos O’Donell. 2021. The GNU C Library version 2.34 is now available – [28033] libc: Need to check RTM_ALWAYS_ABORT for RTM. <https://sourceware.org/pipermail/libc-alpha/2021-August/129718.html>.
- [55] Ismail Oukid, Daniel Booss, Wolfgang Lehner, Peter Bumbulis, and Thomas Willhalm. 2014. SOFORT: A Hybrid SCM-DRAM Storage Engine for Fast Data Recovery. In *Proceedings of the Tenth International Workshop on Data Management on New Hardware (DaMoN ’14)*. Article 8, 7 pages.
- [56] Ismail Oukid, Johan Lasperas, Anisoara Nica, Thomas Willhalm, and Wolfgang Lehner. 2016. FPTree: A Hybrid SCM-DRAM Persistent and Concurrent B-Tree for Storage Class Memory. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD*. 371–386.
- [57] Jinsu Park and Woongki Baek. 2018. Quantifying the Performance and Energy-Efficiency Impact of Hardware Transactional Memory on Scientific Applications on Large-Scale NUMA Systems. In *2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 804–813.
- [58] Steven Pelley, Thomas F. Wensich, Brian T. Gold, and Bill Bridge. 2013. Storage Management in the NVRAM Era. *PVLDB* 7, 2 (2013), 121–132.
- [59] Raghu Ramakrishnan and Johannes Gehrke. 2003. *Database Management Systems* (3 ed.).
- [60] David Schwalb, Tim Berning, Martin Faust, Markus Dreseler, and Hasso Plattner. 2015. nvm malloc: Memory Allocation for NVRAM. In *ADMS@VLDB*. 61–72.
- [61] Benjamin Spector, Andreas Kipf, Kapil Vaidya, Chi Wang, Umar Farooq Minhas, and Tim Kraska. 2021. Bounding the Last Mile: Efficient Learned String Indexing (Extended Abstracts). In *3rd International Workshop on Applied AI for Database Systems and Applications, AIDB Workshops*.
- [62] D. B. Strukov, G. S. Snider, D. R. Stewart, and R. S. Williams. 2008. The missing memristor found. *Nature* 453, 7191 (2008), 80–83.
- [63] Alexander van Renen, Viktor Leis, Alfons Kemper, Thomas Neumann, Takushi Hashida, Kazuichi Oe, Yoshiyasu Doi, Lilian Harada, and Mitsuru Sato. 2018. Managing Non-Volatile Memory in Database Systems. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD ’18)*. 1541–1555.
- [64] Shivaram Venkataraman, Niraj Tolia, Parthasarathy Ranganathan, and Roy H. Campbell. 2011. Consistent and Durable Data Structures for Non-Volatile Byte-Addressable Memory. In *Proceedings of the 9th USENIX Conference on File and Storage Technologies (FAST’11)*.
- [65] Viking Technology. 2017. DDR4 NVDIMM. Retrieved August 17, 2021 from <http://www.vikingtechnology.com>.
- [66] Qing Wang, Youyou Lu, Junru Li, and Jiwu Shu. 2021. Nap: A Black-Box Approach to NUMA-Aware Persistent Memory Indexes. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21)*. 93–111.
- [67] Tianzheng Wang and Ryan Johnson. 2014. Scalable Logging through Emerging Non-Volatile Memory. *PVLDB* 7, 10 (2014), 865–876.
- [68] Tianzheng Wang, Justin Levandoski, and Per-Åke Larson. 2018. Easy Lock-Free Indexing in Non-Volatile Memory. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. 461–472.
- [69] Wikimedia Dump Service. 2022. Wikipedia Dump 20220420. <https://dumps.wikimedia.org/enwiki/20220420/enwiki-20220420-all-titles.gz>
- [70] H. S P Wong, S. Raoux, SangBum Kim, Jiale Liang, John P. Reifenberg, B. Rajendran, Mehdi Asheghi, and Kenneth E. Goodson. 2010. Phase Change Memory. *Proc. IEEE* 98, 12 (2010), 2201–2227.
- [71] Fei Xia, Dejun Jiang, Jin Xiong, and Ninghui Sun. 2017. HiKV: A Hybrid Index Key-value Store for DRAM-NVM Memory Systems. In *Proceedings of the 2017 USENIX Annual Technical Conference (USENIX ATC ’17)*. 349–362.
- [72] Jian Yang, Juno Kim, Morteza Hoseinzadeh, Joseph Izraelevitz, and Steven Swanson. 2020. An Empirical Guide to the Behavior and Use of Scalable Persistent Memory. In *Proceedings of the 18th USENIX Conference on File and Storage Technologies (FAST’20)*. 169–182.
- [73] Jun Yang, Qingsong Wei, Cheng Chen, Chungong Wang, Khai Leong Yong, and Bingsheng He. 2015. NV-Tree: reducing consistency cost for NVM-based single level systems. In *13th USENIX Conference on File and Storage Technologies (FAST 15)*. 167–181.
- [74] Lu Zhang and Steven Swanson. 2019. Pangolin: A Fault-Tolerant Persistent Memory Programming Library. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*. 897–912.
- [75] Ping Zhou, Bo Zhao, Jun Yang, and Youtao Zhang. 2009. A Durable and Energy Efficient Main Memory Using Phase Change Memory Technology. In *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA ’09)*. 14–23.
- [76] Xinjing Zhou, Joy Arulraj, Andrew Pavlo, and David Cohen. 2021. Spitfire: A Three-Tier Buffer Manager for Volatile and Non-Volatile Memory. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD/PODS ’21)*. 2195–2207.
- [77] Xinjing Zhou, Lidan Shou, Ke Chen, Wei Hu, and Gang Chen. 2019. DPTree: Differential Indexing for Persistent Memory. *PVLDB* 13, 4 (2019), 421–434.
- [78] Pengfei Zuo, Yu Hua, and Jie Wu. 2018. Write-Optimized and High-Performance Hashing Index Scheme for Persistent Memory. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. 461–476.