

Towards Automatic Lock Removal for Scalable Synchronization

Maya Arbel^{*1,2}, Guy Golan-Gueta¹, Eshcar Hillel¹, and Idit Keidar^{1,2}

¹ Yahoo Labs, Haifa, Israel

² The Technion, Haifa, Israel

Abstract. We present a *code transformation* for concurrent data structures, which increases their scalability without sacrificing correctness. Our transformation takes lock-based code and replaces some of the locking steps therein with optimistic synchronization in order to reduce contention. The main idea is to have each operation perform an optimistic traversal of the data structure as long as no shared memory locations are updated, and then proceed with pessimistic code. The transformed code inherits essential properties of the original one, including linearizability, serializability, and deadlock freedom.

Our work complements existing pessimistic transformations that make sequential code thread-safe by adding locks. In essence, we provide a way to optimize such transformations by reducing synchronization bottlenecks (for example, locking the root of a tree). The resulting code scales well and significantly outperforms pessimistic approaches. We further compare our synthesized code to state-of-the-art data structures implemented by experts. We find that its performance is comparable to that achieved by the custom-tailored implementations. Our work thus shows the promise that automated approaches bear for overcoming the difficulty involved in manually hand-crafting concurrent data structures.

1 Introduction

The steady increase in the number of cores in today’s computers is driving software developers to allow more and more parallelism. An important focal point for such efforts is scaling the concurrency of shared data structures, which are often a principal friction point among threads. Many recent works have been dedicated to developing scalable concurrent data structures (e.g., [8, 20, 40, 12, 15, 10, 5, 22, 11, 32, 25, 38]), some of which are widely used in real-world systems [44].

Each of these projects generally focuses on a single data structure (e.g., a binary search tree [11] or a queue [38]) and manually optimizes its implementation. These data structures are developed by concurrency experts, typically PhDs or PhD candidates, and proving their correctness is painstaking; for example, the proofs of [10, 22] are 31 and 20 pages long, respectively. The rationale behind

* Maya Arbel is supported in part by the Technion Hasso Platner Institute (HPI) Research School.

dedicating so much effort to one data structure is that it is generic and can be used by many applications. Nevertheless, systems often use data structures in unique ways that necessitate changing or extending their code (e.g., [2, 3, 43, 46]), in which cases custom-tailored implementations may not meet the requirements. Here, we propose an approach to facilitate this labor-intensive process, making scalable synchronization more readily available.

Specifically, we present in Section 2 an algorithm for a source-to-source code transformation that takes a lock-based concurrent data structure implementation as its input and generates more scalable code for the same data structure via judicious use of optimism. Our approach combines optimism and pessimism in a practical way. Like some previous hand-crafted solutions [31], we exploit the common access pattern in data structure operations, (e.g., tree insertion or deletion), which typically begin by traversing the data structure (to the insertion or deletion point), and then perform local updates at that location. Our solution replaces locks in the initial read-only traversal with optimistic synchronization, and performs updates using the original lock-based code. It may thus be seen as a form of software lock elision for read-only prefixes of operations (transactions). Combining optimism and pessimism allows us to achieve “the best of both worlds” – while the optimistic traversal increases concurrency and eliminates bottlenecks, the use of pessimistic updates saves the overhead associated with speculative or deferred shared memory updates, (as occurs in *software transactional memory (STM)* [30]).

In the full version of the paper [9] we show that our transformation preserves the external behavior (e.g., linearizability, serializability, and deadlock-freedom) of the original lock-based code; Moreover, our transformation preserves *disjoint access parallelism* [34], (the property that threads that access disjoint data objects do not contend on low level shared memory locations), as it refrains from introducing a shared global clock (as some STM systems do [45]) or other sources of contention.

One important use case for our transformation is to apply it in conjunction with automatic lock-based parallelization mechanisms [26, 36]. The latter instrument sequential code and add fine-grained locks that ensure its safety in concurrent executions. Our evaluation shows that, by themselves, solutions of this sort may scale poorly. This is due to synchronization bottlenecks, e.g., the root of a tree, which is locked by all operations. By subsequently applying our transformation, one can optimize the lock-based code they produce, yielding an end-to-end approach to scalable parallelization of sequential code.

In Section 3 we evaluate our transformation by generating an unbalanced search tree and a treap (randomized balanced search tree). We synthesize these data structures from *sequential implementations* by applying first the algorithm of [26] (*domination locking*) to create lock-based code, and then our transformation. We evaluate the scalability of the resulting code in a range of workload scenarios on a 32-core machine. In all cases, the lock-based implementations do not scale – their throughput remains flat as the number of running threads increases. In contrast, the code generated by our transformation is scalable, and

its throughput continues to grow with the number of threads. We further use the Synchrobench framework [27] to compare our synthesized code to data structures that were recently hand-crafted by experts in the field [20, 11, 1, 22, 15], as well as a state-of-the-art STM [45]. Our results show that the implementations we have generated perform comparably to custom-tailored solutions.

The advantage of our approach is in its *generic* nature, which allows us to parallelize existing code without requiring experts to perform manual optimizations. Other generic approaches we are familiar with are domination locking [26] and STM [45], both of which perform worse than our transformed code in our experiments. Further discussion of related work appears in Section 4.

To conclude, this paper demonstrates that generic synchronization, based on a careful combination of optimism and pessimism, is a promising approach for bringing legacy code to emerging computer architectures. While this paper illustrates the method for tree data structures, we believe that the general direction is more broadly applicable, and maybe used with a variety of locking schemes, such as two phase locking. Section 5 concludes the paper and touches on some directions for future work.

2 Transformation

We present an algorithm for a source-to-source transformation, whose goal is to optimize the code of a given data structure implemented using lock-based concurrency control. In Section 2.1, we detail our assumptions about the given code and the locks it uses. Section 2.2 overviews our general approach to combining optimism and pessimism, while Section 2.3 details how the code is instrumented.

2.1 Lock-based Data Structures

A *data structure* defines a set of *operations* that may be invoked by clients of the data structure, potentially concurrently. Operations have parameters and local (private) variables. The operations interact via *shared memory variables*, which are also called *shared objects*. Each shared object supports atomic *read* (load) and *write* (store) instructions. More formal definitions appear in the full version [9].

In addition, each shared object is associated with a lock, which can be unique to the object or common to several (or even all) objects. The object supports atomic *lock* and *unlock* instructions. Locks are exclusive (i.e., a lock can be held by at most one thread at a time), and blocking. We assume that in the given code every (read or write) access by an operation to a shared object is performed when the executing thread holds the lock associated with that object.

The given code only uses the *lock* and *unlock* instructions, while the transformed code can apply in addition atomic non-blocking *tryLock* and *isLockedByAnother* instructions: *tryLock* returns *false* if the lock is currently held by another thread, otherwise it acquires the lock and returns *true*; *isLockedByAnother* returns *true* if and only if the lock is currently held by another thread.

2.2 Combining Optimism and Pessimism

Optimistic concurrency control is a form of synchronization, which accesses shared variables without using locks in the hope that they will not be modified by others before the end of the operation (or more generally, the transaction). To verify the latter, optimistic concurrency control relies on *validation*, which is typically implemented using *version numbers*. If validation fails, the operation restarts. Optimistic execution of update operations requires either performing roll-back (reverting variables to their old values) upon validation failure, or deferring writes to commit time; both approaches induce significant overhead [13]. We therefore refrain from speculative shared memory updates.

The main idea behind our approach is to judiciously use optimistic synchronization only as long as an operation does not update shared state; we use a standard approach based on version numbers to allow validation of optimistic reads. Once an operation writes to shared memory, we revert to pessimistic (lock-based) synchronization. In other words, we rely on validation in order to render redundant locks that would have been acquired and freed before the first update. This scheme is particularly suitable for data structures, since the common behavior of their operations is to first traverse the data structure, and then perform modifications. Since the read only prefix has no side effects we can sandbox it by catching exceptions and infinite loops, and defer validation to the end of the traversal.

Conceptually, our approach thus divides an operation into three phases: an optimistic *read-only phase*, a pessimistic *update phase* and a *validation phase* that conjoins them. The read-only phase traverses the data structure without taking any locks, while maintaining in thread-local variables sufficient information to later ensure the correctness of the traversal. The read phase is *invisible* to other threads, as it updates no shared variables. The update phase uses the original pessimistic (lock-based) synchronization, with the addition of updating version numbers. The validation phase bridges between the optimistic and pessimistic ones. It first locks the objects for which a lock would have been held at this point by the original locking code, and then validates the correctness of the read-only phase. This allows the update phase to run as if an execution of the original pessimistic synchronization took place. If the validation fails, the operation restarts. In order to avoid livelock, we set a threshold on the number of restarts. If the threshold is exceeded, the code falls back on pessimistic execution. We show below that it is safe to do so, since our semi-optimistic code is compatible with the fully pessimistic one.

Phase Transition In many cases, the transition from the read-only phase to the update phase occurs at a statically-defined code location. For example, many data structure operations begin with a read-only traversal to locate the key of interest, and when it is found, proceed to execute code that modifies the data structure. This is the case in all the examples we consider in Section 3 below.

More generally, it is possible to switch from the optimistic read-only execution (via the validation phase) to pessimistic execution at any point before the first

update. Moreover, the phase transition point can be determined dynamically at run time.

One possible way to dynamically track the execution mode is using a flag **opt**, initialized to true, indicating the optimistic phase. Every shared memory update operation is then instrumented with code that checks **opt**, and if it is true, executes the validation phase followed by setting **opt** to false and continuing the execution from the same location.

2.3 Transforming the Code Phases

We now describe how we synthesize the code for each of the phases. We first describe the regular three-phase flow, and then continue with describing the exceptional cases.

Normal Flow We illustrate the transformation for a simple code snippet that adds a new element as the third node in a linked list. Each node is associated with a lock. The original and transformed code are provided in Figure 1. The latter uses the tracking and validation functions in Figures 2 and 3, resp. For clarity of exposition, we present a statically instrumented version, without tracking the phases using **opt**.

Our transformation instruments each lock with an additional field *version*. We assume each object supports *getVersion* and *incVersion* instruction to read and increment the version number of the lock associated with the object. We invoke *incVersion* when holding the lock, and are therefore not concerned about contention. Note that each lock has its own version, i.e., version numbers of different locks are independent of each other.

Read-only Phase In this phase the executing thread is invisible to other threads, i.e., avoids contention on shared memory both in terms of writing and in terms of locking. During this phase, our synchronization maintains two thread-local multi-sets: *lockedSet* and *readSet*. The *lockedSet* tracks the objects that were supposed to be locked by the original synchronization. The *readSet* tracks versions of all objects read by the operation, in order to allow us to later validate that the operation has observed a consistent view of shared memory.

At the beginning of the read-only phase, we insert code that initializes *lockedSet* and *readSet* to be empty (see lines 2-3 of Figure 1b). Throughout the read-only phase, (i.e., when **opt** is true with dynamic phase transitions), we replace every lock and unlock instruction with the corresponding code in Table 1. A lock instruction on object *o* is replaced with code that tracks the object and the version of its lock in *lockedSet* and *readSet* (see Figure 2). An unlock instruction on object *o* is replaced with code that removes *o* from *lockedSet* (see lines 2-13 of Figure 1b).

In Figure 2 (lines 5-6), we use an eager validation scheme³: If the object already exists in *readSet*, we check that the current version of its lock is equal

³ Eager validation is not required for correctness.

<pre> 1: FUNCTION addThird(List list, Node new) ----- ▷ read-only phase 2: 3: 4: list.lock() 5: Node prev = list.head 6: prev.lock() 7: list.unlock() 8: Node succ = prev.next 9: succ.lock() 10: prev.unlock() 11: prev = succ 12: succ = succ.next 13: succ.lock() ----- 14: 15: 16: 17: 18: 19: 20: 21: ----- ▷ update phase 22: prev.next = new 23: new.lock() 24: new.next = succ 25: 26: prev.unlock() 27: 28: new.unlock() 29: 30: succ.unlock() </pre>	<pre> 1: FUNCTION addThird(List list, Node new) ----- ▷ read-only phase 2: lockedSet.init() 3: readSet.init() 4: if !track(list) then goto 1 5: Node prev = list.head 6: if !track(prev) then goto 1 7: lockedSet.remove(list) 8: Node succ = prev.next 9: if !track(succ) then goto 1 10: lockedSet.remove(prev) 11: prev = succ 12: succ = succ.next 13: if !track(succ) then goto 1 ----- ▷ validation phase 14: read fence 15: for all obj in lockedSet do 16: if !obj.tryLock() then 17: unlockAll() 18: goto 1 19: if !validateReadSet() then 20: unlockAll() 21: goto 1 ----- ▷ update phase 22: prev.next = new 23: new.lock() 24: new.next = succ 25: prev.incVersion 26: prev.unlock() 27: new.incVersion 28: new.unlock() 29: succ.incVersion 30: succ.unlock() </pre>
(a) Code with original locking	(b) The code produced by our transformation

Fig. 1: Code transformation example. The synchronization code is in bold.

```

1: FUNCTION track(obj)
2:   lockedSet.add(obj)
3:   long ver = obj.getVersion()
4:   readSet.add({obj, ver})
5:   if {obj, v} ∈ readSet and v != ver then return false
6:   if obj.isLockedByAnother() then return false
7:   return true

```

Fig. 2: In read-only phase, locking is replaced by tracking locks and read objects' versions.

```

1: FUNCTION validateReadSet()
2:   for all (obj,ver) in readSet do
3:     if obj.isLockedByAnother() then
4:       return false           ▷ validation failed (locked object)
5:     if obj.getVersion() != ver then
6:       return false           ▷ validation failed (different version)
7:   retrun true                 ▷ validation succeed

```

Fig. 3: Read set validation: verify that objects are unlocked and their versions are unchanged.

Original code	Transformed Code
<code>x.lock()</code>	<code>if !track(x) then goto S</code>
<code>x.unlock()</code>	<code>lockedSet.remove(x)</code>

Table 1: Transformation for read-only phase: each locking instruction (left column) is replaced with the corresponding code on the right; S denotes the beginning of the operation.

to the version in *readSet*; and if the versions are different the operation restarts (line 5). Similarly, it is checked to be unlocked, and the operation restarts if it is locked (line 6).

Although it only accesses thread-local data structures, lock tracking induces a certain overhead due to the need to search a lock in the *lockedSet* in order to unlock it. (In our experiments presented below, in large data structures, this overhead slows operations down by up to 40%). We suggest some optimizations to mitigate this cost. First, we observe that the *lockedSet* does not need to be tracked in read-only operations, which a compiler can easily detect. We can further avoid this overhead in update operations in certain cases by relying on the structure of the transformed code. For example, if the lock-based code is created from sequential code using domination locking [26], then at any given time in the read phase, it holds locks on a well-defined set of objects – the ones currently pointed by the operation’s local variables. When applying our transformation to code generated by this scheme, we can optimize it to remove lock-tracking, and instead populate the *lockedSet* with the appropriate locks immediately before executing the validation phase.

Validation Phase The code of the validation phase is invoked between the read-only phase and the update phase (lines 14-21 of Figure 1b). It locks the objects that are left in *lockedSet* and validates the objects in *readSet*. To avoid deadlocks, the locks are acquired using a *tryLock* instruction. If any *tryLock* fails, the code unlocks all previously acquired locks and restarts from the beginning (lines 15-18).

The function *validateReadSet* in Figure 3 verifies that the objects in the read set have not been updated. The function checks that each object in the read

set is not locked by another thread, and that the current version of the lock associated with the object matches the version saved in the *readSet*. This check guarantees that the object was not locked from the time it was read until the time it was validated. Since operations write only to locked objects, it follows that the object was not changed. This *readSet* validation can be viewed as a double collect [4] of all objects accessed by the read-only phase. The operation is restarted if the validation fails (lines 19-21).

We assume that, following standard practice in lock implementations, the function *isLockedByAnother* imposes a *memory fence* (barrier). This ensures that the lock and version are read during *track* before the object's value is read optimistically during the read-only phase. To ensure that the second read of the lock and version, during the validation phase, succeeds the optimistic read of the object's value, we precede the validation phase with a memory fence as well (line 14). Note that it suffices to impose a *read fence* (sometimes called acquire or load fence) prior to the validation as well as during *isLockedByAnother*, because this part of the code does not include writes to shared memory.

Update Phase In this phase our transformation preserves the original locking while maintaining the versions of the objects, i.e., the version of an object *o* is incremented every time *o* is unlocked. Here, (i.e., in case **opt** is false with dynamic phase transitions), before each unlock instruction *x.unlock()* we insert the code *x.incVersion()*. An example is shown in lines 22-30 of Figure 1b.

Exceptions from Regular Flow The read phase does not validate past reads during its executions (other than when re-reading the same variable). As a result, it may observe an inconsistent state of shared memory, which may lead to infinite loops or spurious exceptions (as explained, e.g., in [30]). We avoid such infinite loops using a timeout. If the timeout expires before the read-only phase completes, read set validation takes place (via the function *validateReadSet*). If the validation fails, the operation is restarted. This is realized by inserting code that examines the timeout in every loop iteration in the original code. Similarly, we avoid spurious exceptions by catching all exceptions and performing validation. Here too, if the validation fails, the operation is restarted. Otherwise, the exception is handled as in the original code.

Our sandboxing relies on properties of managed languages like Java or C#:

1. We can identify all instructions that may update shared memory and end the read phase before they occur.
2. The ability to capture all exits from a block via the try-finally mechanism ensures that we never exit the read phase without performing validation.
3. The code is not self-modifying and hence the tracking and validation code is executed as intended.
4. The speculative execution does not alter the references to the thread-local variables we introduce (*readSet*, *lockedSet*) since they are constant references to well-typed objects.

Hence, our tracking and validation code executes correctly.

While recent work [17] has shown that differing validation to the end of a transaction can be unsafe, this problem does not occur in our solution. The key problem shown there is that access to the object on which the conflict is checked (namely the lock) is deferred until after other unchecked shared accesses, which could potentially be inconsistent and cause the lock not to be accessed. In our case, on the other hand, all accesses to shared data are recorded for validation purposes. If an object that should be accessed (like the lock in the lock elision case) is not accessed because of earlier conflicts, these earlier conflicts will be detected and the transaction will abort.

Note that, using our transformation, the shared state at the end of the validation phase is identical to the state that would have been reached had the code been executed pessimistically from the outset. Hence, the three-phase version of the code is compatible with the instrumented pessimistic version. This means that if the optimistic phase is unsuccessful for any reason, we can always fall back on the pessimistic version. Moreover, we can switch from optimistic to pessimistic synchronization *at any point* during the read phase. We use this property in two ways. First, we avoid livelocks by limiting the number of restarts due to conflicts: The validation phase tracks the number of restarts in a thread-local variable. If this number exceeds a certain threshold, we perform the entire operation optimistically.

Second, this property offers the optimistic implementation the liberty of failing spuriously, even in the absence of conflicts, because it can always fall back on the safe pessimistic version of the code. One can take advantage of this liberty, and implement the *readSet* using a constant size array. In case the array becomes full, the optimistic version cannot proceed, but there is no need to start the operation anew. Instead, one can immediately perform the validation phase, which, if successful, switches to a pessimistic modus operandi, after having acquired all the needed locks.

3 Evaluation

We evaluate the performance of our approach on search trees supporting insert, delete, and get operations. We compare the throughput of our approach to fully pessimistic solutions applying fine-grain locking, solutions based on software transactional memory, and hand-crafted state-of-the-art data structure implementations.

Methodology We use the micro-benchmark suite *Synchrobench* [27], configured as follows. Each experiment consists of 5 trials. A trial is a five second run in which each thread continuously executes randomly chosen operations drawn from the workload distribution, with keys selected uniformly at random from the range $[0, 2 \cdot 10^6]$. Each trial is preceded by initiating a new data structure with 10^6 keys and a warm-up of five seconds. Our graphs present the average throughput over all trials. We consider three representative workload distributions: a *read-only* workload comprised of 100% lookup operations, a *write-dominated* workload

consisting of insert and delete operations (50% each), and a *mixed workload* with 50% lookups, 25% inserts, and 25% deletes.

Platform All implementations are in Java. We ran the experiments on a dedicated machine with four Intel Xeon E5-4650 processors, each with 8 cores, for a total of 32 threads (with hyper-threading disabled). We used Ubuntu 12.04.4 LTS and Java Runtime Environment (build 1.7.0_51-b13) using the 64-Bit Server VM (build 24.51-b03, mixed mode).

Implementations We start from textbook sequential implementations of an unbalanced internal binary tree and a treap [7]. We next synthesize concurrent lock-based code by (manually) applying the domination locking technique [26] to the sequential data structures. The resulting algorithms are denoted Lock-Tree and Lock-Treap. Then, we manually apply our lock-removal transformation to the reference implementations by following the algorithm line-by-line (requiring no understanding of the base code) to get our semi-optimistic versions of the code, which we call LR-Tree and LR-Treap, respectively. Note that this solution does not track the *lockedSet* for read-only operations and does not use eager validation of version numbers. Finally, we apply the optimization described in Section 2.3, which eliminates explicit tracking of the *lockedSet* in update operations, and instead locks all objects the thread holds a pointer to in the validation phase; this optimization is applicable since our parallel implementation is synthesized using domination locking. The resulting algorithms are denoted Opt-LR-Tree and Opt-LR-Treap.

For the competition, we parallelize the sequential implementations also using Deuce [24], a Java implementation of TL2 [18]. The resulting algorithms are denoted STM-Tree and STM-Treap. We further compare our implementations to their hand-crafted state-of-the-art counterparts listed in Table 2.

Unbalanced		Balanced	
LO-Tree	Locked-based [20]	LO-AVL	Lock-based relaxed AVL [20]
LF-Tree	Lock-free [22]	Snap-Tree	Lock-based relaxed AVL [11]
		CF-Tree	Contention-friendly tree [15]
		Skiplist	Java lock-free skiplist

Table 2: Hand-crafted state-of-the-art data structures. The code of LO-Tree was provided by the authors, all other implementations provided by Synchrobench.

We also measured the performance of global lock-based implementations. In all workloads, the results were identical or inferior to those achieved by pessimistic fine-grain locking. We hence omitted these results to avoid obscuring the presentation.

Results Figures 4 and 5 show the throughput of unbalanced and balanced data structures, resp. We see that our semi-optimistic solution, both optimized and

unoptimized, is far superior to the fully-pessimistic automated approach; it successfully overcomes the bottlenecks associated with lock contention in Lock-Tree and Lock-Treap.

Our approach also outperforms STM by 1.5x to 2.5x. The additional overhead of STM most likely stems from two reasons: deferring writes to commit time, and using a global clock to ensure a consistent view of the read set. The latter is done in order to satisfy opacity [29], which we avoid by “sandboxing”. In our experiments, the code *never* incurred a spurious exception or timeout due to inconsistent reads, and so the sandboxing was not associated with a performance penalty.

Our solution comes close to custom-tailored implementations, and the optimized version is even superior to some of them. The throughput of our read-only operations is up to 1.5x lower than that achieved by the best-in-class. By profiling the code, we learned that the bulk of this overhead stems from the need to track all read objects, which is inherent to our transformation. This is in contrast with the hand-crafted implementations, which have small overhead on reads that complete without any retries. In workloads that include update operations, our solution is up to 2.2x slower. This stems from tracking read and locked sets and not from retries as the percentage of retries is less than 1%.

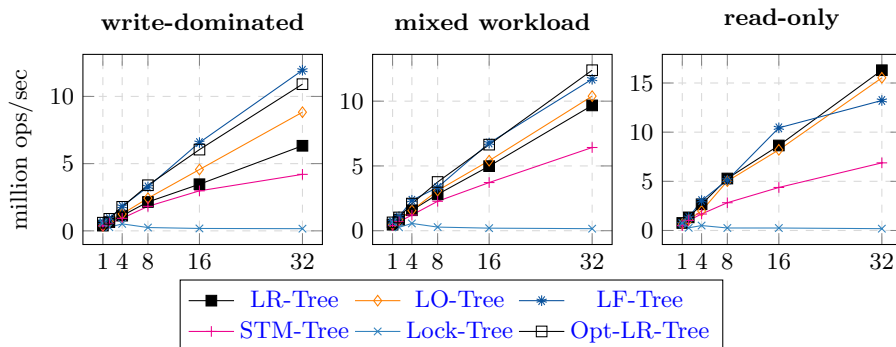


Fig. 4: Throughput of unbalanced data structures.

4 Related Work

Concurrent Data Structures Many sophisticated concurrent data structures (e.g., [8, 20, 40, 12, 15, 10, 5, 22, 11, 32, 25, 38]) were developed and used in concurrent software systems [44]. Implementing efficient synchronization for such data structures is considered a challenging and error-prone task [44, 19, 35]. As a result, concurrent data structures are manually implemented by concurrency ex-




Fig. 5: Throughput of balanced data structures.

perts. This paper shows that (in some cases) an automatic algorithm can produce synchronization that is comparable to synchronization implemented by experts.

Lock Inference Algorithms There has been a lot of work on automatically inferring locks for transactions. Most algorithms in the literature infer locks that follow the two-phase locking protocol [36, 23, 28, 14, 33, 16]. Our approach can potentially be used to optimize the synchronization produced by these algorithms. For example, for algorithms that employ a two-phase variant in which all locks are acquired at the beginning of a transaction (e.g., [28, 14]), our approach may be used to defer the locking (e.g., to just before the first write operation) and even to eliminate some of the locking steps. We demonstrate the benefit of combining our transformation with such algorithms by using the domination locking protocol [26] to produce efficient concurrency control for dynamic data structures.

Transactional Memory Transactional memory approaches (TMs) dynamically resolve inconsistencies and deadlocks by rolling back partially completed transactions. Unfortunately, in spite of a lot of effort and many TM implementations (see [30]), existing TMs have not been widely adopted due to various concerns [21, 13, 37], including high runtime overhead, poor performance and limited ability to handle irreversible operations. Modern concurrent programs and data structures are typically based on hand-crafted synchronization, rather than on a TM approach [44].

Lock Elision Our transformation is inspired by the idea of *sequential locks* [37] and the approach presented in [39], which replace locks with optimistic concurrency control in read-only transactions. But in contrast to these works, we handle read-only prefixes of transactions (operations) that do update the shared memory. In fact, as shown in Section 3, our approach is best suited for update-dominated workloads. Moreover, using these approaches for a highly-contended

data structure (as in Section 3) is likely to provide limited performance, because each update transaction causes many read-only transactions to abort.

Other works have proposed using transactional memory in order to elide locks from arbitrary critical sections, and fall back on lock-based execution in cases of aborts (e.g., [41, 42, 6]). In contrast to our approach, however, lock elision does not combine speculative and non-speculative execution within the same transaction.

5 Discussion

The development of scalable concurrent programs today heavily relies on custom-tailored implementations, which require painstaking correctness proofs. In this paper, we have shown a relatively simple transformation that can facilitate this labor-intensive process, and thus make scalable synchronization more readily available. The input for our transformation is a conventional lock-based concurrent program, which may be either constructed manually or synthesized from sequential code. Our source-to-source transformation then makes judicious use of optimism in order to eliminate principal concurrency bottlenecks in the given program and improve its scalability.

We have illustrated our method for unbalanced and balanced search trees. The transformed code performed significantly better than the original lock-based one, and scaled comparably to hand-crafted implementations that had taken considerably more effort to produce. In these examples, we have applied our transformation manually. An interesting direction for future work would be to create a tool that automatically applies our transformation at compile time.

Our approach makes use of a common pattern in data structures, where an operation typically begins with a long read-only traversal, followed by a handful of (usually local) modifications. A promising direction for future work is to try and exploit similar patterns in order to parallelize or remove locks in other types of code (not data structures), for example, programs that rely on two-phase locking. Furthermore, for programs that follow different patterns, other combinations of optimism and pessimism may prove effective.

Finally, there still remains a gap between the performance achievable by manually optimized solutions and what we could achieve automatically. Our algorithm induces inherent overhead for tracking all operations in the read-only phase for later verification. In specific data structures, these checks might be redundant, but it is difficult to detect this automatically. We believe that it may well be possible to bridge the remaining performance gap using computer-assisted optimizations. For example, a programmer may provide hints regarding certain invariants that are always preserved in the code, in order to eliminate the need for tracking some values for later validation.

References

- [1] Concurrentskiplistmap from java.util.concurrent. docs.oracle.com/javase/7/docs/api/java/util/concurrent/ConcurrentSkipListMap.html.

- [2] A fast and lightweight key/value database library by google. <http://code.google.com/p/leveldb>.
- [3] jmonkeyengine: a 3d game engine for java developers. <http://jmonkeyengine.org/>.
- [4] Y. Afek, H. Attiya, D. Dolev, E. Gafni, M. Merritt, and N. Shavit. Atomic snapshots of shared memory. *J. ACM*, 40(4):873–890, Sept. 1993.
- [5] Y. Afek, H. Kaplan, B. Korenfeld, A. Morrison, and R. E. Tarjan. CBTree: A practical concurrent self-adjusting search tree. In *DISC*, pages 1–15, 2012.
- [6] Y. Afek, A. Levy, and A. Morrison. Software-improved hardware lock elision. In *PODC*, 2014.
- [7] C. R. Aragon and R. Seidel. Randomized search trees. In *FOCS*, pages 540–545, 1989.
- [8] M. Arbel and H. Attiya. Concurrent updates with RCU: search tree as an example. In *PODC*, pages 196–205, 2014.
- [9] M. Arbel, G. Golan-Gueta, E. Hillel, and I. Keidar. Towards automatic lock removal for scalable synchronization (full version). <https://labs.yahoo.com/publications/8476/towards-automatic-lock-removal-scalable-synchronization-full-version>.
- [10] A. Braginsky and E. Petrank. A lock-free B+tree. In *SPAA*, pages 58–67, 2012.
- [11] N. G. Bronson, J. Casper, H. Chafi, and K. Olukotun. A practical concurrent binary search tree. In *PPOPP*, pages 257–268, 2010.
- [12] T. Brown, F. Ellen, and E. Ruppert. A general technique for non-blocking trees. In *PPoPP*, pages 329–342, 2014.
- [13] C. Cascaval, C. Blundell, M. Michael, H. W. Cain, P. Wu, S. Chiras, and S. Chatterjee. Software transactional memory: Why is it only a research toy? *Queue*, 6(5):46–58, Sept. 2008.
- [14] S. Cherem, T. Chilimbi, and S. Gulwani. Inferring locks for atomic sections. In *PLDI*, 2008.
- [15] T. Crain, V. Gramoli, and M. Raynal. A contention-friendly binary search tree. In *19th International Conference on Parallel Processing (Euro-Par)*, pages 229–240, 2013.
- [16] D. Cunningham, K. Gudka, and S. Eisenbach. Keep off the grass: Locking the right path for atomicity. In *CC*, pages 276–290, 2008.
- [17] D. Dice, T. L. Harris, A. Kogan, Y. Lev, and M. Moir. Hardware extensions to make lazy subscription safe. *CoRR*, abs/1407.6968, 2014.
- [18] D. Dice, O. Shalev, and N. Shavit. Transactional locking ii. In *DISC*, pages 194–208, 2006.
- [19] S. Doherty, D. L. Detlefs, L. Groves, C. H. Flood, V. Luchangco, P. A. Martin, M. Moir, N. Shavit, and G. L. Steele, Jr. Dcas is not a silver bullet for nonblocking algorithm design. In *SPAA*, 2004.
- [20] D. Drachler, M. T. Vechev, and E. Yahav. Practical concurrent binary search trees via logical ordering. In *PPoPP*, pages 343–356, 2014.
- [21] J. Duffy. A (brief) retrospective on transactional memory. 2010. <http://joeduffyblog.com/2010/01/03/a-brief-retrospective-on-transactional-memory>.
- [22] F. Ellen, P. Fatourou, E. Ruppert, and F. van Breugel. Non-blocking binary search trees. In *PODC*, pages 131–140, 2010.
- [23] M. Emmi, J. S. Fischer, R. Jhala, and R. Majumdar. Lock allocation. In *POPL*, pages 291–296, 2007.
- [24] G. K. Felber, P. and N. Shavit. Deuce: Noninvasive concurrency with a Java STM. In *MULTIPROG*, 2010.

- [25] K. Fraser. *Practical lock-freedom*. PhD thesis, University of Cambridge, 2004.
- [26] G. Golan-Gueta, N. G. Bronson, A. Aiken, G. Ramalingam, M. Sagiv, and E. Yahav. Automatic fine-grain locking using shape properties. In *OOPSLA*, pages 225–242, 2011.
- [27] V. Gramoli. More than you ever wanted to know about synchronization: Synchrobench, measuring the impact of the synchronization on concurrent algorithms. In *PPoPP*, 2015.
- [28] K. Gudka, T. Harris, and S. Eisenbach. Lock inference in the presence of large libraries. In *ECOOP*. 2012.
- [29] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *PPOPP*, pages 175–184, 2008.
- [30] T. Harris, J. Larus, and R. Rajwar. Transactional memory, 2nd edition. *Synthesis Lectures on Computer Architecture*, 5(1), 2010.
- [31] S. Heller, M. Herlihy, V. Luchangco, M. Moir, W. N. Scherer III, and N. Shavit. A lazy concurrent list-based set algorithm. In *9th International Conference on Principles of Distributed Systems (OPODIS)*, pages 3–16, 2006.
- [32] M. Herlihy, Y. Lev, V. Luchangco, and N. Shavit. A simple optimistic skiplist algorithm. In *14th International Conference on Structural Information and Communication Complexity (SIROCCO)*, pages 124–138, 2007.
- [33] M. Hicks, J. S. Foster, and P. Prattikakis. Lock inference for atomic sections. In *Proceedings of the First ACM SIGPLAN Workshop on Languages, Compilers, and Hardware Support for Transactional Computing*, June 2006.
- [34] A. Israeli and L. Rappoport. Disjoint-access-parallel implementations of strong shared memory primitives. In *PODC*, pages 151–160, 1994.
- [35] G. Jin, W. Zhang, D. Deng, B. Liblit, and S. Lu. Automated concurrency-bug fixing. In *OSDI*, 2012.
- [36] B. McCloskey, F. Zhou, D. Gay, and E. Brewer. Autolocker: synchronization inference for atomic sections. In *POPL*, pages 346–358, 2006.
- [37] P. E. McKenney. Is parallel programming hard, and, if so, what can you do about it? *Linux Technology Center, IBM Beaverton*, Aug. 2012.
- [38] M. M. Michael and M. L. Scott. Simple, fast, and practical non-blocking and blocking concurrent queue algorithms. In *PODC*, 1996.
- [39] T. Nakaike and M. M. Michael. Lock elision for read-only critical sections in java. In *PLDI*, pages 269–278, 2010.
- [40] A. Natarajan and N. Mittal. Fast concurrent lock-free binary search trees. In *PPoPP*, pages 317–328, 2014.
- [41] R. Rajwar and J. R. Goodman. Transactional lock-free execution of lock-based programs. *SIGOPS Oper. Syst. Rev.*, 36(5):5–17, Oct. 2002.
- [42] A. Roy, S. Hand, and T. Harris. A runtime system for software lock elision. In *EuroSys*, 2009.
- [43] O. Shacham. *Verifying Atomicity of Composed Concurrent Operations*. PhD thesis, Tel Aviv University, 2012.
- [44] O. Shacham, N. Bronson, A. Aiken, M. Sagiv, M. Vechev, and E. Yahav. Testing atomicity of composed concurrent operations. In *OOPSLA*, 2011.
- [45] O. Shalev and N. Shavit. Predictive log-synchronization. In Y. Berbers and W. Zwaenepoel, editors, *EuroSys*, pages 305–315. ACM, 2006.
- [46] F. Zylkyarov, V. Gajinov, O. S. Unsal, A. Cristal, E. Ayguadé, T. Harris, and M. Valero. Atomic quake: using transactional memory in an interactive multiplayer game server. In *ACM Sigplan Notices*, volume 44, pages 25–34. ACM, 2009.