



PathCAS: An Efficient Middle Ground for Concurrent Search Data Structures

Trevor Brown
University of Waterloo
Canada
me@tbrown.pro

William Sigouin
University of Waterloo
Canada
wpsigoui@uwaterloo.ca

Dan Alistarh
Institute of Science and Technology
Austria
dan.alistarh@ist.ac.at

Abstract

To maximize the performance of concurrent data structures, researchers have often turned to highly complex fine-grained techniques, resulting in efficient and elegant algorithms, which can however be often difficult to understand and prove correct. While simpler techniques exist, such as transactional memory, they can have limited performance or portability relative to their fine-grained counterparts. Approaches at both ends of this complexity-performance spectrum have been extensively explored, but relatively less is known about the middle ground: approaches that are willing to sacrifice some performance for simplicity, while remaining competitive with state-of-the-art handcrafted designs. In this paper, we explore this middle ground, and present PathCAS, a primitive that combines ideas from multi-word CAS (KCAS) and transactional memory approaches, while carefully avoiding overhead. We show how PathCAS can be used to implement efficient search data structures relatively simply, using an internal binary search tree as an example, then extending this to an AVL tree. Our best implementations outperform many handcrafted search trees: in search-heavy workloads, it rivals the BCCO tree [5], the fastest known concurrent binary tree in terms of search performance [3]. Our results suggest that PathCAS can yield concurrent data structures that are relatively easy to build and prove correct, while offering surprisingly high performance.

CCS Concepts: • **Computing methodologies** → **Concurrent algorithms**; *Shared memory algorithms*.

Keywords: concurrent data structures, search trees, non-blocking algorithms, lock-free, synchronization primitives

1 Introduction

Significant work has been invested in building scalable concurrent variants of fundamental data structures, and fast




Figure 1. AVL trees using PathCAS vs state-of-the-art transactional memory. 10% updates, 1M key trees. x-axis = number of threads. y-axis = millions of operations per second.

implementations are now known for many instances, from search trees to hash tables, or to containers such as queues and stacks [25]. On the one hand, designs based on *fine-grained locking* or *fine-grained lock-free algorithms*, have arguably emerged as the best-performing solution [3]. Yet, such designs tend to have high complexity, and are notoriously difficult to analyze and prove correct.

On the other hand, significant attention has been given to *general* techniques for obtaining fast and simple concurrent data structures. The classic example is transactional memory (TM) [24], which is now available in software, hardware, and hybrid variants, and allows one to derive concurrent implementations from sequential ones with lower programming effort relative to fine-grained designs. When TM is available, such designs can provide excellent performance.

However, TM-based designs still have drawbacks. Software TM (STM) provides a hardware-independent alternative to HTM, but can incur higher overheads. Moreover, although hardware transactional memory (HTM) is technically available on many platforms, via Intel’s TSX/TSX-NI, IBM’s POWER8/9 TM and ARM’s TME [30], it is notably missing from AMD chips (despite the proposal of ASF [10] more than a decade ago), and has been disabled in Intel and recent POWER processors [14, 27, 29], due to various concerns, chief among which is security.

In this context, it is natural to seek a middle ground between the high efficiency, but high complexity, of fine-grained designs, and the relative ease-of-use, but potential pitfalls, of general designs such as the ones based on TM. A number of



This work is licensed under a Creative Commons Attribution International 4.0 License.

PPoPP ’22, April 2–6, 2022, Seoul, Republic of Korea

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9204-4/22/04.

<https://doi.org/10.1145/3503221.3508410>

techniques exploring this trade-off have been investigated over the years, often based on either restricted STMs or extended Multi-Compare Multi-Swap (MCMS) [38] implementations. However, as we illustrate later in the paper, known instances of these techniques often fail to scale in the context of high-performance data structures.

We revisit this question, and present a primitive for building correct and efficient concurrent search data structures from scratch, called *PathCAS*. *PathCAS* combines key ideas from efficient multi-word compare-and-swap (KCAS), and transactional memory, to allow for concurrent data structures which are both efficient, and easier to reason about than hand-crafted data structures using fine-grained primitives.

This mechanism is reminiscent of STM techniques, but it has two **semantic restrictions**: (1) *PathCAS does not guarantee opacity*, and (2) *PathCAS has a bounded read-set*. These restrictions allow for significant performance benefits, as exemplified in Figure 1, which we believe are worth the increase in programming complexity compared to TM.¹ Despite *PathCAS* being less expressive than TM, it is still sufficient to implement useful data structures.

We begin by describing the semantics and rationale behind *PathCAS* in Section 2, and then illustrate how it can be used to implement a simple concurrent unbalanced internal BST, a data structure whose concurrent implementations warranted publication on their own, e.g. [13, 16, 17, 26, 34, 35], in Section 3. To illustrate the difficulty of implementing *correct* variants of these data structures, we note that, during our investigation, we identified a correctness bug in the *lock-based internal BST* of Drachler et al. [16], and that the publicly-available implementations of the *lock-free internal BSTs* of [26] and [35] fail experimental validation tests, and still lack complete correctness proofs. (We defer a detailed discussion of these issues to the full version of this paper.) In this context, *PathCAS* provides an implementation that is both efficient and is easy to prove correct.

To further highlight the expressive power of *PathCAS*, we also show how a lock-free *balanced* version of this tree can be derived, creating an implementation of a *fully-internal relaxed AVL tree* (Section 4), which performs favorably when compared to state-of-the-art solutions (please see Figure 1).

On the practical side, we present two efficient implementations of *PathCAS*: an HTM-enabled one which targets Intel systems, and a software-only variant applicable to AMD systems, and use them to empirically validate the above data structure designs. We perform an in-depth comparison

¹Bounding the read-set size is not strictly necessary, but doing so helps us avoid the overheads associated with dynamically sized data structures. One might imagine, for example, first filling a fixed array with nodes, then appending further nodes to a linked list. Checking whether a read should be added to the array or the linked list would require an extra branch for each read (possibly on a hot code path). Even with branch prediction, we have found this overhead to be significant in unbounded transactional memory implementations.

against previous methods: from HTM- and STM-based designs, to fine-grained lock-free variants, across both Intel and AMD systems with up to 256 threads. We find that *PathCAS* data structures are highly competitive, across the range from read-heavy to update-heavy workloads. In particular, our unbalanced BST implementation manages to outperform the state-of-the-art in unbalanced BSTs, and our balanced implementation matches the performance of the fastest known balanced BST in read-mostly workloads.

In sum, *PathCAS* introduces an additional point in the trade-off between expressiveness and programming effort, on the one hand, and efficiency of the resulting data structures, on the other. Although *PathCAS* builds on known techniques, the resulting mechanisms are novel, and our search tree implementations can achieve state-of-the-art results on a variety of workloads.

2 Related Work

The question of identifying synchronization primitives with the “right” balance between expressivity and efficiency is as old as the field of concurrency. We describe related work with the goal of situating *PathCAS* on the spectrum of synchronization techniques that have been developed.

Treiber [39] gave one of the first illustrations of a non-blocking data structure via CAS, while seminal work by Herlihy [23] showed that CAS is *universal*. Anderson and Moir gave constant-time implementations of *Load Link* (LL) and *Store Conditional* (SC) from CAS [1], which is more expressive than CAS and helps circumvent ABA problems. Luchangco et al. [31] expand on LL/SC by implementing *k-compare-single-swap*, allowing for the change of a single field to be conditional on *multiple* fields containing their expected values. Another extension of LL/SC by Brown et al. [7] introduced *LLX* and *SCX*, which function on *data records*. Data records contain *multiple* related fields which are loaded by LLX, and SCX only succeeds in changing a single memory location if none of the fields loaded by LLX have changed since its invocation. LLX/SCX is less expressive than *PathCAS*, as it atomically: changes *one* field to a *new value*, and *marks* some number of nodes. Brown et al. [7] introduced several search tree design based on LLX/SCX, one of which we compare with in the experimental section (ext-chromatic-lf).

Harris et al. [21], introduced a lock-free version of KCAS, which is the basis of our implementation along with optimizations from Arbel-Raviv and Brown [2]. *KCAS* facilitates atomic multi-word updates, however it does nothing to simplify the arguments around values read but not updated, for example the path followed during a search.

Timnat et al. [38] introduced a direct evolution of KCAS which they called Multi-Compare Multi-Swap (MCMS) [38]. Their proposal has similar goals to our work, attempting to achieve a “middle-ground” between performance and ease

of implementation. MCMS attempts to simplify the implementation of concurrent data structures by increasing the expressiveness of KCAS, allowing fields to *compared* without being *swapped*. When HTM is available, this algorithm attempts to carry out operations in a transaction as a fast-path, similar to our approach.

One key difference is that, in its slow path (or if HTM is not available), the MCMS algorithm incurs high synchronization costs. Specifically, when using MCMS for search data structures, one would need to include the entire search path in the arguments to MCMS. On the software code path, MCMS would write to every node on the entire search path, including the root, both in updates and in searches. This aborts all concurrent hardware transactions, likely causing cascading aborts on NUMA systems. In turn, this induces a global contention bottleneck at the root of the search tree. Another key difference is that, on machines without hardware TM, we offer high performance, whereas MCMS essentially becomes the HFP KCAS algorithm. We present a comparison of MCMS and PathCAS in the experimental section.

The work of Mahreshanski et al. [32] analyzes interplay between HTM and other concurrent designs and how they function together. This work suggests that HTM does not obviate other designs, but can be used to improve them. We apply a similar technique in PathCAS, using HTM as a *fast path* for operations, and falling back to our software algorithm when transactions fail a certain number of times.

Guerraoui and Trigonakis present *Optik* [20], which is a methodology for implementing concurrent data structures using optimistic concurrency and versioned locks. PathCAS is built using similar low-level techniques, and encapsulates their complexities in an expressive primitive.

Herlihy and Moss [24] proposed HTM to provide flexible hardware support for non-blocking data structures. Shavit and Touitou introduced software TM [36], as a software-only alternative. PathCAS has similarities to software transactional memory (STM). STM is easier to use than HTM, as STM does not have the same space limitations, but it suffers from various overheads, such as requiring locks per word, dependency on dynamic data structures, and function call overheads on reads and writes.

Kumar et al. [28] introduced *Hybrid Transactional Memory* (HyTM), which is a combination of HTM and STM. Our experiments include results from state-of-the-art HyTM algorithms. To accelerate TM, other work has attempted to break up transactions to avoid overheads. One example is the speculation-friendly tree [19], which uses ElasticTM [18]. However, this tree has relatively poor performance compared to the state of the art, as we show in Section 5.

3 PathCAS

We provide an overview of the PathCAS primitive, and show how it can be used to implement a concurrent data structure. Data structures implemented with PathCAS should

be *node-based*. (A data structure can have many different *types* of nodes, and PathCAS can be used to modify any or all of them.) PathCAS combines ideas from KCAS and version based validation; the rest of this section will provide a description of these components and how they interact.

3.1 Background

In essence, PathCAS is a generalization of KCAS with additional capabilities. KCAS is semantically similar to compare-and-swap (CAS), with the key difference that it is able to atomically change **multiple** addresses (which do *not* have to be contiguous). KCAS supports a single operation in the form of: $KCAS(addr_1, oldValue_1, newValue_1, \dots, addr_k, oldValue_k, newValue_k)$. KCAS does the following atomically: if $addr_i$ contains $oldValue_i$ for all i , the value stored at $addr_i$ is changed to $newValue_i$ for all i and returns true. If not, false is returned.

Our implementation of PathCAS builds on the lock-free KCAS implementation of Harris, Fraser and Pratt (HFP) [21].

Harris, Fraser and Pratt (HFP) algorithm. A KCAS operation first creates a *KCAS descriptor* D that contains the arguments to the KCAS as well as a *status* word that indicates whether the KCAS is *InProgress*, *Succeeded* or *Failed*. It then performs a sequence of atomic double-compare single-swap (DCSS) operations to change all addresses from their respective old values to point to the KCAS descriptor D **only if** the *status* is still *InProgress*. (DCSS atomically determines whether *two* (potentially non-contiguous) addresses contain their respective old values, and if so, changes one to a new value and returns *true*. Otherwise it returns *false*.)

If all of the DCSS operations are successful, then the *status* is changed to *Succeeded* and the addresses are all changed from the descriptor pointer to their respective new values using CAS. Otherwise, the *status* is changed to *Failed* and the addresses are changed back to their old values. This atomic change to the *status* field decides the outcome of the operation (and dictates the behaviour of *helper* threads).

Since old values are replaced by descriptor pointers, any time a thread reads an address that could be modified by a KCAS, it must use a special *KCASRead* function that knows how to handle descriptor pointers. In particular, whenever *KCASRead* encounters a KCAS descriptor, it will *help* the corresponding KCAS operation to complete (by performing the same set of DCSSs and CASs that would be performed by the thread that initially invoked the KCAS). The use of DCSS avoids ABA problems that could otherwise be introduced by this lock-free helping [21]. Crucially, DCSS prevents any helper from storing new pointers to the KCAS descriptor once the *status* has become *Succeeded* or *Failed* (preventing helpers from resurrecting completed KCAS operations).

The KCAS descriptor pointer behaves conceptually like a **lock** that grants exclusive access of a field to a KCAS *operation*, rather than to a particular *thread*. Once all addresses contain pointers to a KCAS descriptor, they can only be

changed in accordance with the corresponding KCAS operation. If the KCAS descriptor's *status* is *Succeeded*, then all helpers will try to change addresses to their respective new values. The value contained in a memory address *logically* changes when the *status* of the descriptor changes to *Succeeded*, and a successful KCAS is linearized then. If the *status* is *Failed*, helpers will try to revert addresses to their old values. A failed KCAS is linearized when it saw a value that did not match the address' old value. Changing an address from a descriptor pointer to a value conceptually *unlocks* it.

The authors implemented lock-free DCSS in software, using CAS and DCSS *descriptor* objects to facilitate helping. There is no need to pierce the atomic DCSS abstraction in our work, except to mention that DCSS descriptors exist.

3.2 Semantics

Whereas KCAS takes all of the addresses to be modified (and their respective old and new values) as explicit arguments, the PathCAS interface is closer to transactional memory. In the following, we say a node n has been *visited* (resp., an address $addr$ has been *added*) if there has been an invocation of $visit(n)$ (resp., $add(addr, \dots)$) since the last invocation of $start()$. PathCAS offers operations to:

- $start()$ gathering arguments for a PathCAS,
- $read(addr)$ an address that might be modified via PathCAS,
- $add(addr, old, new)$ an address $addr$ to be changed atomically from old to new ,
- $visit(n)$ a node n , and
- $validate()$ to check whether any visited nodes have changed since they were visited. $validate$ succeeds and returns *true* only if no such change has occurred. To allow for implementations with diverse progress properties, $validate$ can fail and return *false spuriously*.
- $exec()$ performs a KCAS according to the arguments passed to invocations of add since the last $start$. That is, if all *added* addresses contain their respective old values, then $exec$ succeeds, changing all *added* addresses to their respective new values and returning *true*. Otherwise it returns *false*.
- $vexec()$ performs $exec$ **only if** $validate$ would succeed.

Behaviour is undefined if an address is *added* multiple times with conflicting old and new values. If a node is *visited* multiple times (after a particular invocation of $start$), then any changes to it after the *earliest* such visit will cause $exec$ and $validate$ to return *false*.

Note that $start$, add and $exec$ can simply be viewed as syntactic sugar for accumulating arguments to a KCAS operation, and $read$ is essentially *KCASRead*. However, $visit$ and $vexec$ have no direct analogue in KCAS. To emulate the behaviour of $visit(n)$ and $vexec$ using KCAS, one could include *every* address in node n in the arguments to the KCAS, “changing” each address from its *current* value v to v .

3.3 Implementation

At a high level, the algorithm differs from HFP in the following ways: we implement the syntactic sugar described above for incrementally accumulating arguments, and we add a new **validation** phase wherein visited nodes are inspected to determine whether they have changed since they were visited. Validation affects progress in subtle ways.

Basic operations: start, read, add, visit A PathCAS descriptor consists of a *status* field, a sequence of $\langle addr, old, new \rangle$ triples denoted *entries*, and a sequence of $\langle node, ver \rangle$ pairs denoted *path*. A $start()$ operation creates a new descriptor, and we refer to it as the *thread's descriptor* (until $start$ is invoked again). Similarly to *KCASRead*, a $read(addr)$ reads $addr$, and if it sees a pointer to a descriptor, then it *helps* the corresponding $exec$ or $vexec$ to complete (more about helping below), and repeats these steps. If it sees a non-descriptor value, that value is returned. An $add(addr, old, new)$ adds a triple to the thread's descriptor's *entries*.

Version numbers are used to track changes to the data structure's nodes. More specifically, each node is augmented with a version number ver that should be incremented every time the node is changed. The programmer using PathCAS is responsible for ensuring that s/he increments the version numbers of any node n that s/he modifies using PathCAS. This simply entails reading $n.ver$ and invoking $add(node.ver, v, v+1)$ to increment the value v that was read from $n.ver$. We discuss the motivation behind the decision further in Section 3.9.

A $visit(n)$ operation reads the version v of node n using $read$, saves $\langle n.ver, v \rangle$ in the thread's descriptor's *path*, and returns v .² The use of $read$ means $visit(n)$ will help any $exec$ or $vexec$ it encounters that is in the process of modifying n .

vexec An invocation of $vexec$ simply passes the thread's descriptor to a subroutine called *help* and returns the result.

Consider the set S of addresses *added* to the thread's descriptor $desc$ (i.e., the addresses that should be *changed* by this PathCAS operation). An invocation of $help(desc)$ first uses DCSS to change all of the addresses in S from their respective old values to point to the PathCAS descriptor. If any of these DCSSs fail, then all of the addresses that *were* changed to point to the PathCAS descriptor are reverted to their old values using CAS. Otherwise, now that all addresses are conceptually locked for this PathCAS operation, we can start *validation*. The two **red** lines of code in Algorithm 1 are the only changes from the HFP KCAS algorithm.

Validation To perform validation, $help$ invokes a subroutine called $validateDesc(desc)$, which rereads the version number of each *visited* node and checks whether it has changed, We

²Since the read-set (i.e., *path*) is bounded, we should mention what happens if the read-set size is exceeded. In our code, exceeding the read set size triggers an assertion. In practice, we imagine that the programmer will either over-allocate a large array for visited addresses, or will implement data structures for which a practical height bound is known.

Algorithm 1 PathCAS::help(desc)

```

1: ▶ Phase 1: “lock” addresses for this PathCAS
2: if desc.state == Undecided then
3:   newState = Succeeded
4:   for each (addr, old, new) in desc.entries do
5:     retry_dcsc:
6:     valueSeen = DCSS(⟨addr, old, desc⟩, ⟨desc.state, Undecided⟩)
7:     if isDescriptor(valueSeen) then
8:       help(valueSeen) ▶ DCSS failed because of other PathCAS
9:       goto retry_dcsc ▶ retry after helping
10:    else if valueSeen ≠ old then
11:      newState = Failed ▶ DCSS failed because old ∉ addr
12:      break ▶ Stop trying to “lock” addresses
13:    end for
14:   if newState == Succeeded and not validateDesc(desc) then
15:     newState = Failed ▶ If validation fails, fail and release “locks”
16:   CAS(desc.state, Undecided, newState)
17: ▶ Phase 2: “unlock” addresses to new or old values according to state
18: result = (desc.state == Succeeded)
19: for each (addr, old, new) in desc.entries do
20:   CAS(addr, desc, (result ? new : old))
21: end for
22: return result

```

Algorithm 2 PathCAS::validateDesc(desc)

```

1: for each (node, visitVer) in desc.path do
2:   currentVer = node.ver
3:   if currentVer = desc then
4:     continue ▶ “locked” for our PathCAS
5:   if isDescriptor(currentVer) and currentVer ≠ desc then
6:     return false ▶ “locked” for a different PathCAS
7:   if currentVer ≠ visitVer or (visitVer & 1) then
8:     return false ▶ node’s version has been changed or marked
9:   end for
10: return true

```

discuss the practical considerations of using version numbers, namely wrapping, in the full version of the paper.

To simplify and optimize the implementation of data structures that *mark* nodes when removing them, we steal the least-significant bit from each node’s version number to indicate whether the node has been marked. (In data structures without marking, this bit is simply not used.) Validation succeeds only if all visited nodes are unmarked. In a data structure that marks nodes, success implies that no visited node has been deleted. Storing the marked bit in the same word as the version number allows a node to be marked as deleted at the same time as its version number is updated with minimal overhead. (Note that *visit* returns the mark along with the version number.)

If validation succeeds, none of the visited nodes have changed (or been deleted, in a data structure with marking) since they were *visited*. In this case, the PathCAS descriptor’s *status* field is changed from *InProgress* to *Succeeded* using CAS. Otherwise, it is changed from *InProgress* to *Failed* using CAS. Once the *status* field changes to either *Succeeded* or

Failed it cannot change again. Finally, if the *status* is *Succeeded*, the addresses in *S* are changed to their new values using CAS. Otherwise, their old values are restored via CAS.

Helping As in the related KCAS algorithms, since old values are replaced by descriptors, a special *read()* function (analogous to *KCASRead()*) **must be used to read any fields that can ever be modified by PathCAS**. A *read()* function that encounters a descriptor pointer will *help* the corresponding PathCAS operation to complete. Helpers perform the same sequence of steps as the thread that first invoked *vexec* for this PathCAS. Note that the validation phase will be performed by all helpers, and slow helpers may fail validation even if a fast helper succeeded. However, a slow helper that fails validation cannot revert addresses to old values, since it will attempt to do so using CAS, and this CAS will fail if the node no longer points to the same PathCAS descriptor (with the same version). Moreover, as long as a node points to the PathCAS descriptor, it cannot cause validation to fail.

Progress and helping At this point, one might wonder why forward progress is guaranteed even though an operation *O* can invoke *read* and begin helping another operation *O'* before *O* has finished invoking *add* on all of its fields: Can this cause *O* and *O'* to abort each other? We note that such helping also occurs the lock-free HFP KCAS (in *KCASRead*). The key observation is: although *O* can help another operation before *O* has finished adding its addresses, the operation *being helped* must have already finished adding all of its own addresses. So, such mutual aborts cannot occur. Progress is discussed in greater detail below.

exec The *exec* operation is just a stripped down version of *vexec* that does not perform validation. It can be implemented simply by removing all pairs for visited nodes from the thread’s descriptor before invoking *help*. The intention of including *exec* in the interface is to allow nodes to be *visited* during a data structure traversal *in case validation will be needed*, and then to decide *not* to validate (reducing overhead) at the end of the traversal.

validate The *validate* operation simply passes the thread’s descriptor *desc* to *validateDesc* and returns the result.

3.4 Correctness and Progress

Correctness The *exec* operation is the same as the linearizable lock-free HFP KCAS algorithm, and is linearized in the same way. In other words, for a successful *exec*, we linearize at the change to the descriptor’s *status* field, and for a failed *exec*, we linearize at the read (of an unexpected, non-descriptor value) that caused the failure. Of course, if a *vexec* is performed but *no nodes were visited*, then *vexec* is the same as *exec*, and is linearized the same way.

The case where a *vexec* is performed after some nodes were *visited* is more nuanced. Recall that many *helper* threads can participate in a single *vexec* operation *O* by invoking

help(desc), where *desc* is *O*'s descriptor. The helpers will collaborate to first “lock” all addresses, then perform validation, then use CAS to change the descriptor's *status*, then “unlock” all addresses. Only one helper will successfully change the descriptor's *status*, and we call that helper the *decider*. Once the *status* field is changed, the behaviour of all helpers is dictated by its contents. Two cases arise.

If no visited node has its version number changed (or marked) between when it was visited and when the decider rereads its version number during validation, then validation succeeds. Given that validation succeeds, *vexec* behaves the same as a successful HFP KCAS (matching the PathCAS semantics). We linearize just before the decider invokes *validate(desc)*, at which point all added addresses are “locked” and no visited node had changed.³

However, if some visited node has its version *changed* (or marked) between when it was visited, and when the decider rereads its version number during validation, then *vexec* will behave like a failed HFP KCAS, restoring old values and returning *false* (matching the PathCAS semantics). We linearize when the value that caused the failure was read.

Progress The progress guarantees for PathCAS are subtle. The *start* and *add* operations are wait-free. The *visit*, *exec* and *vexec* operations only perform a constant number of steps in addition to an invocation of *help*, but *help* can invoke itself recursively. The latter is also true in the HFP KCAS algorithm, and it manages to guarantee *lock-free* progress with an assumption that the addresses passed to KCAS are *sorted*. If we make the same assumption, then it is possible to argue that *visit*, *exec* and *vexec* operations are lock-free.

However, lock-freedom only guarantees that infinitely many operations will *terminate* in an infinite execution—not that any of them will *succeed*. To see why this could be a problem, consider a data structure with two nodes, *A* and *B*. Suppose thread t_1 *visits A* and *adds B* (to change *B*'s value), and thread t_2 *visits B* and *adds A* (to change *A*'s value). If t_1 and t_2 both “lock” their respective *added* nodes, then both perform validation, both will fail validation and “unlock,” *terminating*, and hence satisfying lock-freedom, but perhaps preventing the data structure *using* PathCAS from making progress. The problem here is that both *vexec* operations can fail **spuriously**, even though the non-descriptor values that are *semantically* contained in *A* and *B* have not changed.

3.5 Avoiding spurious failures

It is impossible to avoid *vexec* failures altogether. One can always invoke *vexec* after *adding* addresses with unreasonable old values that they have *never contained*. However, the above implementation allows every *vexec* to fail *spuriously*,

³Just as in the HFT KCAS algorithm, at this linearization point, since all added addresses are “locked,” and threads cannot read their values without first helping, no thread can read one of the added addresses and obtain an old value. Instead, a new value will be obtained (after helping).

simply because a *visited* node contained a descriptor pointer. To be able to implement lock-free data structures using PathCAS, we need to change this. Without loss of generality, in the following, we focus on *vexec* operations (since *exec* operations are just a special case).

We say a thread t invokes a **reasonable** *add(addr, old, new)* if the *old* value was read from *addr* at some point *since* the last invocation S_t of *start* by t . If a thread invokes *start* followed by a sequence of *reasonable add* operations, followed by a *vexec*, then we call the *vexec* *reasonable*. With a small modification to *vexec*, we can guarantee the following.

Property P1. If each thread t invokes only *reasonable vexec* operations, then whenever a *vexec* V_t fails, another *vexec* has *succeeded* since V_t 's *start* operation, S_t .

Strong vexec In the implementation described previously, a *vexec* fails *validation* simply because it sees a descriptor, and “unlocks” all of its nodes. Rather than failing spuriously, *vexec* can fall back to a slower lock-free code path on which it creates a *new copy* of its descriptor with slightly different contents. This new descriptor contains all of the *added* fields of the old one, but crucially, all of the *visited* $\langle node, ver \rangle$ pairs in the old descriptor are converted into *added* $\langle node.ver, ver, ver \rangle$ triples. These triples are then *sorted*. Finally, this new descriptor is passed as the argument to an *exec* operation, which will effectively “lock” all of the visited nodes' version numbers rather than simply validating them.

In practice, to reduce overhead, before switching to this slow path, *vexec* can repeatedly try again (a bounded number of times) using an exact copy of its descriptor and performing validation as usual. Since the slow path has high overhead, the number of retries can be tuned to avoid invoking the slow path except where it is really necessary. One can also try contention management strategies such as bounded exponential backoff to further reduce slow path usage.

Note that the choice of *vexec* or strong *vexec* **does not affect performance** in our experiments, as spurious failures are sufficiently infrequent that there is no need to switch to the slow path.

How strong vexec helps Strong *vexec* is not vulnerable to the progress problem described above. Suppose thread t_1 *visits A* and *adds B*, and thread t_2 *visits B* and *adds A*. If t_1 and t_2 both “lock” their respective *added* nodes, then both perform validation, both will fail validation and “unlock,” *but they will not terminate*. Rather, they will retry. They can retry only a bounded number of times before executing the slow path. Once both are executing the slow path, they will each try to lock *A then B* (because of address sorting), and one of them will succeed.

Let us sketch why P1 is satisfied. A *reasonable vexec* V_t does not fail when it fails validation. Rather, it fails only if (a) one of its *reasonable added* addresses contains an unexpected non-descriptor value, or (b) one of its visited nodes' version numbers has been incremented. (In both cases, V_t might help

one or more other *vexec* operations to complete before it can read a non-descriptor value.) In case (a), since the *added* address contained its reasonable old value at some time since S_t , and it can be changed to a different non-descriptor value only by a successful *vexec*, P1 holds. Similarly, in case (b), the visited node’s version number was read since S_t , and a node’s version number is incremented only when the node is changed by a successful *vexec*, so P1 holds.

3.6 Optimizing descriptor management

Arbel-Raviv and Brown [2] showed how to transform the HFP algorithm to *eliminate* the need to allocate and free descriptors for DCSS and KCAS. The same transformation can be applied to PathCAS, allowing each thread to reuse one PathCAS descriptor (and *we do this* in our experiments). The transformation in [2] is straightforward and mechanical, but it makes the pseudocode much more difficult to read, so we presented *pre-transformation* code. Similarly, to avoid complicating the code, we treated DCSS as an atomic primitive. (In reality it is implemented in software as in [2].)

3.7 Optimizing with hardware TM

On systems with support for hardware TM, the PathCAS algorithm above can be used as a *fallback* code path, and a faster hardware TM based algorithm can be used as a *fast path*. In other words, we can use a hardware transaction to perform *vexec/exec* atomically without the overhead of synchronizing via DCSS and CAS.

Our hardware TM based fast path is simply obtained by taking the software algorithm above, wrapping it in a transaction, and then performing a sequence of sequential optimizations (which do not affect correctness because of the atomicity of hardware transactions).

3.8 Comparison to transactional memory

PathCAS is most similar to a lock-free, non-opaque, bounded, object-based TM that is compiled directly into the data structure (rather than being compiled as a library). Such a highly restricted TM implementation could avoid many of the same traditional TM overheads that we also avoid: incremental validation to guarantee opacity, locks per word (instead of version numbers per node), dynamic data structures such as hash tables with intrusive lists (instead of a simple array for our visited nodes), and function call overhead for reads and writes. However, such a TM would be no easier to use than PathCAS, and to our knowledge no such TM exists. Moreover, it would be a substantial undertaking to design an efficient TM with these properties.

3.9 Design Decision: Manual Version Numbers

We contemplated building the incrementing of version numbers into the abstraction, so that it would be automatic. However, we decided that requiring addresses passed to *add* to be fields of nodes might be overly restrictive. We do not

want to rule out applications wherein PathCAS could be used to atomically validate a set of nodes, and also modify arbitrary fields that do not belong to a data structure node (such as a size variable). Therefore, we only require nodes that are **passed to visit** to have version numbers to track changes, and leave it to the programmer to manage them. Note that our interface supports debugging mechanisms to catch errors in managing version numbers.⁴ In an application where it is acceptable to restrict PathCAS so that it only accesses *nodes*, one could easily change *add()* to also take a node pointer in addition to the field pointer, and automate version increments.

4 Application: Lock-free Internal BST

In this section we provide a concrete example of how to create a data structure using PathCAS, namely, a concurrent set implemented as an *internal* binary search tree.

Operations The tree supports the following operations. *contains(key)* returns *true* if *key* is in the tree, and *false* otherwise. *insert(key, val)* returns *false* if *key* is in the tree. Otherwise, it inserts *key* and *value* returns *true*. *delete(key)* returns *false* if *key* is not in the tree. Otherwise, it deletes *key* and its associated *value* and returns *true*.

Data structures Tree nodes have fields for *left* and *right* children, a *key*, a *value*, and a version number *ver* as required by PathCAS.

To avoid special cases, the tree always contains two *sentinel* nodes with keys $-\infty$ and $+\infty$. Consequently, every node with key $k \in (-\infty, +\infty)$ always has both *predecessor* and *successor* nodes. The sentinel with key $+\infty$, which we call the *maxRoot*, is the root of the entire tree. The sentinel with key $-\infty$, which we call the *minRoot*, is the left child of *maxRoot*. No field of *maxRoot* is ever changed. All keys in $(-\infty, +\infty)$ are always found in the *right subtree* of *minRoot*.

Implicit read() Our pseudocode exemplifies a feature of our PathCAS implementation in C++: implicit *read* invocations. Whereas *KCASRead()* calls must be explicitly added by the programmer, in C++, templates and operator overloading can be used to invoke PathCAS *read()* calls automatically.⁵ Thus, in our BST pseudocode we do not explicitly invoke the PathCAS *read* function, but the reader should note: any field that is ever modified by PathCAS is accessed using *read*.

⁴For example, *visit* can save the address ranges of all visited nodes, and *exec* can then check for intersections between the *visited* nodes and *added* addresses that do not have a corresponding *node.ver* increment. This introduces overhead, but can be enabled only when debugging.

⁵The programmer need only annotate the *types* of fields that can be modified by KCAS in the data structure *node* type definition, by *wrapping* each field’s type in a special PathCAS template type. For example, `int key` becomes `casword<int> key`. This requires very little effort, and can even help us catch some types of programmer errors, such as unsafe writes to fields that can be modified with PathCAS.

Algorithm 3 BST::search(*key*)

```

1: while true do
2:   parent = maxRoot
3:   parentVer = visit(parent)
4:   curr = minRoot
5:   currVer = visit(curr)
6:   while curr ≠ NIL do
7:     currKey = curr.key
8:     if key == currKey then
9:       return ⟨true, curr, currVer, parent, parentVer⟩
10:    parent = curr
11:    if key > currKey then curr = curr.right
12:    else key < currKey curr = curr.left
13:    parentVer = currVer
14:    currVer = visit(curr)
15:  return ⟨false, curr, currVer, parent, parentVer⟩

```




Figure 2. Error in *contains* operation without validation

Search The *search(key)* procedure (Algorithm 3) is invoked by *contains*, *insert* and *delete*. It performs a traditional BST search until it encounters a *NIL* pointer, or finds a node containing *key*. *search* returns a tuple of five items with types: $\langle \text{Boolean}, \text{node}, \text{version}, \text{node}, \text{version} \rangle$. If the key *key* is found, this tuple contains *true*, followed by the node that contains *key* and its version number (observed during *search*), followed by its parent and its version number. If *key* was not found, then *search* returns *false*, followed by the final node it encountered (before seeing a *NIL* pointer) and the version number of that node. The remaining two fields are ignored in this case. We return these two nodes (and their versions) to be used by *insert* and *delete*. The key difference between this search and a *sequential* BST search is that each node is passed to an invocation of *visit*.

Contains The *contains(key)* operation invokes *search(key)*, followed by *validate()*. If validation succeeds, then the entire search was effectively *atomic* (since the entire path was *visited*), so we return *true* if *search* found *key* and *false* otherwise. If validation fails, we retry the *contains* from scratch.

One might wonder why *contains* performs validation. Figure 2 depicts an error that can occur without validation in an internal BST with atomic updates. In that execution, *contains(50)* reaches node 60 then sleeps. Then, *delete(40)* atomically deletes 40, promoting the successor key 50 in its place. When *contains(50)* wakes up and continues its search, it will conclude that 50 is not in the tree and return *false*. This is incorrect, as 50 has been in the tree throughout the entirety of the *contains(50)* operation. Validation would catch

Algorithm 4 BST::insert(*key, val*)

```

1: while true do
2:   start()
3:   ⟨foundKey, -, -, parent, parentVer⟩ = search(key)
4:   if foundKey and validate() then return false;
5:   newLeaf = createNode(key, val)
6:   parentKey = parent.key
7:   ptrToChange = (key < parentKey) ? &parent.left : &parent.right
8:   add(ptrToChange, NIL, newLeaf)
9:   add(&parent.ver, parentVer, parentVer + 2) ▶ Increment version
10:  if vexec() then return true

```

this error, since *delete(40)* changes a node that *contains(50)* has already visited.

Validation makes arguing correctness trivial (validated searches are *atomic*), and only incurs a small amount of overhead. We discuss an optimized implementation that performs less validation in Section 4.1.

Insert The *insert(key, val)* operation (Algorithm 4) first invokes *search(key)* to determine whether *key* is already in the tree, and to locate the *parent* whose child pointer should be changed to insert a new node containing *key* (if *key* is not already in the tree).

If the *search* finds *key*, then *validate* is invoked to determine whether any the nodes visited by *search* changed since they were visited. If validation succeeds, it establishes a time *t* during the *insert* operation when *key* was already in the tree, so *false* is returned (the *insert* is linearized at time *t*).

If *search()* does not find *key*, then a new node containing *key* and *val* is created, and *add* is invoked so that the appropriate child pointer of *parent* will be changed (by a subsequent *vexec*) to point to this new node. Since we are trying to change *parent*, *add* is invoked to cause the parent’s version number to be incremented.

Finally, *vexec* is invoked to (attempt to) atomically change the *added* addresses **only if** none of the visited nodes have changed since they were visited. If it succeeds, we linearize at the *vexec*. Otherwise, we retry the *insert* from scratch.

Delete The *delete(key)* operation (Algorithm 6) first searches for the key to be deleted, similar to *insert*.

If *search* does not find *key*, the path followed in *search* is validated. If this validation is successful, *false* can be returned as a time has been established when the entire path traversed by *search*, which did *not* contain *key*, was atomically contained in the tree. Thus, the tree did not contain *key* at some time during the *delete*, and we can linearize at that time. If validation fails, *delete* is retried from scratch.

If *search* finds *key*, the node *curr* containing *key* is returned, along with its *parent*. *delete()* then checks whether these nodes are marked (and hence deleted already). If either is marked, the *delete* is retried from scratch.

Next, *delete* reads *curr.left* and *curr.right* to determine how many children *curr* has. It does not matter if the number of children is counted incorrectly, for example, because *curr.left*

Algorithm 5 BST::getSuccessor(*start*, *startVer*)

```

1: succP = start
2: succPVer = startVer
3: succ = startNode.right
4: succVer = visit(succ)
5: while true do
6:   next = succ.left
7:   if next == NIL then return (succ, succVer, succP, succPVer)
8:   succP = succ
9:   succPVer = succVer
10:  succ = next
11:  succVer = visit(next)

```

is changed between these two reads. If *curr* changes, our subsequent *vexec* will fail and the *delete* will retry. We would not have to consider this possibility at all if we were using opaque transactional memory instead of PathCAS, but we would argue that this reasoning is not onerous to avoid the overheads that come along with opacity.

As in a *sequential* internal BST, three cases arise. In each case, we use *vexec* to perform the sequential update *atomically*. If *vexec* succeeds, *delete* returns *true*, and we linearize at the *vexec*. If *vexec* fails, the *delete* is retried from scratch.

Leaf deletion If *curr* has no children, *vexec* is invoked to unlink and mark it, and to increment the *parent.ver*.

One child deletion If *curr* only has a single child, *vexec* is invoked to replace *curr* by its child, marking *curr* and incrementing *parent.ver*.

Two child deletion If *curr* has two children, we will try to replace its key and value with those of its *successor succ*, then delete the node *succ* (exactly as one does in a sequential internal BST). We first locate the successor *succ* and its parent *succP* using *getSuccessor*, which *visits* each node it traverses and returns the version numbers it saw. Note that the successor cannot have a left child (or else it is not the successor). So, *succ* has at most one child, which means it can be deleted using one of the previous two cases.

If it has a child, *succR*, then we change the appropriate pointer in the parent *succP* to *succR*. If it has no children, *succR* is NIL, so changing the appropriate pointer in *succP* to *succR* simply unlinks *succR*. We *mark succR* since it is being removed, and *increment* the versions of *curr* and *succP* since they are being changed. (If the successor happens to be the right child of *curr*, then *succP* and *curr* are the same node, so we only need to increment one of *succP* and *curr*.) Note that the success of *vexec* implies that *succ* actually is the successor of *curr* when the *delete* is linearized.

4.1 Optimizing to reduce validation

In *contains*, if *foundKey* is *true*, then it is unnecessary to *validate*, because the key can only be found if it was actually in the tree at some time during the *contains*, and we can linearize the *contains* at that time. (If a node was unlinked *before contains began*, then *contains* cannot reach it.)

Algorithm 6 BST::delete(*key*)

```

1: while true do
2:   start()
3:   (foundKey, curr, currVer, parent, parentVer) = search(key)
4:   if not foundKey then
5:     if validate() then return false;
6:     else continue
7:   if currVer & 1 or parentVer & 1 then continue    ▶ if marked
8:   currLeft = curr.left
9:   currRight = curr.right
10:  if currLeft == NIL and currRight == NIL then    ▶ Leaf deletion
11:    ptrToChange = (curr == parent.left) ? &parent.left : &parent.right
12:    add(ptrToChange, curr, NIL)
13:    add(&parent.ver, parentVer, parentVer + 2)
14:    add(&curr.ver, currVer, currVer + 1)          ▶ mark curr
15:    if vexec() then return true
16:  else if currLeft == NIL or currRight == NIL then ▶ One child
17:    childToKeep = (currLeft == NIL) ? currRight : currLeft
18:    ptrToChange = (curr == parent.left) ? &parent.left : &parent.right
19:    add(ptrToChange, curr, childToKeep)
20:    add(&parent.ver, parentVer, parentVer + 2)
21:    add(&curr.ver, currVer, currVer + 1)          ▶ mark curr
22:    if vexec() then return true
23:  else                                            ▶ Two child deletion
24:    (succ, succVer, succP, succPVer) = getSuccessor(curr, currVer)
25:    if succ == NIL or (succVer & 1) or (succPVer & 1) then continue
26:    succR = succ.right                            ▶ succ has at most one child
27:    if succR != NIL then                          ▶ if it does
28:      succRVer = visit(succR)
29:      if succRVer & 1 then continue
30:    ptrToChange = (succP.right == succ) ? &succP.right : &succP.left
31:    add(ptrToChange, succ, succR)
32:    add(&curr.val, curr.val, succ.val)
33:    add(&curr.key, key, succ.key)
34:    add(&succ.ver, succVer, succVer + 1)          ▶ mark succ
35:    add(&succP.ver, succPVer, succPVer + 2)
36:    if succP != curr then add(&curr.ver, currVer, currVer + 2)
37:    if vexec() then return true

```

Readers familiar with the lazy linked list [22] might wonder why we do not have to consider the case where a node is *marked* as *logically deleted before* our *contains* began, but not *unlinked* until later. Note that, unlike typical concurrent data structures with marking, where nodes are marked *before* they are removed, in our tree nodes are unlinked and marked in the *same atomic PathCAS* operation. Thus, reachability in our tree is *equivalent* to being unmarked (and hence logically contained in the tree).

This optimization can also be applied to *insert* and *delete* operations that return *false*. Additionally note that validation is not required in the leaf deletion and one child deletion cases, so *exec* can be used instead of *vexec*. A detailed explanation is deferred to the full version of this paper.

4.2 Extension: AVL Trees with PathCAS

As a second example application of PathCAS, we extend our internal BST to perform relaxed AVL tree balancing. Due to lack of space, we only give an overview here. Complete details appear in the full version of the paper.

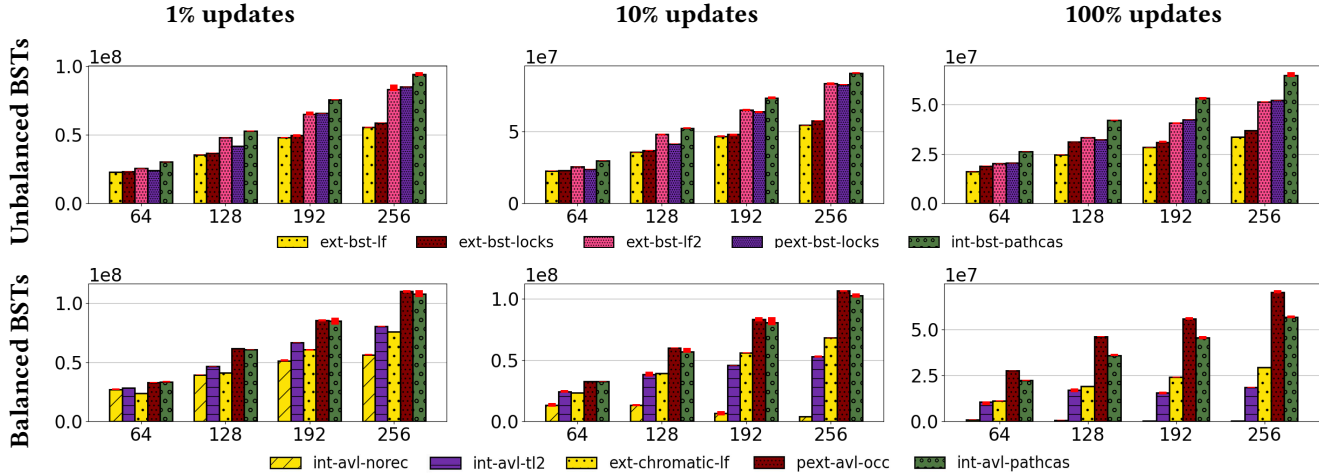


Figure 3. Selected results: AMD system, 10 million keys. x-axis = # of threads. y-axis = operations/sec.

We augment the BST nodes described in the previous section with two new fields: *parent* and *height*. The former points to the node’s parent and the latter contains the *logical height* of the node, which can differ from the node’s *actual height* when the tree is unbalanced.

In *insert*, we initialize newly created nodes’ *parent* pointers to point to the node they are being inserted under. In *delete*, whenever we perform a *vexec* that removes an internal node, and hence changes the *parent* of a child, we *add* that child’s *parent* pointer to the *vexec* as appropriate.

A balance violation exists at node *n* when:

- $n.left.height - n.right.height > 2$; **or**
- $n.left.height - n.right.height < -2$; **or**
- $n.height \neq 1 + \max(n.left.height, n.right.height)$

A violation can be created by any operation that causes node to *gain* or *lose* children. Whenever an operation creates a violation, it performs *rebalancing steps* to fix the violation.

We implement the relaxed AVL tree rebalancing steps of Bougé [4]. Bougé’s proved that, starting from an arbitrarily unbalanced tree, after performing a bounded number of atomic rebalancing steps (wherever they can be applied in the tree, and in any order), the tree will converge to a balanced state. Rebalancing steps are local modifications that affect a small number of nodes, and they do not need to be performed atomically at the same time as a search. The rebalancing steps, namely *rotateLeft*, *rotateRight*, *rotateLeftRight*, *rotateRightLeft* and *fixHeight*, are very similar to the familiar AVL tree rotations. Rebalancing steps eliminate violations, or move them towards the root where they will be eliminated. A tree with no violations is balanced.

Whenever a thread creates a violation at a node, it takes responsibility for repairing that violation, and any subsequent violations it creates while repairing that violation. More specifically, after performing a successful *insert* or *delete*, a thread traverses towards the root, fixing any violations it sees until either: it fixes a violation at the root, it observes a node on the path towards the root that has no violation, or

it encounters a *deleted* node on the path to the root (which means another thread has taken responsibility for any violations further along the path to the root). This is why we augment nodes with parent pointers: they allow us to easily “follow” violations up the tree.

4.3 Freeing data structure nodes

In unmanaged languages like C++, PathCAS manages its own memory, but the programmer must still manually reclaim memory for the *data structures they implement* using PathCAS. Reclaiming nodes that are deleted by a *vexec* (or *exec*) is quite simple using an algorithm such as DEBRA or NBR [6, 37]. The C++ implementation of DEBRA used in [9] offers operations *getGuard()* and *retire(node)*. The former is invoked at the beginning of each data structure operation. The latter can be invoked on any *node* after it is unlinked using *vexec* (or *exec*). This will perform a *delayed free* once no thread has a pointer to *node*. We use DEBRA to reclaim memory for all data structures in our experiments.

Using DEBRA is so mechanical that the necessary invocations of *retire* could even be integrated directly into *vexec* (and *exec*), by having a successful *vexec retire* each node whose *version* it *marks* just before returning.

5 Evaluation

Our experiments follow the methodology of [9], and we use the authors’ publicly available benchmark, *Setbench*. We compare against state-of-the-art hand-crafted trees, as well as several TM-based trees (see Figure 4). We experimented with update rates (1%, 10% and 100%) and uniform key ranges (2×10^5 , 2×10^6 , 2×10^7). Each trial pre-filled the data structure to contain half of the keyrange, then ran for 10 seconds. Data is averaged over six trials with min/max bars shown in red.

Our AMD system has two EPYC 7662 CPUs, each with 64 cores and two hardware threads per core, for a total of 256 hardware threads, and a 256MB shared L3 cache. Threads are *pinned* such that thread counts up to 128 run on one

Unbalanced BSTs		
ext-bst-lf	External, lock-free, CAS	[17]
ext-bst-lf2	External, lock-free, CAS and BTS	[34]
ext-bst-locks	External, ticket locks	[13]
pext-bst-locks	Partially-external, locks, logical ordering	[16]
int-bst-pathcas	Internal, lock-free, PathCAS	this work
int-bst-pathcas+	Internal, lock-free, PathCAS + HTM fast-path	this work
Balanced BSTs		
ext-chromatic-lf	External, LLX/SCX	[8]
pext-avl-occ	Partially-external, locks, logical ordering	[5]
int-avl-pathcas	Internal, lock-free, PathCAS	this work
int-avl-pathcas+	Internal, lock-free, PathCAS + HTM fast-path	this work
Transactional Memory Algorithms		
norec	STM	[12]
hynorec	HTM fast-path + STM slow-path	[11]
rhynorec	HTM fast-path + HTM/STM slow-path	[33]
tle	HTM + Global Lock fallback	this work
tl2	STM	[15]

Figure 4. The list of algorithms in our experiments.

	LLC miss per op	Cycles per op	Instr. per op	Avg. Key Depth	Peak Mem. Usage
ext-bst-lf	26.6	25872	2047	32.9	1137 MiB
ext-bst-locks	24.6	23790	1527	30.5	1038 MiB
ext-bst-lf2	15.3	15226	1739	30	720 MiB
pext-bst-locks	16.3	16184	1011	28.6	2047 MiB
int-bst-pathcas	13.7	12371	2661	28.7	539 MiB
int-avl-norec	141	8677653	789316	21.7	620 MiB
ext-chromatic-lf	31.5	27123	3204	24.3	2441 MiB
int-avl-tl2	44.2	55514	7935	21.7	3642 MiB
pext-avl-occ	11.6	11234	2398	23.2	844 MiB
int-avl-pathcas	12.9	13752	3976	21.7	717 MiB

Figure 5. Detailed analysis for 100% updates, 256 threads.

socket. We used GCC 10.1.0 -O3, the fast allocator jemalloc, and interleaved memory pages across sockets with *numactl*.

Many more experiments can be found in the full version of this paper, including Intel results and additional algorithms and workloads.

Comparing unbalanced trees The top three plots in Figure 3 present results comparing our unbalanced BST (*int-bst-pathcas*) to a variety of leading hand-crafted unbalanced BST implementations. Our code includes all fixes recommended by Arbel-Raviv et al. for obtaining reliable BST performance results [3]. Our PathCAS BST often significantly outperforms its competitors. We explain why using Figure 5.

As expected, out of the unbalanced BSTs, *int-bst-pathcas* performs the largest number of instructions per tree operation. However, it has the smallest cycle count per operation, suggesting that its instructions can be pipelined more efficiently. Crucially, *int-bst-pathcas* incurs the smallest number of last-level cache misses, because of its low average key depth and peak memory usage. The *ext-bst-** BSTs are *external*: the keys that are semantically contained in the dictionary are stored in the *leaves*, and internal nodes contain dummy *routing* keys. External trees contain more nodes than internal trees and are taller. *pext-bst-locks* is *partially* external, which means it is somewhat closer to an internal tree. Despite its low average key depth, its nodes are quite large, containing multiple locks and many pointers, as they participate in both a tree and a doubly linked list. It underperforms because of its peak memory usage and LLC misses.

Comparing balanced trees The bottom three plots in Figure 3 compare our AVL tree (*int-avl-pathcas*) with other balanced BSTs, including *pext-avl-occ* [5], which is known to be the fastest concurrent BST in many workloads [3]. In read-mostly workloads, *int-avl-pathcas* is competitive with *pext-avl-occ*, and outperforms the other trees.

In the 100% update workload, *int-avl-pathcas* is at most 20% slower than the fastest algorithm. The fact that *int-avl-pathcas* is not *drastically* outperformed by the highly tuned and intricate *pext-avl-occ* tree is remarkable.⁶ The *ext-chromatic-lf* tree, which is implemented using the LLX and SCX primitives, does not fare nearly as well.

TM-based trees It should be noted that in an attempt to be generous to these TM approaches, in our implementations we compiled each TM in the same compilation unit as the data structure (rather than as a linked library), and force-inlined all TM code, eliminating the overhead of function calls to the TM code from the data structure. This optimization would be unrealistic in practice, however should give the TM implementations the best performance possible for comparison. Despite this, the TM based algorithms in Figure 3 still suffer from high instruction counts and LLC miss rates (Figure 5). In particular, the extremely high instruction counts for *int-avl-norec* are due to contention on the global version lock and repeated read set validation to guarantee opacity.

The results in the introduction (Figure 1), from our Intel system, include more algorithms, since the system has hardware transactional memory support. In those results, our trees outperform the next fastest algorithm, TLE, by nearly 2x. Moreover, those results are “generous” to TLE, since its global locking fallback code path degrades performance dramatically in workloads with more updates.

5.1 Comparison with MCMS

To compare PathCAS with MCMS, we extended Timnat’s original C++ code for the MCMS linked list. We note that we found some bugs in his implementation of MCMS, one of which only affects the source code, and one of which affects the algorithm in the MCMS paper. Details are to appear in the full version of this paper. We fixed these bugs, and applied the same lock-free descriptor optimizations that we use, and implemented an internal BST using MCMS to validate the entire search path (similarly to how we validate the entire search path using PathCAS). The implementation is optimized to the best of our ability: for instance, it avoids performing MCMS in cases where searches return true or inserts return false. Moreover, deletes that return true perform their modifications in small MCMS operations that do

⁶The *pext-avl-occ* tree is intricate, and carefully engineered, using sophisticated sequence locks that encode whether ongoing rebalancing operations are *shrinking* or *expanding* the key range, and allowing a key that was marked as deleted to be reinserted simply by changing a bit. This may inflate its performance in the types of workloads used in our experiments.

# threads	100% update			100% search		
	PathCAS	MCMS+	MCMS-	PathCAS	MCMS+	MCMS-
1	2.50	0.99	1.12	3.19	1.25	1.54
2	4.42	0.69	0.70	6.33	0.78	0.77
4	7.94	0.56	0.54	12.44	0.63	0.62
8	13.97	0.50	0.50	24.07	0.55	0.54
48	46.67	0.35	0.31	77.23	0.40	0.31
190	79.71	0.04	0.03	263.00	< 0.1	< 0.1

Figure 6. Intel results for an internal BST implemented with PathCAS, MCMS with HTM (MCMS+), and MCMS without HTM (MCMS-), respectively. Tree initially contains 100,000 keys. Results are in millions of operations per second.

# threads:	24	48	96	144	190
<i>ext-bst-elastic</i>	17.24	29.54	52.25	71.59	90.43
<i>ext-bst-lf2</i>	29.28	53.07	129.02	195.71	267.44

Figure 7. Intel results for elastic transactions with 1% updates and 99% searches. The trees initially contain 10,000,000 keys. Results are in millions of operations per second.

not include the search path. The hardware TM code path in our MCMS implementation is faithful to the MCMS paper in that it does not write to nodes on the search path unless the operation is an update that intends to modify that node.

Results appear in Figure 6. We find that the resulting tree is orders of magnitude slower than our PathCAS tree. For example, for 100% searches and 100,000 keys, on a system with hardware TM, the throughput of the MCMS tree is 0.4M at 48 threads, whereas the throughput of PathCAS is 77M. According to `perf stat -e tx_commit,tx_abort`, half of MCMS transactions abort at 8 threads, 84% abort at 48 threads, and there are many capacity aborts even with a single thread. As we suspected, even transient aborts in MCMS can quickly turn into cascading aborts and software path executions. At 190 threads, PathCAS is thousands of times faster than MCMS.

5.2 Comparison with Elastic TM

We also compared PathCAS with *elastic transactions*, which use sophisticated synchronization techniques to split transactions into smaller atomic pieces to improve performance. Specifically, we compared with *ext-bst-elastic*, the “speculation-friendly BST” from Synchronobench [19] that is built using elastic transactions. We were unable to port the Elastic STM library that *ext-bst-elastic* depends on into Setbench, as it would require overhauling Setbench to support the background rebalancing threads used by *ext-bst-elastic*, and would require us to completely redesign its memory reclamation.

So, we obtained performance numbers for *ext-bst-elastic* using Synchronobench, instead of Setbench. We also obtained performance numbers for *ext-bst-lf2* on the same workload, as it is included in Synchronobench as well as setbench. Results appear in Figure 7. This gives us a sort of limited point of comparison between our results in Setbench and our results in Synchronobench. Although the results do not allow for a rigorous comparison between *ext-bst-elastic* and the other trees in our experiments, we note that it is *much* slower

than *ext-bst-lf2*, which is in the middle of the pack in our Setbench experiments. And, we also note that *ext-bst-elastic* is being evaluated in an environment that is presumably most favourable to it, as *ext-bst-elastic* was developed by the authors of Synchronobench, and integrated therein by the authors. Additionally observe that this 1% update workload is quite favourable to *ext-bst-elastic*, as its performance degrades faster than *ext-bst-lf2* as the update rate increases.

6 Conclusion

This paper introduced PathCAS, a primitive used to implement efficient concurrent data structures while maintaining lower complexity than hand-crafted techniques. PathCAS utilizes an HTM fast path combined with an efficient fallback path that relies on KCAS, version numbers, and search path validation.

We implemented a set of historically difficult data structures using PathCAS, including an internal balanced tree, and have shown them to achieve competitive performance with the best fine-grained variants. We compared the performance of our data structures with both TM-based and hand-crafted variants of these structures, showing that PathCAS surpasses TM based approaches, and can rival the state-of-the-art in hand-crafted designs.

While in this work we focus on the two trees presented, we emphasize that one can use PathCAS in a direct way to implement many data structures wherein an operation consists of a read phase followed by a write phase. The construction is similar to our trees: *visit* each node that will be read or written, then *add* and *exec* the necessary modifications. Using this approach, we were able to implement the following: (a,b)-trees, lists, hash-tables, hash-lists, stacks, queues, skip-lists and dynamic graph connectivity data structures. We chose to focus on trees in the paper, since they have traditionally been difficult to implement and prove correct. However, we plan to share more of our implementations as part of future work.

Acknowledgments

This work was supported by: the Natural Sciences and Engineering Research Council of Canada (NSERC) Collaborative Research and Development grant: CRDPJ 539431-19, the Canada Foundation for Innovation John R. Evans Leaders Fund with equal support from the Ontario Research Fund CFI Leaders Opportunity Fund: 38512, Waterloo Huawei Joint Innovation Lab project “Scalable Infrastructure for Next Generation Data Management Systems”, NSERC Discovery Launch Supplement: DGEGR-2019-00048, NSERC Discovery Program under the grants: RGPIN-2019-04227 and RGPIN-04512-2018, and the University of Waterloo. We would also like to thank the reviewers for their insightful comments.

References

- [1] James H. Anderson and Mark Moir. 1995. Universal Constructions for Multi-Object Operations. In *Proceedings of the Fourteenth Annual ACM Symposium on Principles of Distributed Computing* (Ottawa, Ontario, Canada) (PODC '95). Association for Computing Machinery, New York, NY, USA, 184–193. <https://doi.org/10.1145/224964.224985>
- [2] Maya Arbel-Raviv and Trevor Brown. 2017. Reuse, don't Recycle: Transforming Lock-free Algorithms that Throw Away Descriptors. In *Proceedings of the 31st ACM Symposium on Distributed Computing*.
- [3] Maya Arbel-Raviv, Trevor Brown, and Adam Morrison. 2018. Getting to the Root of Concurrent Binary Search Tree Performance. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*. USENIX Association, Boston, MA, 295–306. <https://www.usenix.org/conference/atc18/presentation/arbel-raviv>
- [4] Luc Bougé, Joaquim Gabarró Vallés, Xavier Messeguer Peypoch, and Nicolas Schabanel. 1998. Height-relaxed AVL rebalancing: a unified, fine-grained approach to concurrent dictionaries.
- [5] Nathan G. Bronson, Jared Casper, Hassan Chafi, and Kunle Olukotun. 2010. A Practical Concurrent Binary Search Tree. In *Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Bangalore, India) (PPOPP '10). ACM, New York, NY, USA, 257–268. <https://doi.org/10.1145/1693453.1693488>
- [6] Trevor Brown. 2015. Reclaiming Memory for Lock-Free Data Structures: There Has to Be a Better Way. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing* (Donostia-San Sebastián, Spain) (PODC '15). ACM, New York, NY, USA, 261–270. <https://doi.org/10.1145/2767386.2767436>
- [7] Trevor Brown, Faith Ellen, and Eric Ruppert. 2013. Pragmatic primitives for non-blocking data structures. In *ACM Symposium on Principles of Distributed Computing, PODC '13, Montreal, QC, Canada, July 22-24, 2013*. 13–22. <https://doi.org/10.1145/2484239.2484273>
- [8] Trevor Brown, Faith Ellen, and Eric Ruppert. 2014. A General Technique for Non-blocking Trees. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Orlando, Florida, USA) (PPOPP '14). ACM, New York, NY, USA, 329–342. <https://doi.org/10.1145/2555243.2555267>
- [9] Trevor Brown, Aleksandar Prokopec, and Dan Alistarh. 2020. Non-Blocking Interpolation Search Trees with Doubly-Logarithmic Running Time. In *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (San Diego, California) (PPOPP '20). Association for Computing Machinery, New York, NY, USA, 276–291. <https://doi.org/10.1145/3332466.3374542>
- [10] Jaewoong Chung, Luke Yen, Stephan Diestelhorst, Martin Pohlack, Michael Hohmuth, David Christie, and Dan Grossman. 2010. ASF: AMD64 Extension for Lock-Free Data Structures and Transactional Memory. In *Proceedings of the 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '43)*. IEEE Computer Society, USA, 39–50. <https://doi.org/10.1109/MICRO.2010.40>
- [11] Luke Dalessandro, François Carouge, Sean White, Yossi Lev, Mark Moir, Michael L. Scott, and Michael F. Spear. 2011. Hybrid NOrec: A Case Study in the Effectiveness of Best Effort Hardware Transactional Memory. In *Proceedings of the 16th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*. ACM, 39–52.
- [12] Luke Dalessandro, Michael F. Spear, and Michael L. Scott. 2010. NOrec: streamlining STM by abolishing ownership records. In *Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2010, Bangalore, India, January 9-14, 2010*. 67–78. <https://doi.org/10.1145/1693453.1693464>
- [13] Tudor David, Rachid Guerraoui, and Vasileios Trigonakis. 2015. Asynchronous 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 (Istanbul, Turkey) (ASPLOS '15)*. Association for Computing Machinery, New York, NY, USA, 631–644. <https://doi.org/10.1145/2694344.2694359>
- [14] Linux Kernel development community. 2021. TSX Async Abort (TAA) mitigation documentation. (2021). https://www.kernel.org/doc/html/latest/x86/tsx_async_abort.html
- [15] Dave Dice, Ori Shalev, and Nir Shavit. 2006. Transactional locking II. In *International Symposium on Distributed Computing*. Springer, 194–208.
- [16] Dana Drachler, Martin Vechev, and Eran Yahav. 2014. Practical concurrent binary search trees via logical ordering. In *PPOPP '14 Proceedings of the 19th ACM SIGPLAN symposium on Principles and practice of parallel programming*. 343–356.
- [17] Faith Ellen, Panagiota Fatourou, Eric Ruppert, and Franck van Breugel. 2010. Non-blocking Binary Search Trees. In *Proceedings of the 29th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (Zurich, Switzerland) (PODC '10)*. ACM, New York, NY, USA, 131–140. <https://doi.org/10.1145/1835698.1835736>
- [18] Pascal Felber, Vincent Gramoli, and Rachid Guerraoui. 2009. Elastic Transactions. In *Distributed Computing, Idit Keidar (Ed.)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 93–107.
- [19] V. Gramoli. 2015. More Than You Ever Wanted to Know about Synchronization: Synchrobench. In *Proceedings of the 20th Annual ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP)*.
- [20] Rachid Guerraoui and Vasileios Trigonakis. 2016. Optimistic Concurrency with OPTIK. In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Barcelona, Spain) (PPOPP '16). ACM, New York, NY, USA, Article 18, 12 pages. <https://doi.org/10.1145/2851141.2851146>
- [21] Timothy L. Harris, Keir Fraser, and Ian A. Pratt. 2002. A Practical Multi-Word Compare-and-Swap Operation. In *Proceedings of the 16th International Conference on Distributed Computing (DISC '02)*. Springer-Verlag, Berlin, Heidelberg, 265–279.
- [22] Steve Heller, Maurice Herlihy, Victor Luchangco, Mark Moir, William N. Scherer III, and Nir Shavit. 2006. A Lazy Concurrent List-Based Set Algorithm. In *Proceedings of the 9th International Conference on Principles of Distributed Systems (OPDIS 2005), Revised Selected Papers (Lecture Notes in Computer Science, Vol. 3974)*, James H. Anderson, Giuseppe Prencipe, and Roger Wattenhofer (Eds.). Springer, 3–16.
- [23] Maurice Herlihy. 1991. Wait-Free Synchronization. *ACM Trans. Program. Lang. Syst.* 13, 1 (Jan. 1991), 124–149. <https://doi.org/10.1145/114005.102808>
- [24] Maurice Herlihy and J. Eliot B. Moss. 1993. Transactional Memory: Architectural Support for Lock-free Data Structures. In *Proceedings of the 20th Annual International Symposium on Computer Architecture (San Diego, California, USA) (ISCA '93)*. ACM, New York, NY, USA, 289–300. <https://doi.org/10.1145/165123.165164>
- [25] Maurice Herlihy and Nir Shavit. 2008. *The art of multiprocessor programming*. Morgan Kaufmann. I–XX, 1–508 pages.
- [26] Shane Howley and Jeremy Jones. 2012. A non-blocking internal binary search tree. *ACM Symposium on Parallelism in Algorithms & Architectures* (06 2012). <https://doi.org/10.1145/2312005.2312036>
- [27] IBM. 2020. Power ISA™ Version 3.1. (2020). https://wiki.raptors.com/w/images/f/f5/PowerISA_public.v3.1.pdf
- [28] Sanjeev Kumar, Michael Chu, Christopher J. Hughes, Partha Kundu, and Anthony Nguyen. 2006. Hybrid Transactional Memory. In *Proceedings of the Eleventh ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (New York, New York, USA) (PPOPP '06). New York, NY, USA, 209–220. <https://doi.org/10.1145/1122971.1123003>
- [29] Michael Larabel. 2021. Intel To Disable TSX By Default On More CPUs With New Microcode. (2021). https://www.phoronix.com/scan.php?page=news_item&px=Intel-TSX-Off-New-Microcode
- [30] Arm Ltd. 2021. ARM C Language Extensions. <https://developer.arm.com/documentation/101028/0013/?lang=en>

- [31] Victor Luchangco, Mark Moir, and Nir Shavit. 2003. Nonblocking K-Compare-Single-Swap. In *Proceedings of the Fifteenth Annual ACM Symposium on Parallel Algorithms and Architectures* (San Diego, California, USA) (SPAA '03). Association for Computing Machinery, New York, NY, USA, 314–323. <https://doi.org/10.1145/777412.777468>
- [32] Darko Makreshanski, Justin Levandoski, and Ryan Stutsman. 2015. To Lock, Swap, or Elide: On the Interplay of Hardware Transactional Memory and Lock-Free Indexing. *Proc. VLDB Endow.* 8, 11 (jul 2015), 1298–1309. <https://doi.org/10.14778/2809974.2809990>
- [33] Alexander Matveev and Nir Shavit. 2015. Reduced Hardware NOrec: A Safe and Scalable Hybrid Transactional Memory. *SIGPLAN Not.* 50, 4 (March 2015), 59–71. <https://doi.org/10.1145/2775054.2694393>
- [34] Aravind Natarajan and Neeraj Mittal. 2014. Fast Concurrent Lock-free Binary Search Trees. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Orlando, Florida, USA) (PPoPP '14). ACM, New York, NY, USA, 317–328. <https://doi.org/10.1145/2555243.2555256>
- [35] Arunmozhi Ramachandran and Neeraj Mittal. 2015. A Fast Lock-Free Internal Binary Search Tree. In *Proceedings of the 2015 International Conference on Distributed Computing and Networking* (Goa, India) (ICDCN '15). Association for Computing Machinery, New York, NY, USA, Article 37, 10 pages. <https://doi.org/10.1145/2684464.2684472>
- [36] Nir Shavit and Dan Touitou. 1995. Software Transactional Memory. In *Proceedings of the Fourteenth Annual ACM Symposium on Principles of Distributed Computing* (Ottawa, Ontario, Canada) (PODC '95). Association for Computing Machinery, New York, NY, USA, 204–213. <https://doi.org/10.1145/224964.224987>
- [37] Ajay Singh, Trevor Brown, and Ali Mashtizadeh. 2021. NBR: Neutralization Based Reclamation. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (Virtual Event, Republic of Korea) (PPoPP '21). Association for Computing Machinery, New York, NY, USA, 175–190. <https://doi.org/10.1145/3437801.3441625>
- [38] Shahar Timnat, Maurice Herlihy, and Erez Petrank. 2015. A Practical Transactional Memory Interface. In *Euro-Par 2015: Parallel Processing*, Jesper Larsson Träff, Sascha Hunold, and Francesco Versaci (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 387–401.
- [39] R. K. Treiber. 1986. *Systems programming: Coping with parallelism*. Technical Report RJ 5118. IBM Almaden Research Center.

7 Artifact

We have provided a docker container in which you can download and start running the experimental benchmark.

- The docker container was prepared on Ubuntu 20.04, which is also the OS the actual container is running
- In particular, we build our container using Docker version 20.10.2, build 20.10.2-0ubuntu1 20.04.2.
- Our experiments are run on a machines with 384 GB of RAM and 192 threads total, while possible to run these experiments on a different machines with less memory or threads:
 - If you have less memory, you may need to use smaller data structure sizes
 - If you have fewer threads, you will simply need to use less threads

Inside the docker container, in the `experiments` folder, you will find a script called `run.sh` that builds all data structures and runs a simple test (using each data structure). This script is driven by the experimental configuration described

in `experiments/_common.py`. By default, the contents of `experiments/_common.py` cause `run.sh` to reproduce Figure 3. More details on how to modify this experimental configuration appear below.

Note that `run.sh` performs experiments for both transactional memory and non-transactional memory data structures. This can take quite some time, so you may want to limit the number of trials in `_common.py` to a relatively small number.

7.1 Step By Step Instructions

1. Install docker (ideally the same version we use) on your system with your preferred package manager
2. Download the docker image from Zenodo (<https://zenodo.org/record/5728166>). For example, on Linux you might do this by executing:


```
$ wget https://zenodo.org/record/5728166/files/ppopp2022pathcas.tar.gz
```
3. Add the image to your local Docker images (running in the same directory as step 2):


```
$ docker load -i ppopp2022pathcas.tar.gz
```
4. Launch the docker container


```
$ docker run -i -t --privileged ppopp2022pathcas /bin/bash
```

Note that `privileged` is *required* in order to ascertain the proper thread pinning strategy for the experiments, and to record performance counters for, e.g., cache statistics. You may also need to set `kernel.perf_event_paranoid` to `-1` on Linux.

5. Go to the experiments folder


```
$ cd experiments
```
6. The experiments can be configured in `_common.py`. Be sure to edit the `HOST CONFIGURATION` and `EXPERIMENTAL CONFIGURATION` sections of `experiments/_common.py` to match the machine you are running on, and to reflect the types of experiments you would like to run. Note that the various configuration parameters are described in the comments.

By default, `_common.py` is configured to reproduce Figure 3 in the paper. Running such experiments can take many hours (at least 12). **If you would like to run a shorter test to get started**, please uncomment the more restrictive *testing* values of `ins_del_fractions`, `max_keys`, `exp_duration_millis`, `thread_counts` and `num_trials` provided alongside the default values in `_common.py` before proceeding. (Once you have run your small test, you can comment those testing values out again.)

Additionally, Smaller example test values are provided if you wish to run a quick proof of concept test (commented out next to the actual set values).
7. To run the experiments described in `_common.py`, simply execute `./run.sh`. It will run experiments, store the results in a `sqlite3` database, process those results

using various SQL queries, and produce several text data files and plot images (more on this below).

If the *testing* values in `_common.py` are **uncommented**, this should take less than five minutes.

If the *testing* values are **commented out**, then `run.sh` will perform all necessary experiments to reproduce Figure 3 in the paper. **This should take approximately 12 hours.**

While experiments are being performed, the scripts will produce output of the form “step 000001 / 000336. . .”. Each such line contains the command being executed, as well as a very rough *estimate of the remaining running time* for all experiments.

8. After the runs are complete, **PNG format plot images** for the experiments can be found in directories `experiments/data_tm` and `experiments/data_non_tm`. Note that, by default, a large super set of the plots in the paper are generated.

To more easily view these images, you may want to copy them from the docker container to your host machine using `docker cp`. For example, *while the docker container is running*, you can run the following command *on the host machine*, where `<CONTAINER_NAME>` is the name of the running container.⁷

```
$ docker cp <CONTAINER_NAME>:/root/tmbench/experiments .
```

This will copy all experimental results to the host machine. Alternatively, you can browse the results in text format directly inside the docker container:

- a. A detailed summary of the numerical data can be obtained by starting in the `experiments` directory and running `./basic_info.sh` to produce a pretty-printed table, with columns for update rate, data structure name, key range size, thread count and throughput. If you ran both TM and Non-TM experiments, this script will print a table for each type of results.
- b. Note that you can also perform your own arbitrary SQL queries on the `sqlite` database, either by modifying the queries in `basic_info.sh`, or by entering an interactive `sqlite` console in either of the `data_*` directories in `experiments/`:

```
$ sqlite3 output_database.sqlite
```
- c. If you want to view the complete raw text data (which contains considerably more information than the `sqlite` database), you can find files for each experimental “step” stored in `experiments/data_tm` and `experiments/data_non_tm`, with file names of the form: `data0*<STEP>.txt`.

There are also CSV-format tables of data that are converted into various plots in file names of the form:

```
{DATA_STRUCTURE_TYPE}_{METRIC}_u{UPDATE RATES}-k{MAX_KEY}. {FILE_TYPE}
```

For example, `data_tm/bst_tm_total_throughput_sec-u50.0_50.0-k20000000.txt` contains a table of data that is used to produce a throughput comparison plot for TM-based binary search trees, with 50% inserts and 50% deletes, and a key range size of 20 million (i.e., initial data structure size of 10 million). This example is of particular note, as it corresponds to one of the plots in Figure 3.

9. Note that, due to the nature of docker containers, **all data will be lost if you exit the docker run**. If you want to save any of your generated data, you can do so using `docker cp` from a different terminal on the host machine to copy the relevant data from the container to the host.

Alternatively, you can use `docker commit` to save a new version of the docker image that includes all of the data you’ve generated. You can then launch *that* image in a docker container to access your data again.

10. Once you have saved any data you want to keep, you can *exit* the docker container by running `exit` in the container, or `docker stop <CONTAINER_NAME>` on the host.

LLC misses on AMD Note the artifact fails to track LLC misses on recent AMD processors. Instead, they must be tracked manually using `perf stat`.

For example, on our machine, we used a command of the form: `[usual command prefix up until the binary] perf stat -e l3_comb_clstr_state.request_miss [actual binary being run] [args]`.

Also note that will produce a raw number of LLC misses for the entire execution, so one will need to divide by the number of operations completed. This number will be somewhat inflated because it includes experimental setup and tear-down as well as prefilling. To get around this, use `perf record -k CLOCK_MONOTONIC`, which timestamps all cache misses with a clock that is compatible with our benchmark. Then one can take a pair of timestamps emitted by our benchmark, `REALTIME_START_PERF_FORMAT` and `REALTIME_END_PERF_FORMAT`, and plug them into `perf report -ns -time <start>, <stop>` where `<start>` and `<stop>` are our timestamps.

Generating Figure 5 Figure 5 is generated by running the following SQL queries on the results databases (starting in the same directory as `_common.py` and `run.sh`).

```
for exp in tm non_tm ; do
  \
  ../setbench/tools/data_framework/run_experiment.py \
  _exp_${exp}.py -q "select maxkey, TOTAL_THREADS, \
  alg, mem_maxresident_kb, PAPI_L3_TCM, PAPI_TOT_CYC, \
  PAPI_TOT_INS, tree_stats_avgKeyDepth from data \
  where ins_del_frac = '50.0 50.0' \
  order by maxkey desc, total_threads desc, alg" ; done
```

⁷To obtain the name of the running container, run `$ docker container list` and look in the `NAMES` column.