

DLHT: A Non-blocking Resizable Hashtable with Fast Deletes and Memory-awareness

Antonios Katsarakis*
Huawei Research

Vasilis Gavrielatos*
Huawei Research

Nikos Ntarmos
Huawei Research

ABSTRACT

This paper presents DLHT, a concurrent in-memory hashtable. Despite efforts to optimize hashtables, that go as far as sacrificing core functionality, state-of-the-art designs still incur multiple memory accesses per request and block request processing in three cases. First, most hashtables block while waiting for data to be retrieved from memory. Second, open-addressing designs, which represent the current state-of-the-art, either cannot free index slots on deletes or must block all requests to do so. Third, index resizes block every request until all objects are copied to the new index. Defying folklore wisdom, DLHT forgoes open-addressing and adopts a fully-featured and memory-aware closed-addressing design based on bounded cache-line-chaining. This design offers ① lock-free operations and deletes that free slots instantly, ② completes most requests with a single memory access, ③ utilizes software prefetching to hide memory latencies, and ④ employs a novel non-blocking and parallel resizing. In a commodity server and a memory-resident workload, DLHT surpasses 1.6B requests per second and provides 3.5× (12×) the throughput of the state-of-the-art closed-addressing (open-addressing) resizable hashtable on Gets (Deletes).

CCS CONCEPTS

• **Computing methodologies** → **Concurrent algorithms**; • **Information systems** → **Data structures**; **Hashed file organization**.

KEYWORDS

high throughput; in-memory; non-blocking; hashtable; prefetching; consistency; memory-aware; deletes; resize; close-addressing

ACM Reference Format:

Antonios Katsarakis, Vasilis Gavrielatos, Nikos Ntarmos. 2024. DLHT: A Non-blocking Resizable Hashtable with Fast Deletes and Memory-awareness. In *The 33rd International Symposium on High-Performance Parallel and Distributed Computing (HPDC '24)*, June 3–7, 2024, Pisa, Italy. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3625549.3658682>

1 INTRODUCTION

Concurrent in-memory hashtables are essential and versatile data structures in the modern cloud. They are responsible for storing and accessing large amounts of data in main memory via thread-safe *Get*,

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 the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, and to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HPDC '24, June 2024, Pisa, Italy

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0413-0/24/06...\$15.00

<https://doi.org/10.1145/3625549.3658682>




Figure 1: Throughput of state-of-the-art hashtables and DLHT with 64 threads in a memory-resident workload (100M objects).

Put, *Insert*, and *Delete* requests. To ensure requests complete rapidly as the dataset expands, hashtables must be also able to efficiently *Resize* their index. In-memory hashtables serve a wide spectrum of applications, including in-memory storage, online services, caching, key-value stores, and transactional databases [4, 12, 13, 17, 41].

To meet the ever-growing performance demands [4, 14], state-of-the-art hashtables from industry and academia offer designs that attain close to a billion requests per second on a single server [10, 36, 38, 45, 47, 57, 60]. Problematically, their evaluation hints that such high throughput is reachable only under *cache-resident* workloads where accesses are served by hardware caches and seldom reach main memory – i.e., due to small datasets [10], data partitioning [45, 57, 60], or highly skewed accesses [36, 38, 47]. So we pose the following question: *Can state-of-the-art in-memory hashtables attain a billion requests per second under a memory-resident workload?*

To answer this question, we evaluate eight state-of-the-art designs over a *memory-resident* workload of 100 million objects accessed uniformly on a commodity server (details in § 4). As shown in Figure 1, almost all hashtables are more than 2× slower than a billion requests per second. The most recent work, DRAMHiT [51], is the only one close to the target (in Gets), but its open-addressing design hinders Deletes and Resizes. Hence, achieving a billion requests per second without forfeiting core functionality on a commodity server remains a challenge for memory-resident workloads.

In a deeper inspection (detailed in Table 1 and § 2), state-of-the-art designs offer lock-free accesses but sacrifice core functionality, incur multiple memory accesses per request, and block processing in three cases. First, most hashtables stall processing on every request when accessing memory. Second, open-addressing designs offer impaired Deletes that either cannot reclaim index slots or must cease processing and rebuild the entire index to do so. Third, those that support index Resizes block every request until all objects are copied to the new index. These stalling factors impede the throughput of state-of-the-art hashtables, rendering them *practi-*

cally blocking under memory-resident workloads.

In this work, we introduce DLHT, a concurrent hashtable that is *memory access aware* and *practically non-blocking* (i.e., alleviates stalling) to transcend a billion requests per second in memory-resident workloads. Defying folklore wisdom [47, 51, 61], DLHT forgoes open-addressing and adopts a closed-addressing approach. Its design is based on bounded cache-line-chaining and has the following features. First, it enables lock-free index operations, including deletes with immediate index slot reclamation. Second, it minimizes memory traffic and completes most requests with a single memory access. Third, it exploits software prefetching to overlap the memory latency of a request with productive work on other requests. Finally, it incorporates a novel, non-blocking (but not lock-free) Resize where requests complete with strong consistency while a multi-threaded index migration occurs in parallel.

Unlike state-of-the-art designs that trade core functionality for throughput [10, 15, 38, 47], DLHT provides a complete set of implemented features to accommodate its clients' needs. Beyond core functionality, this includes namespaces, variable-sized key-value pairs, efficient single-thread and hashset variants, as well as pointer APIs that minimize copies.

We extensively evaluate DLHT on a commodity server using micro-benchmarks, sensitivity studies, application examples, and standard single- and multi-key OLTP benchmarks (YCSB, TATP, and Smallbank). We compare DLHT with eight state-of-the-art concurrent in-memory designs. DLHT surpasses 1.6B Get (1.4B Inserts/Deletes, 1B Gets/Puts) requests per second. This is more than 3.5× (3×, 2.7×) the performance of the fastest closed-addressing design and an order of magnitude faster Deletes than open-addressing designs. Finally, the parallel and non-blocking resize of DLHT allows for a population that is 3.9× faster than the state-of-the-art. In short, the contributions of this work are as follows:

- **We spot core performance and functionality features for memory-resident hashtables** and detail how state-of-the-art concurrent hashtables fall short on those. (§ 2)
- **We introduce DLHT, a closed-addressing hashtable** that is fully-featured and maximizes throughput via a lock-free bounded cacheline-chaining design. DLHT minimizes memory traffic per request, leverages software prefetching to mask memory accesses, ensures good occupancy, offers fast Deletes that free index slots instantly, and is equipped with a parallel and non-blocking index resizing. (§ 3)
- **We extensively evaluate the performance of DLHT** on memory-resident workloads over a commodity server. DLHT provides 1.66B requests/second and 3.5× (12×) the throughput of the state-of-the-art resizable closed-addressing (open-addressing) designs on Gets (Deletes). (§ 5)

2 MEMORY-RESIDENT HASHTABLES

In Table 1, we summarize two key performance features for hashtables targeting memory-resident workloads: ① *memory access awareness* and ② *practically non-blocking* operations. We break down those features next and subsequently follow on why state-of-the-art falls short on them.

2.1 Key features for performance

Memory access awareness. This includes three guidelines. First, the design must ① *overlap accessing memory with useful work*. This can be achieved by software prefetching and then switching to doing useful work, while memory is being accessed [21]. We observe that once this is achieved, the bottleneck tends to shift towards memory bandwidth. Therefore the second guideline is to ② *minimize the memory traffic per request*. To do so, the design must strive to approach the ideal limit of accessing a single cache-line per request, while avoiding extraneous write-backs to memory (e.g., Gets should not write). This calls for designs where small values are inlined inside the index and Gets are lock-free.

Reduced memory traffic can be achieved by prematurely resizing the index. To ensure this is not the case, the third rule is that, given a state-of-the-art hash function, the design must ③ *maintain high occupancy* – i.e., a high percentage of occupied to total slots before a resize must be performed.

Practically non-blocking. A Get, Put, Insert or Delete to key K_A is deemed to ④ *be practically non-blocking* when its algorithm guarantees that it does not impede the progress of any concurrent Get, Put, Insert or Delete to a different key K_B . The Resize operation is deemed *non-blocking* when it does not block all other operations until all objects are copied to a new index. The latter is crucial for memory-resident workloads with hashtables that can span gigabytes of data.

2.2 Shortcomings of state-of-the-art hashtables

From Figure 1, we focus on the five fastest concurrent in-memory hashtables: MICA [38, 40], CLHT (lock-free) [10], GrowT [47], Meta's Folly [15] and DRAMHiT [51]. In Table 1, we show how each design falls short on satisfying the performance features, by handling memory accesses inefficiently, blocking excessively, and even sacrificing core functionality.

Sacrificed functionality. As shown in Table 1, four out of five fastest hashtables sacrifice Puts, Deletes, or Resizes. Moreover, only MICA supports keys or values beyond 8 bytes (not shown Table 1). CLHT assumes that values across different keys in the index are unique. Notably, GrowT, Folly, and DRAMHiT are *open-addressing* hashtables [25]. It is well documented that there is no cheap way to perform Deletes in open-addressing [2, 20, 25, 31, 32, 53]. All three (GrowT, DRAMHiT, and Folly) perform Deletes through tombstones, which permanently occupy space in the index. These are very fast, but will fill the index after a while. DRAMHiT and Folly do not address that. GrowT moves all live elements to a new index, incurring a large performance penalty (§ 5).

Memory access awareness. We observe that only MICA and DRAMHiT overlap memory accesses with useful work. Detrimentally, MICA does not inline values in the index, thus mandating multiple memory accesses for a request, even for small values and no collisions. DRAMHiT overlaps the memory accesses of a batch of requests provided by a client. However, it performs the requests of the batch out-of-order. Problematically, this can lead to errors in some use cases (e.g., deadlocks in § 5.3.3). Finally, CLHT has low occupancy, as it cannot chain buckets and must Resize after any 4 collisions.

	Collision handling	Practically non-blocking operations					Memory access awareness		
		Gets	Puts	Inserts	Deletes (that free index slots)	Resizes (non-blocking: safe Get,...,Del during resize)	Overlaps memory accesses (i.e., exploits s/w prefetch)	Minimizes memory traffic (e.g., supports index inlining)	Occupancy until resize (from § 5.1.5 with wyhash)
GrowT	open-addressing	✓	✓	✓	Blocking (until all index is transferred)	Parallel, Blocking	✗	✓	30-50%
Folly	open-addressing	✓	✓	✓	✗	✗	✗	✓	—
DRAMHiT	open-addressing	✓	Only upserts	✗	✗	✗	✓ Reorders requests	✓	—
MICA	closed-addressing	✓	Blocking	✗	Blocking	✗	✓	✗	—
CLHT	closed-addressing	Unique values	✗	✓	✓	Serial, Blocking	✗	✓	1-5%
DLHT	closed-addressing	✓	✓	✓	✓	✓	✓	✓	61-72%

Table 1: Key features for memory-resident performance across state-of-the-art concurrent in-memory hashables and DLHT.

Practically non-blocking. MICA adopts a blocking lock-based scheme for updates. In contrast, CLHT, GrowT, and Folly support lock-free index modifications. DRAMHiT offers lock-free Gets and Upserts, but an application cannot express a pure Put or Insert. DRAMHiT, GrowT, and Folly struggle with non-blocking Deletes, as previously discussed. Finally, only CLHT and GrowT can grow on-demand but come with a blocking resize, during which all index operations are stalled.

3 DANDELION HASHTABLE (DLHT)

This section describes the design of DLHT, which addresses these shortcomings. First, we present the architecture of DLHT (§ 3.1) and its core algorithms (§ 3.2). Then we detail how DLHT overlaps memory accesses (§ 3.3) and discuss its extra features, including the iterator and namespaces (§ 3.4).

3.1 Architecture and terminology

In this section, we describe the architecture of DLHT (shown in Figure 2) and the basic terms we use throughout the paper.

Index. The index is an array of bins.

Bin. A bin is a chain of buckets. Each chain can have up to four buckets. Initially, each bin is comprised of a single (primary) bucket. There is an additional array of link buckets, which can be chained to a bin when needed. There is only a small number of link buckets (by default 8 times fewer than the bins), ensuring that at any time, most bins only have a single bucket – the primary bucket. When we run out of link buckets, we trigger a non-blocking resize (§ 3.2.5).

Bucket (64 B). A primary bucket has an 8-byte bin header followed by 8-byte link metadata and three 16-byte slots for key-value pairs. Link buckets just have four 16-byte key-value slots. Buckets are always cache-line aligned.

Bin Header (8 B). The first 8-bytes of a primary bucket is the bin header, which stores concurrency metadata. It contains a 32-bit version, a 2-bit bin state, and a 2-bit state for each of the 15 slots – i.e., for the primary bucket and up to three link buckets per bin. Fitting that metadata in 8-bytes enables non-blocking operations (e.g., Inserts or Deletes) in the bin via compare-and-swaps (CASes)

on the header. In later sections, we detail how these pieces are used.

Link Meta (8 B). The second 8-bytes of a primary bucket contains two 32-bit link-bucket indexes, which are used for chaining up to 3 buckets from the link bucket array to the bin. The first index chains one bucket while the second index chains two (consecutive in link array) buckets to the bin.

Slot (16 B). A key-value pair *slot* comprises two consecutive 8-byte segments, one for the key and one for the value.

DLHT has three modes of operation. Each mode uses the slot differently and offers a slightly different API. The three configurations are a result of DLHT’s real-world use cases.

1. Inlined. In this mode, key-value pairs are stored in the slot, i.e., they are *inlined*. Keys and values must be at most 8 bytes, so that they can both fit in their respective positions inside the slot. The Inlined configuration is extensively used by DLHT clients. Two examples of these use-cases are: ① a pointer cache in a query processing engine, ② a map between 4- to 8-byte pointers for a database storage engine.

2. Allocator. In this mode, we allocate additional memory to store a key-value pair. This mode is used when either the key or the value is larger than 8 bytes. The value is always written in the additional memory, and the slot contains a pointer to that memory, instead of the value. If the key is not larger than 8 bytes, it fits in the slot and is not stored externally. When the key is larger than 8 bytes, the slot stores its 8 least significant bytes (for filtering). In this mode, DLHT takes an allocator as input, as in C++ containers. This mode is used as a primary index by a database storage-engine.

3. HashSet. DLHT can be configured as a *HashSet*, which does not store values, only keys. The restriction is that keys must be at most 8 bytes to fit in the slot. The HashSet mode is currently used for semi-/anti-joins and by a database lock manager to lock records, where an insertion of a key, locks one or more records, which are released by deleting the key.

Notably, DLHT is suitable for more use-cases than those mentioned above. We further discuss some in § 3 and § 5.




Figure 2: DLHT layout: bins composed of a primary bucket (a bin header and 3 slots) and up to 3 extra linked buckets (4 slots each).

3.2 Algorithms

In this section, we detail the core algorithms of DLHT, which are practically non-blocking, minimize memory accesses, and ensure the strongest consistency (i.e., linearizability [26]). We focus our discussion on the Inlined mode, and at the end of each algorithm, we note the changes for the other modes.

3.2.1 Get: lock-free + ~ one memory access + ptr API

A Get first hashes the key to retrieve a bin. With closed-addressing, there is a one-to-one relation between a key and a bin: if the key is not in that bin, then the key does not exist in the index. The Get does a linear search over all filled slots of the bin. In the worst case, it searches all 15 slots over four buckets. Most commonly, the key hashes in a single-bucket bin; thus, the entire operation only costs a single memory access. When the key is hashed in a bin with more buckets and is not found in the first bucket, a second memory access is required to fetch the second bucket, and so on.

We use the version of the bucket to read a consistent view of the slot. Specifically, we first read the version of the header, then the 2-bit slot state, the slot, and then again the version. This is a common technique to implement lock-free reads, used in seqlocks [33] and other non-blocking algorithms [58].

Allocator. In Allocator mode, the slot stores a pointer instead of the value. Additionally, in Allocator mode, keys can be larger than 8 bytes; if so, they are stored inside the allocated memory. In this case, we must dereference the pointer to read the key as part of the concurrency algorithm.

Allocator pointer API. In Allocator mode, we offer a *pointer API*, where a Get returns a pointer to the value and there is no Put. Rather, the Get is used to modify the value. The alternative would have been a *Get/Put API*, where a Get returns the value and Put overwrites it. We choose the pointer API because of three scenarios: ① when the new value depends on the old value ② when only a small part of the value needs to be updated and ③ when the concurrency scheme can be optimized using context-specific information. Puts that blindly lock and overwrite the entire value would incur unnecessary copies or index probes in those cases.

HashSet. Gets return only if a key exists, as it has no values.

Bounded chaining. Recall from the previous section that the default number of available link buckets is 8 times smaller than the bins. This ensures that the average number of memory accesses per Get is bounded and remains close to one.

Summary. DLHT ensures that most Gets require only one memory access and no write-backs. We implement this with a well-established lock-free algorithm. In Allocator mode, Gets return

a pointer to the value. This can be used to efficiently implement custom modification algorithms.

3.2.2 Insert: à la CLHT + bounded chains + shadow API

A bin can have up to 15 slots across four chained buckets. The header of the bin stores a 2-bit state for each slot. This 2-bit state can take the values of Valid, Invalid, and TryInsert. An Insert first checks that the key is not already present, then picks an empty (i.e., Invalid) slot to insert, fills the slot, and changes the state of the slot to Valid.

From a concurrency perspective, the challenge is to ensure that there can never be two successful Inserts for the same key given that two Inserts may work on two different slots. We solve this using the CLHT lock-free algorithm. The key idea is to use the bin header as the synchronization point. This is possible, despite a bin in DLHT having up to four buckets, as all slot states are packed together in the same 8-byte word. Specifically, Inserts CAS the bin header to change the state of a single slot. CASing the header ensures the atomicity of the Insert with respect to other operations (Inserts/Deletes) on the same bin. Notably, because the slot and its state are on separate words, we must perform the Insert in two steps, using an intermediate state called *TryInsert*. The Insert algorithm is as follows.

Algorithm. ① Read the bin header. ② Run the Get algorithm, if the key already exists return its value along with the corresponding flag. ③ Otherwise, find the first slot that is in Invalid state. If no slot can be found, trigger a resize. ④ CAS the header to transition the state of the slot from Invalid to TryInsert. ⑤ If successful, fill the slot. ⑥ If unsuccessful, start over. ⑦ CAS the header again to transition the state of its selected slot from TryInsert to Valid. If unsuccessful we start over from step 1, but we skip steps 3 and 4.

This algorithm ensures that there can never be two successful Inserts for the same key, as only one can succeed in step 5; the other start over and search the bin for the key (step 2), where it will find the successful Insert. Also, recall that the header has a 32-bit version. The version is incremented with every CAS. This is required by the Get algorithm, but also protects from the notorious ABA bug [58].

Chaining buckets. After transitioning a slot to TryInsert state (step 4), it is possible that this slot is on a bucket that has not been chained yet. In this case, we first atomically allocate a link bucket (or two consecutive link buckets) through a Fetch-And-Add and then chain it (them) to the bin with a CAS on the link header – setting the appropriate link-bucket-index. If no link buckets are left, we trigger a resize.

Allocator. When in this mode, the Insert algorithm allocates memory in step 4.1 (after transitioning the slot state to TryInsert). Note that it is possible for the Insert to fail after the allocation, either because another thread has just inserted the same key, or because the bin ran out of slots. In both cases, we free the allocated memory, before continuing.

HashSet. The algorithm skips all actions related to values.

Transactions. Often transactional protocols have two rounds, ① lock some objects and ② commit (or abort) [23]. With the pointer API, locking can be implemented by embedding a lock inside each object value. However, this is not possible for Inserts, as the values do not yet exist in DLHT. To solve this, the client could implement locking through an additional HashSet. To avoid this overhead, we offer two additional API calls: The first is a *shadow Insert* that inserts the key but keeps it hidden from Get/Put/Deletes, essentially locking it. Internally, the only difference of the shadow Insert is that it transitions the inserted key to a *Shadow* state instead of the *Valid* state (in step 5). Later the transaction commits or aborts the shadow Insert using an API call with a corresponding argument, which transitions the slot state to *Valid* or *Invalid*.

Summary. DLHT uses a variant of CLHT's lock-free Inserts. In the common case, an Insert requires a single memory access, and two CASes on the bin header. The Insert is also responsible for chaining buckets and allocating memory. Finally, we offer optimized Inserts for transactional protocols.

3.2.3 Delete: lock-free + immediate slot reclaim

Unlike open-addressing designs, DLHT instantly reclaims slots on Deletes. To delete, we search the bin to locate the matching slot and CAS the header to change the state of the slot from *Valid* to *Invalid*. Hence, Inserts can reuse the slot.

Allocator. In the Allocator mode, DLHT stores pointers in the slots instead of values. Slots are again instantly reused. To free the pointer after a Delete, we offer an epoch-based GC, for which the client can opt-in. Our GC remembers the pointers that must be freed. The client periodically performs a call from all threads to advance the epoch. After moving to new epoch, our GC frees the pointers of the previous epoch.

HashSet. The Delete is the same as in the Inlined mode.

3.2.4 Put: dw-CAS + transfer keys

We offer a Put only for the Inlined mode. We implement Puts using a double-word CAS (dw-CAS) on the slot. Specifically, the client provides a key and the new value to be written. We first execute the Get algorithm to find the corresponding slot. We then use the value that we just read to dw-CAS the entire slot, with the value provided by the client.

It is possible that between reading and dw-CASing, the slot gets deleted and then reused by a subsequent Insert of another key. In this case, the dw-CAS ensures that the Put will either overwrite the entire slot before the Insert, or fail.

Besides concurrent Inserts and Deletes, we must also handle non-blocking Puts racing with an index Resize. It is important that no Put succeeds in the old index after its slot has been copied to the new index. Otherwise, it would be hard to guarantee linearizability.

A straightforward solution to this problem, is reducing inlined keys (or values) to 63 bits and requiring the resize algorithm to set this one bit before a transfer occurs. To avoid sacrificing this bit in every slot, we use the idea of *transfer keys* that we discuss in § 3.2.5.

Allocator & HashSet. In Allocator there are no Puts. Instead Gets return pointers for updates (discussed in § 3.2.1). In HashSet mode, there are no values and hence no Puts.

3.2.5 Resize: parallel growing + concurrent operations

An index resize is triggered by an Insert when either all slots of a bin are used or there are no more link buckets to chain to the bin. The *resizer*, i.e., the thread triggering the resize, executes the entire resize algorithm and then it performs its Insert in the new index. Resizing has three steps: ① allocate a new index, ② transfer all keys to the new index, and ③ GC the old index. Crucially, resizing does not block other threads from executing their requests. If another Insert also triggers a resize, then it joins the effort of resizing.

Allocating new index. The growth factor varies from 8 when the index is small (e.g., < 4K bins) to 4 for medium-sized indexes (e.g., < 64M bins) to 2 for larger sizes. The resizer is responsible to allocate the new index.

Transfer. The index is partitioned into chunks of 16K bins. This allows threads to collaborate with minimum synchronization. Each collaborating thread (resizer or not) picks a not-yet-transferred chunk and transfers it. We detail this collaboration later in this section. To transfer a chunk, we iterate through all bins, and for every valid slot in a bin we: ① read the slot; ② replace the key in the old index with a transfer key; and ③ insert the key into the new index.

Practically non-blocking operations. Instead of blocking all operations for the duration of the resize, we only block operations at a bin granularity for the duration of a bin transfer. Recall that each bin has a 2-bit bin state in its header. The bin state can be in *NoTransfer* (initial state), *InTransfer*, or *DoneTransfer*. Before transferring the bin, we CAS the header, setting the bin state to *InTransfer*. Once the transfer of a bin is done, we set the bin state to *DoneTransfer*.

All index operations first check the bin state. If it is *InTransfer* they wait until it reaches *DoneTransfer*. If it is *DoneTransfer*, they perform the operation in the new index. Note that CASing the bin state of the header ensures that all concurrent Inserts and Deletes will either take place before the transfer begins, or will fail and be retried in the new index. A Get also reads the header to complete a slot read and is similarly retried if the bin state is other than *NoTransfer*.

Puts read the header in the beginning but do not re-read or CAS the header afterwards, only dw-CAS the slot. For this reason, the resize algorithm also replaces the key of the slot with a transfer key, before transferring the slot. A transfer key is a key that can never be hashed in the targeted bin. Our implementation uses one key for odd and another for even bins. This ensures that either the transfer will see the Put, or the Put will fail, encountering the transfer key. In the latter case, the Put is retried and will find the bin in *InTransfer* or *DoneTransfer*. It will wait for the transfer to complete if needed and then it will be performed on the new index.

Collaboration. When a thread tries to trigger a resize while a resize is in progress, it becomes a *helper*. Helpers first wait for the new index to be allocated. Then they start transferring any available chunks of bins until there are no more chunks to transfer. Finally, they execute their Insert in the new index.

GC old index. Once the transfer is done, the resizer updates the index pointer, ensuring that all new requests use the new index. Now, the resizer needs to wait for all threads that have read the old index pointer and are still executing a request, to finish their request. To do this, we mandate that threads notify each other when finishing a request. We implement this with a per-thread pointer. When a thread enters DLHT (e.g., on a Get), we set the pointer to the current index. Just before the thread leaves DLHT, it sets the pointer to null. Once none of the pointers point to the old index, the resizer can GC it. The overhead is two atomic store instructions per request (per batch in practice – § 3.3). Crucially, no programmer intervention is required.

Summary. Resizes only block operations on a single bin at a time. Bin transfers are parallel with minimal synchronization. We GC the old index, without programmer intervention.

3.3 Overlapping memory latencies with useful work

A side-effect of hashing is that accesses to the index do not follow any specific pattern, nor do they exhibit any locality. This renders hardware techniques ineffective: the processor will block on every request waiting for main memory [21]. Software prefetching can be used to tolerate these latencies, either by batching requests or by using coroutines.

Batching. DLHT offers a function that executes an array of requests (i.e., a batch) while respecting the requested order. Before executing the requests, DLHT loops through the array and issues software prefetches for the bin of each request. This overlaps the memory latencies of all requests in the batch. The technique is inspired by MICA [41]. Unlike MICA, our pointer-based API also allows us to prefetch the externally stored values in Allocator mode.

In some software stacks, batching occurs naturally. For instance, one of our clients implements distributed transactions over DLHT, where each network packet contains multiple requests for different keys. Instead of looping through the network packet issuing requests one-by-one, the client packages them in a batch, reaping the benefits of DLHT batching. Note that it may be crucial to respect the order of requests in this and similar scenarios, as locks might need to be grabbed in order to avoid deadlocks (e.g., as in § 5.3.3).

DLHT accepts different request types (i.e., Get, Put, Insert, Delete) on the same batch, respects their order, and amortizes the cost of the index GC (two atomic stores) across the batch. Moreover, DLHT offers the option to terminate the execution of a batch on an operation that does not complete successfully. For example, if an application-level lock cannot be grabbed because a key does not exist, DLHT will not perform any subsequent requests of the batch.

Coroutines. Some software stacks make the use of batching challenging. In this case, coroutines can be used alternatively to mask memory latencies [24]. To interoperate with coroutines, we offer a function that prefetches the bin of a key. The client calls

this function prior issuing a request to DLHT, placing a yield in between. Consequently, the client can do useful work, while the bin is fetched from memory.

3.4 Additional features

DLHT has been in use for several months. During this time, we have accommodated numerous new features. We detailed some in the previous section, e.g., DLHT’s three modes of operation (§ 3.1). We briefly discuss others next. When a feature incurs a performance penalty, the client must explicitly opt-in, such that clients only pay for the features they need.

3.4.1 Variable size keys and values (in a single index)

A client can configure DLHT to hold key-values of any size. This allows having one DLHT instance for different object types and facilitates batching. For instance, the client may insert a 2-byte key with a 5-byte value and a 128-byte key with a 1024-byte value. This feature is available in Allocator mode – the only mode where keys can be larger than 8B. When this feature is enabled, we store the key and value sizes in the allocated memory of each key-value pair.

This poses a performance challenge. For inlined keys (i.e., at most 8 bytes), dereferencing the pointer would double the cost of a Get. Recall that each slot, has 8 bytes for the key and 8 bytes for the value. But in the Allocator mode, we store pointers instead of a value. As pointers only use 48 bits, we safely overload the 16 most significant bits (MSBs). We use the four MSBs for the key size (the rest are used in § 3.4.2). Four bits suffice, as keys larger than 8 bytes anyway need to dereference the pointer, which stores the key size.

3.4.2 Namespaces

When using a single DLHT instance to accommodate different types of keys (e.g., keys from different database tables), it is possible that keys have name conflicts. To address this, we introduced *namespaces*. Specifically, clients can tag a key with a namespace-id, which is an integer that can range from 0 to 4Ki. Keys with different namespace-ids do not conflict inside DLHT. We allow for 4Ki different namespaces, as we overload the 12 remaining most significant bits (MSBs) from the pointer inside the slot (as discussed in § 3.4.1).

3.4.3 Hash functions

The default hashing algorithm is a simple modulo operation:

$$bin_id = key \% number_of_bins$$

DLHT can also be configured to use wyhash[64]. We benchmarked many of the most prominent hash functions, including CityHash, xxHash, Murmur3, and FNV1. We found that wyhash provides the most favorable trade-off between performance and randomness. We chose not to include these benchmarks in the paper, as they are orthogonal to DLHT.

3.4.4 Iterator: weak and strong snapshots

We offer an iterator API that allows the client to iterate through all of DLHT’s key-value pairs. The iterator tracks its current index position. A Next() function fetches the next pair until it has iterated through all bins. We provide a strongly-consistent snapshot for the iterator via an index migration (i.e., a “resize” to a same-size index) that temporarily stalls updates – until transitioning all bins to a

Commodity Server	Two-socket 18-core Intel Xeon Gold 6254 72 h/w threads in total (incl. hyper-threads)
Hardware Caches	36× 1MB L2 2× 24.8MB L3
System Memory	8× 32GB DDR4-2933 (256GB in total)
OS / Kernel	Ubuntu 20.04.3 Linux 5.4.0-90-generic
(Source Code of) Evaluated Baselines	CLHT [62], MICA [40], GrowT [46], Leapfrog [55], DRAMHiT [56], Folly [15], Cuckoo [16], TBB [29]
Software Threads	1, 2, 4, 8, 16 , 32, 64, 71
Key Size, Value Size	Fixed /Variable 8B , 16B, ... , 1.5KB
Hash Function	modulo , wyhash
Allocator	mimalloc (2MB pages) , malloc
Number of Bins	16K, ..., 67M (index size = 4GB), ..., 1B (64GB)
Number of Keys	16K, ..., 100M , ..., 1B, 1.6B
Resizing	disabled , enabled
Access Patterns	uniform , skewed
DLHT modes	Inlined , Allocator, HashSet
Batching	disabled, batch-size = 1, 2, ..., 32 , 64, 128
Workloads/ Application scenarios/ Benchmarks	1. Get , 2. InsDel , 3. PutHeavy, 4. Resizing, 5. Single-thread, 6. Lock manager, 7. CXL emul. 8. YCSB, 9. TATP and Smallbank, 10. Hash Join

Table 2: Experimental configuration (default values in bold)

Snapshot state. However, our clients prefer a weakly-consistent snapshot that is non-blocking and does not need a migration.

3.4.5 Single-thread: synchronization-overhead-free

A client can opt-in to use DLHT with only a single-thread. DLHT includes three sources of overhead for thread-safety: ❶ lock-free algorithms, ❷ checking for concurrent resizes, and ❸ notifying other threads when entering/leaving DLHT. When configured for a single thread, we completely remove the second and third overheads. To alleviate the first overhead, we convert every CAS into a regular store, and downgrade atomic loads/stores into regular loads/stores. We also tried replacing the lock-free algorithms with simplified single-threaded algorithms, but the gain was negligible, and thus we abandoned that avenue.

3.4.6 Summary

At the request of its clients, DLHT offers variable-size keys/values, namespaces, a sophisticated hashing function, an iterator API, and single-threaded optimizations. The features that incur a performance penalty are disabled by default.

4 EXPERIMENTAL METHODOLOGY

Table 2 summarizes the testbed, the evaluated hashables, and the variables we study in our evaluation. Unless stated otherwise, the experiments use the default values (**bold** in **Table 2**). We next clarify aspects that do not fit in the table.

Testbed. We conduct the experiments on a server with two 18-core CPUs and two threads per core (72 hardware threads in total). We spare a thread for the OS; thus, the hashables use up to 71 threads. Across all experiments and for all hashables, we pin threads as follows. When running fewer than 18 software threads, we pin them all in the first socket. When running between 19 and 36, we

CLHT	the lock-free variant of CLHT [10, 11]
MICA	the CRCW variant of MICA2 [38, 39]
GrowT	the uaGrowT variant of GrowT [47]
Folly	the atomic hash map of Meta’s Folly [15]
Cuckoo	the concurrent Cuckoo hash map [17]
Leapfrog	Leapfrog map of Junction’s library [54]
TBB	the concurrent hash map of Intel’s OneTBB [29]
DRAMHiT	a lock-free non-resizable map (concur. work) [51]
DLHT(-NoBatch)	as in § 3 (without batching – § 3.3)

Table 3: Summary of evaluated Hashables.

balance them among the two sockets. When running more than 36, we use hyper-threads and again balance them among the sockets.

By default, we run experiments using 16 threads with 8-byte keys and 8-byte values inlined in the index. We typically use a simple modulo operation for hashing. Otherwise, we use wyhash [64]. We preload the mimalloc allocator [49] and configure it to use (2MB) huge pages. We chose wyhash and mimalloc after extensive experimentation with the available hash functions and allocators. For a fair comparison, we use the same allocator and huge pages in all hashables. When threads are pinned in both sockets, we interleave the bins of the index across both sockets. Unless stated otherwise, we instantiate DLHT, with 67 million bins, 8.3 million chained buckets, and we disable resizing. Before beginning the experiment, we populate the hashables with 100M keys.

Workloads. As shown in **Table 2**, we have two default workloads: 1) *Get* with 100% Gets and 2) *InsDel* with 50% Inserts and 50% Deletes. To execute a Get (similarly a Put), we first select one of the prepopulated keys uniformly at random, using a random-number generator (RNG). Inserts also use the RNG to select a key, but they choose a key that has not been prepopulated. This ensures that Inserts will always incur the full overhead of the insertion. In the *InsDel* workload, an Insert is always followed by a Delete to the same key. Details on non-default workloads are provided inline within each corresponding experimental discussion.

5 EVALUATION

In this section, we first, compare DLHT with state-of-the-art hashables and test our hypotheses (§ 5.1). Then we perform sensitivity studies over DLHT (§ 5.2) and investigate its performance on benchmarks and applications (§ 5.3).

5.1 Claims and comparison with state-of-the-art

In the introduction, we argued for the opportunity to achieve a billion requests per second on a *memory-resident* hashtable whose requests require main memory accesses. We claimed that state-of-the-art concurrent in-memory hashables are unable to reach this goal due to excessive blocking and inefficient handling of memory accesses. Finally, we argued that state-of-the-art designs either lack or offer impractical support for core operations (i.e., Put, Delete, or Resize). In this section, we set these hypotheses to the test.

Evaluated Hashables. In this section, we evaluate eight state-of-the-art concurrent in-memory hashables (baselines) and two variants of DLHT (shown in **Table 3**). We ensure all hashables are measured with the default configuration of **Table 2**. Note that almost all baselines only support up to 8-byte values. Therefore,




Figure 3: Get throughput




Figure 4: Get power-efficiency




Figure 5: InsDel throughput




Figure 6: Put-heavy throughput




Figure 7: Avg. Population throughput: Inserting 800M keys over a growing index.

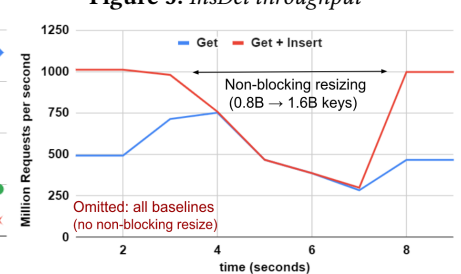


Figure 8: Gets, Inserts, and Non-blocking Resizing (transfer of 800M keys) in DLHT.

we use 8-byte values in this subsection. We extensively explore different sizes and configurations of DLHT in the subsection with sensitivity studies (§ 5.2).

For simplicity, we mainly focus on comparing with GrowT, DRAMHiT, Folly, CLHT, and MICA, as these are the fastest baselines. We analyzed the features of these hashmaps in Table 1. In short, GrowT, DRAMHiT, and Folly are open-addressing designs. Folly does not support Resizes or Deletes with slot reclamation. GrowT offers parallel but blocking Resizes and impractical Deletes that must perform a blocking transfer of all objects to a new index to reclaim index slots. The lock-free variant of CLHT, assumes unique values, does not support Puts, and cannot chain buckets. The latter limits collisions to one cache-line (i.e., three objects), after which it performs a Resize that problematically is blocking and single-threaded. MICA uses software prefetching for its memory accesses but is lock-based, lacks support for Resizes, and cannot inline keys and values inside its index. Thus, MICA needs at least two memory accesses to access or update a value and incurs (de)allocation overheads on all (Deletes) Inserts. Finally, DRAMHiT is the only design that combines frugal memory accesses and prefetching like DLHT. However, DRAMHiT lacks Resizing, its Deletes cannot reclaim index slots, its Puts (Inserts) may silently insert (update) key-values, and its batching may reorder a client's requests.

5.1.1 Get throughput and power-efficiency

Figure 3 illustrates the throughput of the Get workload in million requests per second (M. reqs/s) as we vary the number of threads. Our DLHT implementation outperforms all baselines, offers beyond 1B Gets/s, and scales almost linearly up to 32 cores across both available sockets. With hyper-threading enabled (i.e., 64 and 71 threads), DLHT saturates the memory bandwidth and peaks at 1.66B Gets/s. Its high performance mainly stems from serving each Get with a single, prefetched memory access and no write-backs.

In contrast, (DLHT-NoBatch,) GrowT, Folly, and CLHT can serve

Gets with a single memory access but do not leverage prefetching. Thus, they are over (2.2×) 3.5× slower than DLHT. Similarly, MICA uses prefetching, but its non-inlined approach requires at least two memory accesses for every Get, resulting in up to 4.8× lower throughput than DLHT. DRAMHiT, which combines an inlined index with software prefetching, is only 1.7× slower than DLHT. However, unlike DLHT, DRAMHiT may reorder the execution of requests inside a client's batch. We partially attribute DRAMHiT's lower performance to the design that does not let application threads to directly access the hashtable and the overheads of its asynchronous engine (e.g., producer-consumer queues).

Finally, Cuckoo, TBB, and Leapfrog mandate more than one memory access and do not use prefetching; hence their throughput is below 250 M. reqs/s. For simplicity, we omit those baselines from the rest of our graphs and focus our comparisons on GrowT, DRAMHiT, Folly, CLHT, and MICA.

Figure 4, shows the power efficiency on the Get workload in M. reqs/s per watt. The power efficiency of DLHT sees a steady increase up to 32 threads and peaks at 3.35M. reqs/s per watt, which is up to 1.7× higher than (the non-resizable) DRAMHiT and 3.6× higher than (the resizable) GrowT, the two most efficient baselines on Gets. In short, DLHT fulfills its purpose. It offers substantially higher throughput than the competition, surpassing 1.6B requests per second with its memory-access-aware design without sacrificing functionality. Next, we focus on functionality, starting with Deletes, which is a key point of criticism for open-addressing hashmaps.

5.1.2 Delete throughput

Deleting elements in a concurrent open-addressing scheme is cumbersome. In fact, neither DRAMHiT nor Folly support Deletes that free index slots. In this section, we investigate Deletes in open-addressing and contrast them with DLHT.

GrowT supports Deletes through tombstones. As the client is issuing Deletes, over time, the index fills with tombstones. When




Figure 9: Varying value size




Figure 10: Varying key size




Figure 11: Varying index size

the index gets filled beyond a threshold, GrowT creates a new index, where it must move all the alive keys. We perform an experiment with the common pattern of an Insert followed by a Delete to the same key (i.e., the default InsDel workload). We start with an empty hashtable that can fit 100 million keys. Figure 5 shows the completed Inserts and Deletes in M. reqs/s for increasing thread counts.

DLHT achieves up to 12.8× higher throughput than GrowT. This is because GrowT must move to a new index roughly every 100 million Deletes. Notably, this is still a favorable experiment for GrowT, because we allocate a very large table compared to the number of alive keys to amortize the cost of moving to a new table every 100M Deletes, and when it needs to move to a new table, only a handful of keys are actually moved. The latter is because there is always only at most one alive key per thread at any given moment.

The closed-addressing but inlined incapable scheme of MICA also suffers in this workload. This is because MICA requires at least two memory accesses and two write-backs to memory for each Insert and Delete to a key. As expected, CLHT performs comparably to DLHT-NoBatch, as both perform the Insert and Delete to the object via accessing and updating the same single cache-line. On top of that, DLHT with prefetching overlaps the blocking of memory accesses, hence achieving up to almost 3× the throughput of CLHT.

5.1.3 Put throughput

In Figure 6, we evaluate a Put-heavy workload with 50% Gets and 50% Puts over the default configuration of Table 2 and while varying threads. DLHT reaches up to 1042 M. reqs/s and outperforms both open-addressing baselines that do not utilize software prefetching up to 2.7×. DLHT provides a smaller benefit over DRAMHiT, which also exploits prefetching. However, Puts in DRAMHiT may silently insert an item in the index and may be performed out-of-order. As in the InsDel experiment, MICA requires multiple memory accesses for every Get and Put hence cannot maximize performance despite utilizing prefetching. Finally, recall that CLHT does not support Puts and is thus omitted from the graph.

5.1.4 Insert and Resizing throughput

To evaluate DLHT’s non-blocking resizing, we perform two experiments. First, we populate 800 million keys in an initially small index that grows on demand. Figure 7 shows the throughput as the thread count increases and includes the baselines that support resizing. As expected, we observe that inserting in a growing index is detrimental to CLHT, which cannot increase its population

throughput beyond 8 threads, as the cost of its single-threaded (blocking) resize dominates the population. GrowT’s population with parallel resizing scales better with threads, but its blocking nature and slower Inserts render GrowT’s throughput up to 3.9× lower than the parallel non-blocking population of DLHT.

To better demonstrate the non-blocking nature of DLHT’s resize, we run a second experiment where 32 threads populate the DLHT hashtable to 1.6B keys and 32 additional threads randomly perform Gets on these keys. Figure 8 shows the throughput of Gets and Inserts over time and a full index of 800M keys that starts resizing (to a 1.6B keys index) at approximately 3 seconds. At the early stages of the resize, threads complete Inserts as long as they still find space in the old index. Once they cannot fit their object into the old index, they help the resizer to accelerate the transfer. Helping allows this large transfer to complete in about 4 seconds. For comparison, it takes GrowT (CLHT) more than 5.6 (14) seconds to complete the same transfer in a parallel (sequential) but blocking way, during which all operations are stalled.

In contrast, the non-blocking resize of DLHT allows other operations (Gets in this case) to complete safely without waiting. Gets during a resize that find their bin of the old index in a state other than DoneTransfer, proceed as normal with a single (prefetched) memory access. As bins are transferred, more Gets must pay the overhead of accessing both the old and the new index until the resize completes. Hence, the Get throughput degrades over time during transferring but comes back once the transfer completes.

5.1.5 Index occupancy

To study the occupancy, we use wyhash and populate a growing index. We limit the number of link buckets to one-fifth of the bins to keep serving most operations with a single memory access. Recall that the lock-free CLHT does not support chaining. Meanwhile, DLHT chains up to 3 extra buckets per bin. Not surprisingly, even with a state-of-the-art hash function, we observe that CLHT’s inability to chain buckets results in low occupancy (1-5%). In contrast, DLHT gracefully handles collisions via its bounded chaining and sees 63-72% occupancy. Note that to maintain good performance, open-addressing hashables typically resize when they reach 30-50% occupancy (30% in GrowT’s codebase) [46, 47].

5.2 Sensitivity studies on DLHT

In this section, we isolate our focus on DLHT to characterize its performance while varying its features.




Figure 12: Varying batch size




Figure 13: Skew with 1000 hot keys




Figure 14: Enabling DLHT features

5.2.1 Varying value size and key size

In Figure 9, we vary the value size from 8 to 1500 bytes. When values are 8 bytes, they are inlined in the index, and we do not need to allocate any extra space for them. Apart from our two default workloads (Get and InsDel), we also run a workload called Get-Access. This is different from the Get workload as after every Get, we read the whole value. Recall that when the value is larger than 8 bytes, Gets only return a pointer to the value without accessing it.

For this reason, the Get workload is only slightly affected as the value size increases. The throughput of InsDel gradually decreases as the size of allocations increases. The throughput of Get-Access drops quickly as each get is associated with extra memory accesses to read the entire value from memory.

In Figure 10, we run the two default workloads while varying the key size from 8 to 256 bytes. Note the steep performance drop when the key is larger than 8 bytes. In this case, the key does not fit in the bucket slot and is stored with the value. This means that 1) all Gets must now dereference the pointer to the value to read the entire key (an 8-byte signature is left on the bucket for filtering), 2) we need to allocate more bytes per Insert to fit the key in the value, and 3) each Insert must write the key to the newly allocated value.

5.2.2 Varying index size

We next vary the index size of DLHT from 1MB (with 8K prepopulated keys) to 64GB (with 1B prepopulated keys). To study when prefetching is beneficial, apart from the two default workloads (Get and InsDel), we run a workload called Get-NoBatch, which is the same as Get but without batching.

We make two observations. First, because the smallest index (1MB) can fit inside the private L2 of each core, prefetching is not helpful, and thus batching has no benefit, only overheads. As the index grows, prefetching, and thus batching, becomes increasingly beneficial. Second, we see that unlike Get, InsDel does not benefit from a small cache-resident index. This is because a smaller index increases the chances of conflicts in the same bin. This affects the performance as Inserts and Deletes must CAS the header of the bin.

5.2.3 Varying batch size

Batching is a key feature of DLHT where software prefetching is used to overlap memory accesses for a request with useful work on other requests. Figure 12 shows the throughput of DLHT under the default configuration, while varying the batch degree, i.e., the number of requests grouped in a batch. Apart from our two default

workloads (Get and InsDel), we also show a workload called Get-Resizing, which is the same as Get but with resizing enabled. This does not mean that the workload resizes the index, but rather that we have compiled DLHT with the capability to resize, if needed. The bars denoted as “No batch”, do not use the batching API.

We make three observations. First, batching can be very effective in boosting performance. The gains saturate around 24 accesses; we attribute this limitation to hardware (number of MSHRs and TLB size). Second, batching improves throughput if there are at least four requests (two for Get) that can be overlapped. Otherwise, the overhead of creating the request array trumps the benefit of overlapping memory accesses. Third, enabling resizing hurts performance more when no batching is used. This is because resize requires each request to perform two additional atomic stores (§ 3.2.5). Batching amortizes this overhead amongst the batched requests.

5.2.4 Skew

To study the behavior of DLHT under skew, we perform an experiment with the default configuration, but 1000 keys are hot and receive an increasing percentage of accesses. Figure 13 shows the throughput of DLHT for the Get, InsDel, and Get-NoBatch while we increase the percentage of skewed accesses. As expected, in Get and Get-NoBatch the throughput increases with skew, surpassing 2.2B Gets/s as both benefit from the locality of skewed accesses. When all accesses target hot keys, prefetching stops being useful and Get-NoBatch outperforms the Get workload that includes the overheads of batching. Under high skew, InsDel manifests a high number of conflicts which impacts its performance.

5.2.5 Enabling features

DLHT has several features that can be enabled at runtime or compile-time. Enabling these features typically incurs a performance penalty. Figure 14 shows the throughput in M. reqs/s of the two default workloads (Get and InsDel) while enabling features either in a *stacked* fashion, where features are enabled on top of each other, or one-at-a-time (i.e., *single*). We explain the meaning of each bar below.

Default. This bar shows the throughput of the default configuration (Table 2), with the exception that we use 32-byte values. This serves as the baseline for the rest of the bars.

Resizing. DLHT is compiled with resizing enabled. Note that no resizing actually takes place in this experiment, but DLHT incurs a performance penalty, as each thread must notify other threads when entering and leaving the DLHT index, as well as check if




Figure 15: Latency of Gets and InsDel




Figure 16: Single thread application




Figure 17: Lock Manager

other threads have initiated resizing.

Hashing. DLHT uses the wyhash hash function [64] instead of a simple modulo operation to hash a key to a bin. Naturally, this also further reduces throughput.

Variable value size. DLHT is compiled to allow every value to have a different size. The performance penalty is very small because the only overhead is that the size of each value must be stored in the header of the value.

Variable key size. DLHT allows every key to have a different size. In this experiment, all used keys are smaller than 8 bytes. Hence the only overhead is storing and checking the key size in the most significant bits of the pointer (discussed in § 3.4.1). The overheads for larger keys are studied in § 5.2.1

Namespaces. DLHT is compiled with namespaces, incurring a small penalty, as the namespace is stored and checked in each search in the most significant bits of the value pointer.

No mimalloc. Finally, we use the standard library malloc instead of mimalloc. Naturally, this impacts only the InsDel workload, which calls either malloc or free on every request.

Notably, this is far from an exhaustive evaluation of all possible combinations of runtime and compile-time options, which would have been intractable to evaluate. Rather it is a selection of popular features among our clients.

5.2.6 Latency

While we focus on delivering high throughput, we also study the latency of DLHT in this section. Figure 15 shows the average and tail (99th%) latency in nanoseconds for our Get and InsDel workloads, as a function of the load. Unsurprisingly, the nature of memory-resident workloads translates into average latencies of 100s nanoseconds which increase with the load. We also observe that the tail is below a microsecond, even under high load. Naturally, Get has lower latencies than InsDel which involves slower CAS operations.

5.3 Applications and Benchmarks over DLHT

In this section, we evaluate applications and benchmarks that represent some of our target use-cases.

5.3.1 Optimized single-threaded application

We optimized DLHT to deliver high performance on single-threaded applications. Figure 16 shows the single-threaded throughput with and without our optimizations under four workloads. Apart from our two default workloads, we also run InsDel-Resize, which enables resizing, and InsDel-Resize-NoBatch, which enables resizing but disables batching.

Starting from InsDel, we see a 31% improvement as a concur-

rent Insert requires two CASes, and a Delete requires one CAS. With the optimization, we convert these into regular stores. InsDel-resizing sees a 35% improvement, as resizing requires notifying other threads when entering and leaving DLHT, via atomic stores. The optimization removes the need to notify other threads, as there are no other threads. The impact of this is amplified when batching is disabled (91% improvement), as the enter/exit notification must now happen once per request instead of once per batch. Finally, we note that the optimization has no impact on the Get workload because on x86 8-byte atomic loads have a trivial overhead.

5.3.2 Remote memory: CXL emulation

Memory latency seldom improves [65]. In fact, recent trends increase capacity over cost by introducing slower memories that are denser (e.g., NVM) or reside far (e.g., as in RDMA or CXL). Hence, we believe that memory-aware designs, which minimize accesses and exploit software prefetching, will remain critical in next-generation deployments.

To test this hypothesis, we run the Get workload while emulating a CXL deployment of DLHT. For this, we run 36 threads on the first socket and pin DLHT’s memory on the remote socket – as in recent works [37, 48]. We observe that DLHT (with prefetching) offers 2.9× the throughput of DLHT-NoBatch. As expected, remote memory lowers the throughput (to about half) compared to local.

5.3.3 Lock manager over HashSet

This workload is inspired by a client that used DLHT’s HashSet to implement a database lock manager. Inserting and deleting a key over the HashSet mimics locking and unlocking a key (or a set of keys). Recall that DLHT’s batching interface preserves the order of requests. As such, DLHT’s order-preserving batching is an especially good fit for a lock manager, which typically wants to implement protocols similar to two-phase-locking that grab locks in some order and later release them. This is also an application example where DRAMHiT’s out-of-order batching, which may reorder requests to overlap their accesses, can introduce deadlocks.

Figure 17 shows the throughput of locks and unlocks while varying the threads for DLHT(-NoBatch). When batching is used, the throughput scales well with the number of threads and is able to peak at almost 1.5B locks/unlocks per second. As expected without batching, when the memory accesses cannot be masked, the throughput is up to 2.2× lower.

5.3.4 Single-key benchmarks (YCSB Mixes)

We study the performance of DLHT under YCSB [9], the popular single-key benchmark. Figure 18 shows the throughput of four




Figure 18: Single-key YCSB mixes




Figure 19: Transactional benchmarks




Figure 20: Hash Joins

	Characteristic	Tables	Columns	Txs	Read Txs
TATP	read-intensive	4	51	7	80%
Smallbank	write-intensive	3	6	6	15%

Table 4: Summary of evaluated transactional benchmarks.

YCSB mixes (shown in legend) as the threads increase. The configuration follows the defaults of Table 2.

All four mixes scale well up to 32 threads where cores in both sockets are used and do saturate the memory bandwidth when hyper-threading is enabled. The update-only (YCSB F) mix peaks at about half the throughput of read-only (YCSB C), saturating the memory bandwidth at 32 threads. This is expected, since it incurs double the memory traffic as every accessed cache-line is updated and written back to memory.

5.3.5 Multi-key benchmarks (OLTP Transactions)

To study the performance of DLHT on multi-key transactions, we implement and evaluate two popular OLTP benchmarks, TATP [52], which is read-intensive, and Smallbank [5] that is write-intensive (summarized in Table 4). In this experiment, the same threads run both the logic that generates transactions and the DLHT logic. We ensure the benchmarks are memory-resident by populating their tables with 1M subscribers (TATP) and 10M accounts (Smallbank).

Figure 19, shows the throughput of the two benchmarks while varying threads. We see that both benchmarks scale well with the number of threads and surpass 100M transactions per second. With 64 threads TATP (Smallbank) reaches 175M (129M) transactions per second. Note that both benchmarks only access and/or update just a few objects in a transaction. Thus, TATP, with fewer updates, outperforms Smallbank as it incurs fewer write-backs to memory.

5.3.6 Hash Join (OLAP Application)

We evaluate a non-partitioned join over DLHT. We use workload A from [44] in which both relations are memory-resident. Workload A consists of 16B tuples and a build relation (R) and probe relation (S) with size 2^{27} and 2^{31} similarly to [44, 57]. Figure 20, shows the throughput ($\frac{|R| + |S|}{runtime}$) while varying the number of threads.

In joins, where batching is naturally applicable, DLHT (with software prefetching) achieves up to 1.4B tuples/sec. This translates to a 2.2x benefit over DLHT-NoBatch. Interestingly, the memory-aware and non-blocking design of DLHT delivers a very high throughput CPU-based join – despite large tables that cannot fit in cache and under no partitioning or specializations for joins (e.g., eschewing synchronization for the read-only probing phase). We note that

Baseline	Comparison with DLHT
CLHT	- 3.5x lower Get throughput 8x slower Population - more than 10x worse occupancy (cannot chain buckets) - Blocking resize, no Puts, assumes unique values, $\leq 8B$ keys/values
MICA	- 4.8x lower Get throughput - Non-resizable, lock-based, non-inlining
GrowT	- 3.5x lower Get throughput 3.9x slower Population - 12.8x lower InsDel throughput (tombstone-based deletes)
Folly	- 3.5x lower Get throughput - Non-resizable, deletes cannot reclaim slots, $\leq 8B$ keys/values only
DRAMHiT	- 1.7x lower Get throughput - Non-resizable, deletes cannot reclaim slots, $\leq 8B$ keys/values only - Only Upserts (no pure Insert/Put), batching may reorder requests

Table 5: Comparison summary of DLHT and (fastest) baselines.

partitioning and other such optimizations are synergistic to the features of DLHT. We leave a more elaborate study for future work, as joins are not the main focus of this paper.

6 RELATED WORK

This section extends related work beyond the designs summarized in Table 5, which we also analyzed in § 2 and § 5.1.

Numerous works focus on high-throughput hash tables. Some exploit hardware acceleration via FPGAs [6, 30, 34, 36], GPUs [3, 27, 44, 45, 57, 66], or SIMD [1, 15]. Others propose alternative probing [13, 17, 59], efficient concurrency control [29, 41], or focus on persistency [7, 8, 28, 35, 43, 50, 67]. Finally, some designs offer non-blocking resizes [18, 19, 22, 42, 63]. Unlike DLHT, none of these designs exploit software prefetching to hide latencies and most need two or more memory accesses to serve a request from memory.

7 CONCLUSION

We showed that existing in-memory hash tables are unable to reach a billion requests per second in a commodity server when accessing memory. In particular, we showed that state-of-the-art hash tables forfeit core functionality, block excessively, and handle accesses inefficiently in memory-resident workloads. We presented DLHT, a practical non-blocking and memory-aware hashtable. DLHT addresses the issues of existing designs to offer complete functionality, fast non-blocking Resizes, 1.6B requests per second, and 3.5x (12x) the throughput of state-of-the-art (open-) closed-addressing designs on Gets (Deletes).

REFERENCES

- [1] Abseil. 2023. Abseil C++ library: An open-source collection of C++ libraries from Google. <https://github.com/abseil/abseil-cpp>.
- [2] Michael A. Bender, Bradley C. Kuszmaul, and William Kuszmaul. 2021. Linear Probing Revisited: Tombstones Mark the Demise of Primary Clustering. *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (2021), 1171–1182.
- [3] Alex D. Breslow, Dong Ping Zhang, Joseph L. Greathouse, Nuwan Jayasena, and Dean M. Tullsen. 2016. Horton Tables: Fast Hash Tables for in-Memory Data-Intensive Computing. In *Proceedings of the 2016 USENIX Conference on Usenix Annual Technical Conference (USENIX ATC '16)*. USENIX Association, USA, 281–294.
- [4] Nathan Bronson, Zach Amsden, George Cabrera, Prasad Chakka, Peter Dimov, Hui Ding, Jack Ferris, Anthony Giardullo, Sachin Kulkarni, Harry Li, Mark Marchukov, Dmitri Petrov, Lovro Puzar, Yee Jiun Song, and Venkat Venkataramani. 2013. TAO: Facebook's Distributed Data Store for the Social Graph. In *Proceedings of the 2013 Conference on Annual Technical Conference (ATC'13)*. USENIX, Berkeley, 49–60.
- [5] Michael J Cahill, Uwe Röhm, and Alan D Fekete. 2009. Serializable isolation for snapshot databases. *ACM Transactions on Database Systems (TODS)* 34, 4 (2009), 1–42.
- [6] Sai Rahul Chalamalasetti, Kevin Lim, Mitch Wright, Alvin AuYoung, Parthasarathy Ranganathan, and Martin Margala. 2013. An FPGA Memcached Appliance. In *Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA '13)*. Association for Computing Machinery, New York, NY, USA, 245–254.
- [7] Badrish Chandramouli, Guna Prasaad, Donald Kossmann, Justin Levandoski, James Hunter, and Mike Barnett. 2018. FASTER: A Concurrent Key-Value Store with In-Place Updates. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD '18)*. Association for Computing Machinery, New York, NY, USA, 275–290.
- [8] Zhangyu Chen, Yu Hua, Bo Ding, and Pengfei Zuo. 2020. Lock-Free Concurrent Level Hashing for Persistent Memory. In *Proceedings of the 2020 USENIX Conference on Usenix Annual Technical Conference (USENIX ATC'20)*. USENIX Association, USA, Article 55, 14 pages.
- [9] Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. 2010. Benchmarking Cloud Serving Systems with YCSB. In *Proceedings of the 1st ACM Symposium on Cloud Computing (SoCC '10)*. ACM, New York, NY, USA, 143–154.
- [10] Tudor David, Rachid Guerraoui, and Vasileios Trigonakis. 2015. Asynchronized Concurrency: The Secret to Scaling Concurrent Search Data Structures. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '15)*. Association for Computing Machinery, New York, NY, USA, 631–644.
- [11] Tudor Alexandru David, Rachid Guerraoui, Tong Che, and Vasileios Trigonakis. 2014. Designing ASCY-compliant Concurrent Search Data Structures. (2014), 23. <http://infoscience.epfl.ch/record/203822>
- [12] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. 2007. Dynamo: Amazon's Highly Available Key-value Store. *SIGOPS Oper. Sys.* 41, 6 (2007), 5–20.
- [13] Aleksandar Dragojević, Dushyanth Narayanan, Miguel Castro, and Orion Hodson. 2014. FaRM: Fast Remote Memory. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. USENIX Association, Seattle, WA, 401–414.
- [14] Mostafa Elhemali, Niall Gallagher, Nick Gordon, Joseph Idziorek, Richard Krog, Colin Lazier, Erben Mo, Akhlesh Mritunjai, Somasundaram Perianayagam, Tim Rath, Swami Sivasubramanian, James Christopher Sorenson III, Sroaj Sothikul, Doug Terry, and Akshat Vig. 2022. Amazon DynamoDB: A Scalable, Predictably Performant, and Fully Managed NoSQL Database Service. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*. USENIX Association, Carlsbad, CA, 1037–1048. <https://www.usenix.org/conference/atc22/presentation/elhemali>
- [15] Facebook. 2023. Folly: An open-source C++ library developed and used at Facebook. <https://github.com/facebook/folly>.
- [16] Bin Fan. 2013. libcuckoo: A high-performance concurrent hash table. <https://github.com/efficient/libcuckoo>.
- [17] Bin Fan, David Andersen, and Michael Kaminsky. 2013. MemC3: Compact and Concurrent MemCache with Dumber Caching and Smarter Hashing. In *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation (nsdi'13)*. USENIX Association, Berkeley, CA, USA, 371–384. <http://dl.acm.org/citation.cfm?id=2482626.2482662>
- [18] Panagiota Fatourou, Nikolaos D. Kallimanis, and Thomas Ropars. 2018. An Efficient Wait-Free Resizable Hash Table. In *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures (SPAA '18)*. Association for Computing Machinery, New York, NY, USA, 111–120.
- [19] H. Gao, J.F. Groote, and W.H. Hesselink. 2004. Almost wait-free resizable hash-tables. In *18th International Parallel and Distributed Processing Symposium, 2004. Proceedings.* 50–.
- [20] H. Gao, J. F. Groote, and W. H. Hesselink. 2005. Lock-Free Dynamic Hash Tables with Open Addressing. *Distrib. Comput.* 18, 1 (jul 2005), 21–42.
- [21] K. Gharachorloo, A. Gupta, and J. Hennessy. 1992. Hiding Memory Latency using Dynamic Scheduling in Shared-Memory Multiprocessors. In *[1992] Proceedings of the 19th Annual International Symposium on Computer Architecture.* 22–33.
- [22] Michael Greenwald. 2002. Two-Handed Emulation: How to Build Non-Blocking Implementations of Complex Data-Structures Using DCAS. In *Proceedings of the Twenty-First Annual Symposium on Principles of Distributed Computing (PODC '02)*. Association for Computing Machinery, New York, NY, USA, 260–269.
- [23] Tim Harris, James Larus, and Ravi Rajwar. 2010. *Transactional Memory, 2nd Edition* (2nd ed.). Morgan and Claypool Publishers.
- [24] Yongjun He, Jiacheng Lu, and Tianzheng Wang. 2020. CoroBase: Coroutine-Oriented Main-Memory Database Engine. *Proc. VLDB Endow.* 14, 3 (nov 2020), 431–444.
- [25] Maurice Herlihy and Nir Shavit. 2008. *The Art of Multiprocessor Programming*. Morgan Kaufmann Publishers Inc., USA.
- [26] Maurice Herlihy and Jeannette Wing. 1990. Linearizability: A Correctness Condition for Concurrent Objects. *ACM Trans. Program. Lang. Syst.* 12, 3 (July 1990), 463–492.
- [27] Tayler H. Hetherington, Mike O'Connor, and Tor M. Aamodt. 2015. MemcachedGPU: Scaling-up Scale-out Key-Value Stores. In *Proceedings of the Sixth ACM Symposium on Cloud Computing (SoCC '15)*. Association for Computing Machinery, New York, NY, USA, 43–57.
- [28] Daokun Hu, Zhiwen Chen, Jianbing Wu, Jianhua Sun, and Hao Chen. 2021. Persistent Memory Hash Indexes: An Experimental Evaluation. 14, 5 (jan 2021), 785–798.
- [29] Intel Corporation. 2023. oneTBB: Intel oneAPI Threading Building Blocks. <https://github.com/oneapi-src/oneTBB>.
- [30] Zsolt István, Gustavo Alonso, Michaela Blott, and Kees Vissers. 2015. A Hash Table for Line-Rate Data Processing. *ACM Trans. Reconfigurable Technol. Syst.* 8, 2, Article 13 (mar 2015), 15 pages.
- [31] Rosa M. Jiménez and Conrado Martínez. [n.d.]. *On Deletions in Open Addressing Hashing.* 23–31.
- [32] Donald E. Knuth. 1998. *The Art of Computer Programming, Volume 3: (2nd Ed.) Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., USA.
- [33] Christoph Lameter. 2005. Effective synchronization on Linux/NUMA systems.
- [34] Maysam Lavasani, Hari Angepat, and Derek Chiu. 2014. An FPGA-Based In-Line Accelerator for Memcached. *IEEE Comput. Archit. Lett.* 13, 2 (jul 2014), 57–60.
- [35] 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)*. Association for Computing Machinery, New York, NY, USA, 462–477.
- [36] Bojie Li, Zhenyuan Ruan, Wencong Xiao, Yuanwei Lu, Yongqiang Xiong, Andrew Putnam, Enhong Chen, and Lintao Zhang. 2017. KV-Direct: High-Performance In-Memory Key-Value Store with Programmable NIC. In *Proceedings of the 26th Symposium on Operating Systems Principles (SOSP '17)*. ACM, New York, NY, USA, 137–152.
- [37] Huaicheng Li, Daniel S. Berger, Lisa Hsu, Daniel Ernst, Pantea Zardoshti, Stanko Novakovic, Monish Shah, Samir Rajadnya, Scott Lee, Ishwar Agarwal, Mark D. Hill, Marcus Fontoura, and Ricardo Bianchini. 2023. Pond: CXL-Based Memory Pooling Systems for Cloud Platforms. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS 2023)*. Association for Computing Machinery, New York, NY, USA, 574–587.
- [38] Sheng Li, Hyeontaek Lim, Victor W. Lee, Jung Ho Ahn, Anuj Kalia, Michael Kaminsky, David Andersen, Seongil O, Sukhan Lee, and Pradeep Dubey. 2016. Full-Stack Architecting to Achieve a Billion-Requests-Per-Second Throughput on a Single Key-Value Store Server Platform. *ACM Trans. Comput. Syst.* 34, 2, Article 5 (April 2016), 30 pages.
- [39] Sheng Li, Hyeontaek Lim, Victor W. Lee, Jung Ho Ahn, Anuj Kalia, Michael Kaminsky, David G. Andersen, O. Seongil, Sukhan Lee, and Pradeep Dubey. 2015. Architecting to Achieve a Billion Requests per Second Throughput on a Single Key-Value Store Server Platform. *SIGARCH Comput. Archit. News* 43, 3S (jun 2015), 476–488.
- [40] Hyeontaek Lim. 2017. MICA2: A fast in-memory key-value store. <https://github.com/efficient/mica2>.
- [41] Hyeontaek Lim, Dongsu Han, David Andersen, and Michael Kaminsky. 2014. MICA: A Holistic Approach to Fast In-memory Key-value Storage. In *Proceedings of the 11th Networked Systems Design and Implementation (NSDI'14)*. USENIX Association, USA, 429–444.
- [42] Yujie Liu, Kunlong Zhang, and Michael Spear. 2014. Dynamic-Sized Nonblocking Hash Tables. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing (PODC '14)*. Association for Computing Machinery, New York, NY, USA, 242–251.
- [43] Baotong Lu, Xiangpeng Hao, Tianzheng Wang, and Eric Lo. 2020. Dash: Scalable Hashing on Persistent Memory. *Proc. VLDB Endow.* 13, 8 (apr 2020), 1147–1161.

- [44] Clemens Lutz, Sebastian Breß, Steffen Zeuch, Tilmann Rabl, and Volker Markl. 2020. Pump Up the Volume: Processing Large Data on GPUs with Fast Interconnects. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD '20)*. Association for Computing Machinery, New York, NY, USA, 1633–1649.
- [45] Clemens Lutz, Sebastian Breß, Steffen Zeuch, Tilmann Rabl, and Volker Markl. 2022. Triton Join: Efficiently Scaling to a Large Join State on GPUs with Fast Interconnects. In *Proceedings of the 2022 International Conference on Management of Data (SIGMOD '22)*. Association for Computing Machinery, New York, NY, USA, 1017–1032.
- [46] Tobias Maier. 2023. growt: A concurrent growing hash table. <https://github.com/TooBiased/growt>.
- [47] Tobias Maier, Peter Sanders, and Roman Dementiev. 2019. Concurrent Hash Tables: Fast and General(?)! *ACM Trans. Parallel Comput.* 5, 4, Article 16 (feb 2019), 32 pages.
- [48] Hasan Al Maruf, Hao Wang, Abhishek Dhanotia, Johannes Weiner, Niket Agarwal, Pallab Bhattacharya, Chris Petersen, Mosharaf Chowdhury, Shobhit Kanaujia, and Prakash Chauhan. 2023. TPP: Transparent Page Placement for CXL-Enabled Tiered-Memory. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3 (ASPLOS 2023)*. Association for Computing Machinery, New York, NY, USA, 742–755.
- [49] Microsoft. 2023. mimalloc: A general purpose allocator with excellent performance. <https://github.com/microsoft/mimalloc>.
- [50] Moohyeon Nam, Hokeun Cha, Young-Ri Choi, Sam H. Noh, and Beomseok Nam. 2019. Write-Optimized Dynamic Hashing for Persistent Memory. In *Proceedings of the 17th USENIX Conference on File and Storage Technologies (FAST'19)*. USENIX Association, USA, 31–44.
- [51] Vikram Narayanan, David Detweiler, Tianjiao Huang, and Anton Burtsev. 2023. DRAMHIT: A Hash Table Architected for the Speed of DRAM. In *Proceedings of the Eighteenth European Conference on Computer Systems (EuroSys '23)*. Association for Computing Machinery, New York, NY, USA, 817–834.
- [52] Simo Neuvonen, Antoni Wolski, Markku Manner, and Vilho Raatikka. 2009. Telecom Application Transaction Processing Benchmark. <http://tatbenchmark.sourceforge.net/>.
- [53] W. W. Peterson. 1957. Addressing for Random-Access Storage. *IBM Journal of Research and Development* 1, 2 (1957), 130–146.
- [54] Jeff Preshing. 2016. New Concurrent Hash Maps for C++. <https://preshing.com/20160201/new-concurrent-hash-maps-for-cpp/>.
- [55] Jeff Preshing. 2018. junction: A library of concurrent data structures in C++. <https://github.com/preshing/junction>.
- [56] Mars Research. 2023. DRAMHIT. <https://github.com/mars-research/DRAMHIT>
- [57] Stefan Schuh, Xiao Chen, and Jens Dittrich. 2016. An Experimental Comparison of Thirteen Relational Equi-Joins in Main Memory. In *Proceedings of the 2016 International Conference on Management of Data (SIGMOD '16)*. Association for Computing Machinery, New York, NY, USA, 1961–1976.
- [58] Michael L. Scott. 2013. Shared-Memory Synchronization.
- [59] Ori Shalev and Nir Shavit. 2006. Split-Ordered Lists: Lock-Free Extensible Hash Tables. *J. ACM* 53, 3 (may 2006), 379–405.
- [60] Panagiotis Sioulas, Periklis Chrysogelos, Manos Karpathiotakis, Raja Appuswamy, and Anastasia Ailamaki. 2019. Hardware-Conscious Hash-Joins on GPUs. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*.
- [61] Alex Stivala, Peter J. Stuckey, Maria Garcia de la Banda, Manuel Hermenegildo, and Anthony Wirth. 2010. Lock-Free Parallel Dynamic Programming. *J. Parallel Distrib. Comput.* 70, 8 (aug 2010), 839–848.
- [62] Vasileios Trigonakis. 2021. CLHT: Cache-Line Hash Table. <https://github.com/LPD-EPFL/CLHT>.
- [63] Junchang Wang, Xiong Fu, Fu Xiao, and Chen Tian. 2020. DHash: Enabling Dynamic and Efficient Hash Tables. [arXiv:cs.DC/2006.00819](https://arxiv.org/abs/2006.00819)
- [64] Yi Wang. 2023. wyhash: ultra fast hash. <https://github.com/wangyi-fudan/wyhash>.
- [65] Wm A Wulf and Sally A McKee. 1995. Hitting the memory wall: Implications of the obvious. *ACM SIGARCH computer architecture news* 23, 1 (1995), 20–24.
- [66] Kai Zhang, Kaibo Wang, Yuan Yuan, Lei Guo, Rubao Lee, and Xiaodong Zhang. 2015. Mega-KV: A Case for GPUs to Maximize the Throughput of in-Memory Key-Value Stores. *Proc. VLDB Endow.* 8, 11 (jul 2015), 1226–1237.
- [67] Xingsheng Zhao, Chen Zhong, and Song Jiang. 2023. TurboHash: A Hash Table for Key-Value Store on Persistent Memory. In *Proceedings of the 16th ACM International Conference on Systems and Storage (SYSTOR '23)*. Association for Computing Machinery, New York, NY, USA, 35–48.