Every Data Structure Deserves Lock-Free Memory Reclamation

NACHSHON COHEN, EPFL, Switzerland

Memory-management support for lock-free data structures is well known to be a tough problem. Recent work has successfully reduced the overhead of such schemes. However, applying memory-management support to a data structure remains complex and, in many cases, requires redesigning the data structure. In this paper, we present the first lock-free memory-management scheme that is applicable to general (arbitrary) lock-free data structures and that can be applied automatically via a compiler plug-in. In addition to the simplicity of incorporating to data structures, this scheme provides low overhead and does not rely on the lock freedom of any OS services.

CCS Concepts: • **Software and its engineering** → **Allocation / deallocation strategies**; **Concurrency control**; **Concurrent programming structures**;

Additional Key Words and Phrases: Lock-free memory management, Hazard Pointers, Optimistic Access

ACM Reference Format:

Nachshon Cohen. 2018. Every Data Structure Deserves Lock-Free Memory Reclamation. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 143 (November 2018), 24 pages. <https://doi.org/10.1145/3276513>

1 INTRODUCTION

Lock-free data structures simplify reasoning about concurrent code. They are immune to dead locks, live locks, and they aid programmers in achieving high scalability and responsiveness. However, providing memory-reclamation support for such data structures, without foiling their progress guarantee, is difficult. Indeed, many lock-free data structures were published without any memory-reclamation support, as it was assumed that the application is short lived enough so that all memory is allocated in advance. This prevents the use of such data structures in real applications.

Early solutions to this problem include the lock-free reference-counting scheme [Michael and Scott 1995; Valois 1995], the epoch-based reclamation scheme [Harris 2001] and the hazard-pointers scheme [Michael 2002a, 2004]. However, the epoch-based reclamation scheme foils the lock-freedom guarantee of the data structure; and the reference-counting and hazard-pointers schemes add significant overhead [Hart et al. 2007]. Recently, there has been a significant interest in designing lock-free reclamation schemes that support the lock-freedom progress guarantee, yet they do not add significant overhead [Alistarh et al. 2017, 2015; Balmau et al. 2016; Braginsky et al. 2013; Brown 2015; Cohen and Petrank 2015a,b; Dice et al. 2016].

Although the above-mentioned solutions substantially reduce the overhead of lock-free memory-reclamation schemes, they address only specific (interesting) data structures. There is still a wide discrepancy between the expectations of the data-structure designer and the state-of-the-art memory-reclamation schemes. The pervasive assumption of data-structure designers is that memory reclamation is a non-issue: it is possible to publish a paper describing a novel design of a lock-free data structure without considering memory reclamation. It remains the responsibility of

Author's address: Nachshon Cohen, EPFL, Lausanne, Switzerland, nachshonc@gmail.com.

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

© 2018 Copyright held by the owner/author(s).

2475-1421/2018/11-ART143

<https://doi.org/10.1145/3276513>

a (possibly non-expert) practitioner to apply a memory-reclamation scheme when incorporating the data structure into a real application. However, state-of-the-art memory reclamation schemes require the data structure to satisfy certain algorithmic properties that are not met by many practical data-structures. If a data structure does not satisfy the required properties, there is no clear recipe for modifying it so that it would be amenable to memory reclamation. In general, applying a lock-free memory-reclamation scheme to a given data structure might require expertise that is not available to all practitioners, might modify the properties of the data structure, or might not be possible at all.

Consider, for example, Harris’s linked list [Harris 2001] that stands at the core of many lock-free data structures. The original hazard pointers paper [Michael 2002a] states that Harris’s linked list is “inherently incompatible” with the hazard-pointers scheme¹. Instead, they designed a variation of this data structure — the Harris-Michael linked list [Michael 2002b] — on which hazard pointers can be applied. The differences between these data structures are subtle, hence it is easy to confuse them. Such confusion can lead even experienced researchers to apply hazard pointers on the wrong data structure, resulting in an incorrect design. Furthermore, it was later shown that the Harris-Herlihy-Shavit linked list [Herlihy and Shavit 2012] is also incompatible with the hazard-pointers scheme. This data structure offers stronger progress guarantees (i.e., wait-free searches) and better performance than the Harris-Michael linked list [Cohen and Petrank 2015a]. Other recent memory-reclamation schemes [Alistarh et al. 2015; Balmau et al. 2016; Braginsky et al. 2013; Brown 2015; Cohen and Petrank 2015b; Dice et al. 2016] are similarly incompatible with the Harris and Harris-Herlihy-Shavit linked-list data structures.

In this paper, we present Free Access: the first memory-reclamation scheme that stands up to the expectations of the data-structure designers. The Free Access scheme does not require the data structure to satisfy any algorithmic properties, hence making it generally applicable for existing and future lock-free data structures. Instead, it requires only that no instruction simultaneously reads a new pointer and writes to the shared memory. All lock-free data structures we are aware of satisfy this property without any modification. We cannot preclude a future novel design that uses an unsupported instruction, but even such a design could be modified easily by replacing the unsupported instruction with a Compare-And-Swap (CAS) loop, enabling Free Access to be applied². The main impact of the Free Access scheme is the decoupling of the design of lock-free data structures from the memory-reclamation effort; this decoupling enables designers of lock-free data structures to concentrate on designing the data structure without worrying whether their design is amenable to memory reclamation.

The Free Access scheme is based on the concept of data-structure-specific garbage collection, introduced by Cohen and Petrank [2015a]. During the execution of the program, Free Access automatically determines the set of unreachable data-structure nodes and reclaims them for future allocations. The reclamation algorithm of Free Access is very similar to the Automatic Optimistic Access (AOA) scheme [Cohen and Petrank 2015a] and is loosely based on the mark&sweep garbage-collection technique. It starts by gathering the *global roots*, data-structure nodes that are reachable by all threads (e.g., the head of a linked list); and the *local roots*, the data-structure nodes residing in the registers and stack of each thread. Then, it traces these roots and marks all data-structure nodes that are reachable from one of the roots. Finally, it reclaims all data-structures nodes that are unmarked. (More background information is provided in Section 2).

Although the reclamation algorithm of Free Access is very similar to the AOA scheme, there is a major difference between the two. The AOA scheme is applicable only for data structures written

¹Section 2.2.1 provides more background on this incompatibility

²Section 6 discusses this unsupported instruction and presents the replacement algorithm.

in a special normalized form of Timnat and Petrank [2014]. As most standard data-structures are not written in this form, applying the AOA scheme requires the practitioner to understand the data-structure algorithm well enough to re-write the data structure in this normalized form (or might not be possible at all). In contrast, the Free Access scheme does not require the data structure to be presented in any special format, which enables a (possibly non-expert) practitioner to easily apply Free Access, without understanding the internal data-structure details.

Furthermore, we implemented an LLVM compiler pass that automatically modifies, according to the Free Access scheme, the data structure code. The programmer needs only to add the `freeaccess` attribute to each function that corresponds to a data-structure operation and run the compiler pass. This stands in contrast to the existing lock-free memory-reclamation schemes that require the programmer to manually modify all loads (and possibly stores) in the data-structure code. Automatically applying the required modifications reduces the programming effort, keeps the code simple, and avoids bugs from imprecise modifications. We note that implementing a similar compiler pass for other lock-free reclamation schemes is complicated due to the need to inform the compiler about data-structure-specific mechanisms (e.g., how to restart an operation or how to correctly interpret the normalized form). This need is eliminated by the Free Access scheme.

The main challenge in designing an automatic lock-free reclamation scheme is collecting local roots in a lock-free manner. Although many garbage collectors attempt to provide partial guarantees of progress [Auerbach et al. 2008; Herlihy and Moss 1992; Hudson and Moss 2001; Pizlo et al. 2007, 2008, 2010], they still require all threads to perform a synchronization barrier (e.g., handshake) for gathering local roots. If a thread fails to participate in the synchronization barrier, the local roots cannot be gathered and memory cannot be reclaimed. Eventually, threads will run out of memory and fail to make progress, thus foiling the lock-freedom guarantee [Petrank 2012]. The AOA algorithm gathers local roots without any synchronization barrier, but heavily relies on a special format of the data structure, which strictly specifies how local roots can be used.

The Free Access scheme efficiently gathers local roots by partitioning the execution into read-only and write-only periods. Read-only periods can load new pointers but never modify the shared memory; write-only periods modify the shared memory but never load new pointers to a local variable. The Free Access scheme incurs a low overhead for write-only periods by noticing that the set of local roots is not modified during this period. It incurs an even lower overhead for read-only periods by designing a restart mechanism for arbitrary read-only periods. Using this restart mechanism, during read-only periods the thread is not required to publish the set of local roots at all; it only needs to restart if one of its local roots is reclaimed by concurrent threads.

Our contributions in this paper are as follows.

- A new algorithm for gathering local roots for general lock-free data structures, without depending on any algorithmic properties of the data structure.
- The Free Access scheme, which combines our root gathering algorithm with the lock-free tracing algorithm of the automatic optimistic access scheme. The Free Access scheme can be applied to general lock-free data structure, thus enabling a separation of concerns between designing the data structure and applying it to a real application.
- The implementation of an LLVM compiler pass that automatically applies the Free Access modifications, which reduces programming efforts and buggy implementations.

Organization. In Section 2, we provide preliminaries and discuss some background. We present an overview of the Free Access scheme in Section 3 and details in Section 4. In Section 5 we present measurements of the algorithm. We discuss the assumption of separating code into read-only and write-only periods in Section 6 and related work in Section 7. We conclude in Section 8.

2 BACKGROUND AND PRELIMINARIES

2.1 System Model and Progress Guarantees

Lock-freedom for memory reclamation. We use the standard computation model of Herlihy [1991]. A shared memory is accessible by all threads. The threads communicate through memory-access instructions on the shared memory; a thread makes no assumptions about the status of any other threads, nor about the speed of its execution.

A lock-free progress guarantee ensures that while the system executes a bounded number of steps, at least one thread progresses. Fair scheduling must not be assumed and the system must progress even if the scheduler is adversarial. Lock-free data structures are fast, scalable and widely used. They eliminate deadlock, livelock, and provide guaranteed system responsiveness. However, they permit a single thread to starve (even if it is scheduled enough steps) if other threads in the system progress. Wait-free progress guarantee ensures that a thread progresses after it executes a finite number of steps, regardless of the behavior of other threads in the system. A series of work [Fatourou and Kallimanis 2011; Kogan and Petrank 2011, 2012; Timnat et al. 2012], culminating in Timnat and Petrank method [Timnat and Petrank 2014], shown that wait-free algorithms can be designed efficiently.

Providing a lock-freedom guarantee for memory reclamation requires that under any schedules, including adversarial ones, at least one thread is able to progress in a finite number of steps. Waiting until all threads respond before reclaiming memory violates lock-freedom. Consider, for example, a schedule where a thread never executes any instructions. As this thread cannot respond, memory cannot be reclaimed. Eventually, other threads will not be able to allocate memory and will fail to progress, hence violating lock-freedom.

We assume that the number of nodes available for allocation (i.e., available memory) is sufficiently larger than the maximum number of nodes that are simultaneously used by the data structure. Otherwise, a data structure that allocates memory, but never frees, would violate lock-freedom in a non-interesting manner.

Read reclaimed nodes. The Free Access scheme assumes that the memory allocator is strictly lock-free and that it does not return memory to the OS, which results in lock-acquisition for the current Linux implementation. Consequently, all allocated memory resides in the application's address space and accessing it does not generate a segmentation fault, *even if the accessed node was already reclaimed*. Alternatively, it assumes that a segmentation fault can be caught by a signal handler and ignored. Specifically, we follow the AOA scheme and make the following assumption.

ASSUMPTION 1. *If a thread reads the memory of a reclaimed node, the read value is arbitrary. However, the thread is allowed to proceed to the next step of execution.*

2.2 Lock-Free Memory-Reclamation Schemes

In this subsection, we provide some necessary background regarding lock-free memory-reclamation schemes. We begin by briefly presenting the hazard-pointers scheme and describe why it does not support the Harris-Herlihy-Shavit linked-list data structure. The Free Access scheme extends the AOA [Cohen and Petrank 2015a] scheme and inherits some of its properties. In this subsection, we sketch the design of the AOA scheme and present some parts of the algorithm that Free Access inherits. We also sketch the normalized-form representation, on which AOA depends. Recall that the main difference between Free Access and AOA is that AOA requires the data structure to be presented in a special normalized form, whereas Free Access does not pose this restriction.

2.2.1 The Hazard-Pointers Scheme. The hazard-pointers scheme is considered the most popular memory-reclamation scheme for lock-free data structures. The hazard-pointers scheme associates

Listing 1. Reading a pointer in the hazard-pointers scheme

```

1 Node *HPRead(Node **addr, int i){
2   while(true){
3     Node *newnode = *addr; //read address of the new node
4     thread.HP[i].store(newnode, memory_order_seq_cst); //ensure that newnode is visible. Expensive
5     if(*addr == newnode) //validate that new node is still reachable from previous node
6       return newnode;
7     else continue; //failed to ensure that newnode is protected. Restart.
8   }
9 }

```

a set of *hazard pointers* to each thread; these pointers are used by the thread to announce that a set of data-structure nodes are currently accessed by the thread and must not be reclaimed. A hazard pointer can be written only by the owner thread, but can be read by other threads in the system. The hazard-pointers scheme modifies the data-structure code by replacing all reads of a new data-structure node with the code of Listing 1. Reclaiming a node requires three steps. First, the node must be disconnected from the data structure. Second, the application passes the node to the *retire* method, which announces the node as a candidate for reclamation. Third, the scheme waits until no hazard pointer points to this node then reclaims the node (make it reusable). Note that a candidate for reclamation (Step 2) is not immediately reclaimed as a thread could access it concurrently.

The hazard-pointers scheme cannot be applied to all data structures and specifically cannot be applied to Harris’s linked list or Harris-Herlihy-Shavit linked list. To explain why, consider Harris-Herlihy-Shavit linked list and suppose that no memory reclamation is used³. Consider a snippet of the linked list with 4 nodes $\boxed{1} \Rightarrow \boxed{2} \Rightarrow \boxed{3} \Rightarrow \boxed{4}$ and a thread T that searches whether key 4 exists, while currently traversing node $\boxed{2}$. While T is asleep, concurrent threads remove $\boxed{2}$ and $\boxed{3}$ and disconnect them from the list. When T wakes up, a possible way to continue the traversal is to restart from the head pointer, then traverse $\boxed{1}$ and $\boxed{4}$. There exists, however, a better way. Even though $\boxed{2}$ is not part of the linked list, the target node of T , $\boxed{4}$, is still reachable from $\boxed{2}$. Thus, T can continue navigation (ignoring the fact that $\boxed{2}$ and $\boxed{3}$ were removed) until it reaches the target node $\boxed{4}$. See the illustration in Figure 1 (left). If $\boxed{4}$ is marked as deleted (or simply does not exist in the search path), the argument become subtle as there might be another node $\boxed{4}$ that is not marked as removed and is reachable from $\boxed{1}$ (but not from $\boxed{2}$). Nonetheless, it is possible to show that returning a “key does not exist” is correct since there exists a time during the execution of the operation when the key did not exist in the list.

However, this mechanism does not work when the hazard-pointer scheme (or a similar scheme) is applied. When $\boxed{2}$ and $\boxed{3}$ are removed from the list, both become candidates for reclamation. As T traverses $\boxed{2}$, there exists a hazard pointer guarding $\boxed{2}$ and it will not be reclaimed. But T does not traverse $\boxed{3}$ when it is removed, hence $\boxed{3}$ is reclaimed; after re-allocation, $\boxed{3}$ ’s content becomes arbitrary. When T continues its traversal to find $\boxed{4}$, it reads $\boxed{3}$ ’s next pointer that returns an arbitrary address (instead of $\boxed{4}$). An attempt to read from the arbitrary address would crash the application.

³Harris’s linked list is discussed in Appendix A. Harris-Herlihy-Shavit linked list was designed for Java, which provides correct memory reclamation without progress guarantees.

Fig. 1. Illustrating why the hazard-pointers scheme cannot be applied directly to Harris-Herlihy-Shavit linked list. The black arrows represent T 's current pointer and hazard pointer, protecting node $\boxed{2}$. Nodes $\boxed{2}$ and $\boxed{3}$ were removed from the list. T continue traversal from $\boxed{2}$ until finding $\boxed{4}$. If the hazard-pointers scheme is used, this optimization cannot be used since $\boxed{3}$ is not protected by a hazard pointer and can be reclaimed.

The Harris-Michael algorithm does not use the wait-free traversal optimization (nor does it use the single-CAS-multiple-removes optimization used by the Harris's algorithm, see Appendix A). When the Harris-Michael algorithm encounters a removed node, it immediately attempts to disconnect the node from the list. If disconnecting fails, the operation is *restarted*. Thus, in the example above, T must restart its traversal from the head pointer. Nevertheless, this variant is compatible with the hazard-pointers scheme.

We note that restarting an operation does *not* generally mean restarting the operation from scratch. For example, if the thread already applied the operation to the data structure, restarting the operation from scratch can lead to re-application of the same operation twice, resulting in an incorrect behavior of the data structure.

2.2.2 Normalized Form of Lock-Free Data Structures. The normalized form of Timnat and Petrank [2014] was originally designed for upgrading lock-free data structures into wait-free ones. This form is used by the AOA scheme for formally defines restarts for lock-free data structures. If a data structure is presented in a normalized form, each operation consists of three parts: generator, executor, and wrap-up. Loosely speaking, the executor is responsible for applying the operation and must not be applied twice, and the generator and wrap-up do not modify the abstract representation of the data structure and can be restarted an arbitrary number of times⁴. The executor has a very rigid structure. It executes a list of CASes (produced by the generator); stops on the first CAS failure or completion of all CASes; and returns the number of successfully executed CASes. The generator and wrap-up are permitted to modify the shared memory (assuming this modification can be applied multiple times without harming correctness) only via the CAS instruction. In addition, the execution of generator and wrap-up must only depend on the operation's input. The wrap-up part can additionally depend on the CAS list executed by the executor and on the number of successfully executed CASes (i.e., input and output of the executor).

The normalized form requires that all modifications to the shared memory be executed by the CAS instruction. In recent years, there has been a growing trend of using the Fetch-And-Add (FAA) instruction for synchronization, in addition to CAS [Basin et al. 2017; Braginsky et al. 2016; Yang and Mellor-Crummey 2016]. Under a high contention, the CAS instruction is likely to “fail”, hence necessitating additional attempts to apply the operation and increasing the contention on the memory bus. In contrast, the FAA instruction always “succeeds” in providing each thread a unique token, which reduces contention. Data structures using FAA cannot be represented natively in the normalized form.

⁴Disconnecting a marked node is an example of a step that modifies the shared memory without changing the abstract representation of the data structure.

2.2.3 Automatic Optimistic-Access Scheme. The AOA scheme was designed to be easily applicable and to support Harris and Harris-Herlihy-Shavit data structures. At runtime, it automatically identifies the nodes that are unreachable and reclaims them, without imposing any kind of synchronization barrier for initiating or finishing reclamation phases. Thus, the AOA scheme reclaims memory, even in the presence of a thread that does not respond during multiple reclamation phases.

The AOA scheme consists of four parts: programmer interface, initiating mechanism, gathering local roots, and tracing (which are discussed next). The Free Access scheme inherits the programmer interface, the initiating mechanism and the tracing algorithm of the AOA scheme; and the reader should consider these background parts as necessary for this paper. In contrast, the Free Access scheme gathers local roots in a completely different manner, which is described in the rest of the paper. The reader can consider the AOA method mainly for the sake of contrasting it with this work.

Programmer Interface. The programmer has to specify the class or struct used in representing data-structure nodes. Specifically, the programmer specifies the size of each node and the offset of pointer fields that point to other data-structure nodes (pointers to other objects are ignored during reclamation). At runtime, the program specifies the size of the memory pool; it also specifies the location of the *global roots* to the data structure that is visible to all threads (e.g., the head pointer in a linked list). Finally, all memory allocations must use a special AOAlloc function. Note that no free-equivalent interface is provided: all memory is reclaimed automatically by the scheme, similarly to GC languages. The Free Access scheme uses exactly the same interface as AOA.

Initiating Mechanism. A memory reclamation *phase* is initiated when the memory available for allocation is exhausted. Any thread in the system can initiate a new phase and many threads can simultaneously attempt to initiate a new phase. A phase is similar to the concept of a GC cycle. But there is a difference: To satisfy lock-freedom, no synchronization barrier is executed at the beginning of a new phase, hence some threads might lag behind and still operate on an older phase.

A new phase starts by signaling all threads. It is desirable to avoid signaling each thread multiple times, even if multiple threads attempt to initiate a new phase simultaneously. Toward this end, each phase has a unique, always increasing, *index*. Each thread has a 8-bit dirty flag that co-resides with a 56-bit index on a single machine word. Signaling a new reclamation phase is done by setting the dirty flags of all threads. This is done exactly once for each thread by using a CAS instruction that sets dirty while increasing the thread's index to the current index. A thread can efficiently check whether a new phase begun by simply checking whether its dirty flag is set. A pseudo code is provided in Listing 2. Recall that this mechanism is also used by Free Access.

Listing 2. Start reclamation phase

```

1  union udirty{
2      atomic<bool> dirty; //fast access to the dirty flag.
3      atomic<uint64_t> dirty_phase; //8-bits dirty, 56-bits phase index
4      pair<bool dirty,uint64_t phase> get()
5          long val = dirty_phase.load();
6          return make_pair(val&1, (val>>8));
7      bool CAS(pair &ex, pair de)
8          return dirty_phase.compare_exchange_strong(ex, de, memory_order_seq_cst);
9  };
10 atomic<uint64_t> phaseIndex;
11 thread_local uint64_t lPhaseIndex;
12 void InitReclamation(){
13     //advance shared phase#; one thread will succeed.
```

```

14  phaseIndex.compare_exchange_strong(IPhaseIndex, IPhaseIndex+1);
15  IPhaseIndex = phaseIndex; //current thread might lag, just read most updated number.
16  foreach(thread T) //Signal all threads by setting the dirty flag
17      val = T.dirty.get(); //read dirty flag and the associated phase index.
18      while(val.phase < IPhaseIndex) //dirty was not set yet in current phase
19          T.dirty.CAS(val, {IPhaseIndex, true}); //set dirty, advance to current phase.
20  }

```

Gathering Local Roots. To gather local roots in a lock-free manner, the AOA scheme requires data-structure operations to be presented in the normalized form of Timnat and Petrank [2014]. Recall from Section 2.2.2 that each such operation consists of three parts: generator, executor, and wrap-up. Loosely speaking, the AOA scheme considers local roots only for the executor part of each operation. The generator and wrap-up can be restarted an arbitrary number of times, hence the thread’s local state is “unimportant”. Before the executor begins, all local pointers are stored in the hazard pointers of the thread. The executor and the wrap-up parts must only depend on nodes in the CAS list (as generated by the generator part) and some additional non-pointer fields. Thus, specifically these nodes are stored in the hazard pointers of the thread. Then, the thread checks, via the dirty flag, whether a new reclamation phase has started. If dirty is set (i.e., a new phase has started), the local state might be unsafe: it is abandoned and the operation restarts from the beginning of the generator. Recall that according to the normalized form structure, the generator can be restarted an arbitrary number of times without affecting the correctness of the data structure. If dirty is clear (i.e., no new phase has started), the thread proceeds to the executor.

The generator and wrap-up parts are considered to have no local roots, hence a memory reclamation phase might reclaim nodes accessed by these parts. Thus, every load from memory is followed immediately by checking the dirty flag. If the dirty flag is set, a new phase has started and the local state might contain pointers to reclaimed nodes. The thread immediately abandons its local state and restarts the current part (generator or wrap-up). Correctness results from the normalized-form property that permits the generator or wrap-up to restart an arbitrary number of times.

The hazard pointers, which are assigned at the beginning of the executor, are kept until the wrap-up is completed. The reason is that the wrap-up might depend on the input to the executor that contains only the nodes protected via hazard pointers. To correctly restart the wrap-up part, these nodes must not be reclaimed while the wrap-up is being executed.

Tracing and Reclamation. Once local roots (hazard pointers of all threads) and global roots (as reported by the programmer interface) are gathered, the AOA algorithm traces these roots and marks all nodes that are reachable from the roots. The AOA algorithm finishes by reclaiming all unmarked nodes for future allocations. The tracing and reclamation algorithm of the AOA scheme is based on a mark&sweep GC with a side markbit table. But it also handles many subtleties that arise when attempting to implement mark&sweep in a lock-free manner. For example, the algorithm must protect against a thread that wakes up after a reclamation phase is finished and that attempts to mark a stale node. The AOA algorithm solves this issue by using a versioning scheme for the side markbit table. Another issue that the AOA algorithm deals with is enabling tracing to continue even if a thread is stuck while it processes a node. Solving this issue requires that the currently processed node is always visible to other threads and a helping mechanism for helping the stuck thread. For the sake of brevity, we do not list here the full details of the AOA algorithm. The Free Access scheme modifies only the way local roots (i.e., hazard pointers) are set and gathered; it uses exactly the same tracing algorithm as AOA, thus avoiding the need to re-consider the above subtleties in our context. A high-level pseudo code for the reclamation algorithm appears in Listing 3.

Fig. 2. Illustrating read-only periods. On ①, the thread is immediately after a write-only period, hence the restart information matches the local variables. PC represents the program counter and ctr some local integer variable. On ②, the thread traversed two nodes in the list. The restart information is unmodified. On ③, concurrent threads delete L and M and invoke a GC. The GC gathers K as a root from the restart information (`RestartInfo`); M is thread-local and inaccessible to the GC. It reclaims L and M as they are unreachable from any root. On ④, the thread restarts using the restart information (`RestartInfo`). The thread state is valid: there are no pointers to reclaimed nodes. The performance benefits are due to the infrequent changes to the root, which reduces the number of fences.

Listing 3. Reclamation

```

1 void Reclamation(){
2   InitReclamation(); //Listing 2 above
3   Roots = gatherLocalRoots(); //For freeaccess, Listing 7 below
4   Roots += gatherGlobalRoots(); //e.g., head of linked list. As reported by the programmer.
5   trace(Roots); //mark nodes reachable from LocalRoots
6   sweep(); //reclaim unmarked nodes
7 }

```

3 OVERVIEW

The main challenge that Free Access addresses is gathering local pointers of all threads in a lock-free manner. According to the definition of lock freedom (Section 2.1), once the memory reclamation-scheme begins reclamation, it must gather the local roots of all threads while the system executes a finite number of steps. This problem is difficult. Gathering local roots seems to follow two paths: cooperative (asking each thread to report its local roots) or uncooperative (without getting any help from the thread). Collecting local roots in a cooperative manner is problematic because it foils the lock-freedom guarantee. Under the lock-freedom model, it is not possible to assume that the thread would execute any instructions before the system executes a finite number of steps (or even before reclamation finishes). But collecting local roots in an uncooperative manner is also problematic as current architectures do not enable a thread to access registers of another thread. Another seemingly possible solution is to prevent the thread from using local roots: all local pointers must always be visible in the shared memory. But such a solution is also very problematic as the thread needs to update the shared copy (and issue an expensive memory fence) anytime a new pointer is loaded, which results in an extremely high overhead.

Free Access solves this problem by partitioning the execution into write-only and read-only periods. For write-only periods, the assumption is that while the thread is busy writing, it does not load new pointers into its local variables. Hence, it is sufficient to publish all local pointers at the beginning of the write-only period. These local pointers are valid throughout the execution of the entire write-only period, thus reducing the publishing overhead.

For read-only periods, the main observation is that such periods can be restarted in a local manner because they never modify the shared state of the program. Returning to the beginning of

the current read-only period requires only loading the local state of the thread at the beginning of the read-only period, without requiring cooperation from other threads. Hence, it is sufficient to consider the state of the thread — and the set of local pointers it uses — only at the beginning of the read-only period. An illustration of this idea appears on Figure 2.

Furthermore, the beginning of a read-only period is actually the end of the previous write-only period. The Free Access scheme publishes the set of local roots at the beginning of the previous write-only period, and we assume that these local roots are not modified during the write-only period. Thus, the set of local roots at the beginning of the read-only period is already visible and no additional publication step is required. Overall, the Free Access scheme never requires the thread to publish its state during read-only periods, thus allowing the system to achieve good performance.

Finally, consider the problem raised at the beginning of the section regarding the hardness of gathering local pointers of a thread. The Free Access scheme solves this problem for read-only periods by using two properties. First, it ensures that the local pointers of the thread *at the beginning of the period* are visible in the shared memory and that they can be gathered. Second, it ensures that the thread restarts its execution from the beginning of the period. These two properties enable threads to use local state without frequent publishing while also enabling the memory reclamation scheme to gather roots in a lock-free manner.

4 THE FREE ACCESS ALGORITHM

In this section, we provide the details of the Free Access scheme for gathering the set of local roots of all threads in a finite system-wide steps. We begin by defining the data structures used by each thread in the system.

4.1 Thread Local Variables

Each thread has a *dirty flag*, two *frames of hazard pointers*, an *arbiter flag*, and a non-pointer *checkpoint*. The *dirty flag* is inherited from the AOA scheme, as mentioned in Section 2.2.3. A dirty flag signifies, when it is set by other threads, that a new reclamation phase started and nodes might be reclaimed. The dirty flag is read and cleared by the thread and is set by other threads (Listing 2).

A hazard-pointers frame is written only by the thread and can be read by other threads. It contains all local pointers during a write-only period. During a read-only period, the hazard-pointer frame is used to store information about the local pointers at the beginning of the read-only period. If a restart is required, the local pointers are restored from the hazard-pointer frame. Each thread has an *arbiter flag* (*ar*), that is used to decide which of the two hazard-pointer frames is currently used. Two frames of hazard pointers are used for enabling the thread to publish the local roots in a new hazard pointers frame without foiling the (still important) information stored in the current hazard pointers frame. The arbiter flag is local to the thread and cannot be accessed by other threads. During a read-only period, the current hazard pointers frame is denoted $HP[ar]$ and the other frame is denoted $HP[!ar]$ (pronounced as *not ar*).

The *checkpoint* stores all local information required for returning to the beginning of the current read-only period, including all local (non-pointers) variables and the program counter. However, it does not contain local pointers, because these are restored from the current hazard pointers frame.

4.2 Write-Only Periods

Beginning a write-only period (and ending the previous read-only period) starts with publishing the set of local pointers of the thread in $HP[!ar]$. This includes writing all local pointers and issuing a memory fence to ensure that these pointers are visible to other threads. Then, the thread checks the dirty flag. If the dirty flag is off, then the thread flips the meaning of the arbiter flag and proceeds to the write-only period. But If the dirty flag is on, a new reclamation phase is started and might have

already gathered the local roots of the thread. Recall that gathering local state is done by reading the hazard pointers of the thread; and these pointers correspond to the state of the thread at the beginning of the previous read-only period. Thus, the thread restarts its execution at the beginning of the previous read-only block. To restart the previous read-only block, the thread replaces the content of its local variables with the content stored in the checkpoint (for non-pointers) and the current hazard-pointer frame (for pointers); and it jumps to the program counter stored at the checkpoint.

When a write-only period ends, the thread stores all of its non-pointer local variables in the *checkpoint*. It does not store the local pointers as they are already stored on HP[ar] (with the new arbiter value). Note that a write-only period does not read pointers from the heap, hence the set of local pointers at the beginning of the write-only period are the same as the local pointers at the end of the write-only period. Therefore, even though the current hazard-pointer frame was written at the beginning of the write-only period, it does not need to be updated.

The pseudo code for the restart mechanism is provided in Listing 4, and the pseudo code for *begin-write-only* and *end-write-only* is in Listing 5.

Listing 4. Restart

```

1 void restart(){
2     helpReclamation(); //help if there is reclamation phase running, as in [Cohen and Petrank 2015a]
3     dirty.store(false, memory_order_relaxed); //no need to restart again, we are already restarting
4     for each  $p$  in local pointers do:
5          $p = \text{HP}[\text{ar}].p$ ;
6     for each  $v$  in local variables do:
7          $v = \text{checkpoint.v}$ ;
8     goto checkpoint.program_counter;
9 }
```

Listing 5. *begin-write-only* and *end-write-only*

```

1 void beginWriteOnly(){ //=endReadOnly
2     For each  $p$  in local pointers do:
3          $\text{HP}[\text{!ar}].p = p$ ;
4     atomic_thread_fence(memory_order_seq_cst); //read dirty after setting HPs
5     if(dirty.load(memory_order_acquire))
6         restart();
7      $\text{ar} = \text{!ar}$ ;
8 }
9 void endWriteOnly(){ //=beginReadOnly
10    for each  $v$  in local variables do:
11        checkpoint.v =  $v$ ;
12    checkpoint.program_counter = &PC; //PC is the program counter of the next instruction
13 }
```

4.3 Read-Only Periods

If a reclamation phase starts when a thread executes a read-only period and the thread does not respond, the Free Access scheme gathers the set of roots published in the hazard-pointer frame of the thread. This represents the set of roots at the beginning of the read-only period of the unresponsive thread. To ensure that the behavior of the memory-reclamation algorithm is correct, the unresponsive thread must restart at the beginning of the read-only period. Furthermore, it must not execute any instructions that are visible to other threads before the restart. Therefore, the

system behaves as if the thread never started the read-only period for the entire duration that its roots were gathered; in which case, it is clearly correct to gather the set of roots at the beginning of the read-only period. An illustration of this mechanism is provided in Figure 2.

Ensuring that the thread is not visible until it is restarted requires that the thread does not write to the shared memory; this follows immediately from the definition of read-only periods. However, it is also required that the thread behaves “reasonably” so that it does not cause segmentation faults or other runtime errors that are visible, even though the thread does not explicitly write to the shared memory [Dice et al. 2014]. As memory reclamation does not wait for the thread to respond and does not gather the set of current local pointers of the thread, it is possible that a currently accessed node is reclaimed. In general, this can lead to an arbitrary behavior of the thread and cause the thread to be visible before it manages to restart.

A core invariant maintained by the Free Access scheme is that the content of a reclaimed node is never used. To understand how a reader thread T maintains this invariant, we first consider an abstract view of how reclamation works. Reclamation starts by setting T ’s dirty flag to true (Listing 2 Line 8) and then reclaims some nodes and sets their content to arbitrary. To ensure that T never uses an arbitrary value, every read by T is followed by reading the dirty flag. Only if the dirty flag is clear, is the read value used. Otherwise, T restarts the operation, using the algorithm in Listing 4. As reclamation sets the content to arbitrary only after setting T ’s dirty flag, and as T uses the read value only if the dirty flag is off after reading the content, then T never uses the arbitrary content.

There are two ordering constraints that a reader thread must maintain. First, it must read the dirty flag after reading the node content. Second, it must use the read node content after checking that the dirty flag is clear. Ensuring that the compiler and the CPU do not reorder the reads requires the use of fences. Specifically, the first ordering constraint is satisfied by using an acquire fence after reading the node content and before reading the dirty flag. Correctness is similar to [Boehm 2012]. The second ordering constraint is satisfied by using an additional acquire fence when reading the dirty flag. The validation code appears in Listing 6 (the macro `VALIDATE_READ`).

A simple way to maintain the read invariant is to execute the validation code immediately after every read (thus, before using the read value). But in some cases, a node has multiple fields that are read simultaneously and then used. In such a case, the validation code can be placed once, after reading both fields and before using both values. The invariant is still maintained as these two reads are independent and as each read is validated before the read value is used. An example code for applying read validation also appears in Listing 6.

Listing 6. Applying Free Access to read-only period

```

1  #define VALIDATE_READ  atomic_thread_fence(memory_order_acquire); \
2                        if(dirty.load(memory_order_acquire)) restart();
3      q=p.load(memory_order_relaxed);
4      VALIDATE_READ; //validate q is not an arbitrary value
5      v1 = q.load(memory_order_relaxed); //use read value, also another load from memory.
6      v2 = p.load(memory_order_relaxed); //another independent load
7      VALIDATE_READ; //validate v1 and v2 are not arbitrary values
8      //use v1, v2.

```

4.4 Starting and Ending a Data-Structure Operation

At the beginning of a data-structure operation, the thread’s local state is initialized as follows. The arbiter flag is set to zero. To avoid publishing an uninitialized pointer, local pointers are initialized to `NULL`. If a data-structure operation starts with a read-only period, as is common in many

algorithms, a dummy write-only period is executed at the beginning of the function to ensure that the read-only period can be restarted. In this case, the checkpoint of the thread consists of the input to the function and of the program counter at the beginning of the function. If the input to the function does not contain any pointer to a data-structure node, there is no need to write any hazard pointer and the memory fence can be elided.

Finishing an operation should nullify all hazard-pointer fields to avoid *floating garbage*: memory that cannot be freed even though it is not used by any data-structure operation. However, the Free Access scheme uses two sets of hazard pointers, hence nullifying all of them is inefficient. Instead, the thread writes a special marker (-1) to the first hazard pointer in HP[1]. This marker signifies that all other pointers are implicitly NULL. When the next operation writes to the hazard pointers, it always starts with the first hazard pointer at HP[1] (recall that the arbiter flag is initially zero and is swapped before setting hazard pointers), thus removing the implicit marker.

4.5 Gathering Local Roots

When a reclamation phase starts, the Free Access scheme gathers, by reading the hazard pointers of all threads, their local roots. There are two subtleties to be considered. First, if the first hazard pointer in HP[1] of a thread contains the special marker (-1), all other hazard pointers of this thread are ignored. Second, if the data structure uses tagged pointers (i.e., modifies some bits of pointer fields), these tagged pointers are untagged when gathering the local roots.

We remark that the AOA scheme converts all tagged pointers into untagged pointers before setting them into the thread's hazard pointers. Thus, gathering the local roots is trivial in the AOA scheme. This is not possible in the Free Access case, because hazard pointers are used for restarts. If a pointer is tagged when executing a read-only period and the read-only period requires a restart, the pointer must remain tagged. Hence, the pointer is stored to the hazard-pointer frame in its tagged form and is untagged during local-root gathering. In fact, clearing the tag while gathering the roots provides a minor performance benefit as gathering roots occurs less frequently than setting hazard pointers does. The pseudo code for root gathering is provided in Listing 7.

Listing 7. Gathering local roots

```

1  set<node*> gatherLocalRoots(){
2      set<node*> Roots;
3      foreach( thread T )
4          if(T.HP[1].first == MARKER) continue; //continue with next thread
5          for( ar ∈ {0,1}, p in local pointers )
6              Roots += clearTagging(T.HP[ar].p);
7      return Roots;
8  }
```

4.6 Automatically Applying Free Access Modifications

Although the modifications mentioned above do not require special algorithmic properties, implementing them manually is a tedious effort, risks bugs, and complicates the original data-structure code. To simplify this process, we implement an LLVM compiler pass that automatically applies all these modifications. The programmer needs only to add the `freeaccess` attribute to each data-structure operation. To avoid modifications to the clang front-end, we use the general `annotate` attribute, that enables the programmer to specify an arbitrary string as an attribute. Thus, the programmer simply adds `__attribute__((annotate("freeaccess")))` to each data-structure operation.

The compiler pass starts by parsing the `llvm-ir` annotation header and by extracting the set of functions with the `freeaccess` attribute. Then, for each of these functions, it executes the following

four steps. First, the compiler pass extracts the set of local variables. For each non-pointer variable, it adds an additional local variable that serves as the checkpoint value for original variable. It also adds the arbiter flag and the program-counter (used for restarting) as two additional local variables and inserts code to initialize them. Lastly, Initialization sets all local pointers to NULL. Before returning from the operation, the compiler pass inserts a code to set the special marker (-1) to the first entry in HP[1], which implicitly nullifies all hazard pointers.

Second, the compiler pass creates a new basic block for restarting the operation, denoted *restart basic block*. The code of the restart basic block closely follows the code in Listing 4.

Third, the compiler pass considers reads from the shared memory. It begins by analyzing consecutive reads to find if they are independent. If they are, the first one is discarded, because checking the dirty flag after the second one is sufficient to protect the first one. Then, for each remaining read, the compiler pass inserts a check to the dirty flag and a conditional branch to the restart basic block.

Fourth, the compiler pass handles writes (and CASes) to the shared memory. It analyzes the code and identifies write-only regions. Then, it inserts code for starting a write-only region (*begin-write-only*) at the beginning of the region, and it inserts code for ending a write-only region (*end-write-only*) at the end of the regions. The inserted code closely follows Listing 5.

The current version of the plug-in implements the hazard-pointers set as an array of 15 elements, which forms 2 cache lines, together with the dirty flag. This supports up to 7 pointers for each data-structure operation. We carefully coded the plug-in to ensure that if the pointer to the dirty flag and the arbiter are stored in registers (as in the code produced by the assembler we used), then checking the dirty flag or accessing a hazard pointer is done in a single x86 instruction⁵.

If an operation is read only (as determined by LLVM's `functionattrs` pass), our compiler pass produces an optimized version. As there are no writes, restart is done by jumping to the beginning of the operation, using a non-conditional jump instead of an indirect jump. Hence, there is no need for the restart program-counter. Furthermore, hazard pointers are never accessed, so the arbiter flag is also omitted.

Our LLVM compiler pass is scheduled to run before other optimization passes, enabling it to produce unoptimized code that is later heavily optimized by other compiler passes. Except for the two optimizations mentioned above, the implementation of the compiler pass is quite straightforward and it consists of less than 600 lines of C++ code.

5 MEASUREMENTS

We incorporated the Free Access scheme with the AOA implementation⁶. We measured the effectiveness of our system using the linked-list implementation by Herlihy and Shavit [2012] and a hash-map implementation based on this list. The data-structure code was modified to replace all AOA modifications with the `freeaccess` attribute and run the Free Access compiler pass. The resulting implementation is denoted FA.

The FA implementation is compared against several other reclamation schemes. The baseline for comparison is a leaky implementation that never reclaims memory and incurs no memory reclamation overhead; it is denoted NR. Typically, the leaky implementation cannot be used for real applications. The comparison of FA and NR represents the full cost of adding Free Access to a data structure that did not previously support memory reclamation. NR allocates all memory before execution begins so that memory allocation from the OS does not become the bottleneck. The

⁵We slightly changed the meaning of the `ar` flag in the code: instead of being either 0 or 1, it is either 1 or 9, so the address of the i -th hazard pointer in the current frame (HP[ar]) is $((\text{node} \text{ **})\text{dirty}) + \text{ar} + i$. Flipping `ar` is done by XORing it with 8.

⁶Code was taken from github.com/nachshonc/AutomaticOptimisticAccess

second comparison is against the AOA scheme. As Free Access is based on the AOA scheme but provides further simplicity, a comparison of FA and AOA reveals the direct cost of the additional simplicity. We also consider the lock-free reference-counting scheme [Michael and Scott 1995; Valois 1995], which is also applicable for general code but suffers from high overhead; it is denoted RC. These are the only schemes we are aware of that guarantee strict lock freedom and support the Harris-Herlihy-Shavit linked list [Herlihy and Shavit 2012].

For completeness, we also compare FA to the hazard-pointers scheme; the latter is applied to the Harris-Michael linked list [Michael 2002b] (Section 2.2.1 discusses the main difference from the Harris-Herlihy-Shavit algorithm); it is denoted HP. Finally, we also include a version of hazard pointers that uses fence elision [Dice et al. 2016] that replaces frequent memory fences during data-structure operations by the `mmapbarrier` Linux system call during reclamation. As the `mmapbarrier` system call cannot return until it interacts with all other threads in the system, we do not consider it to be strictly lock-free. Still, it guarantees lock-freedom, except for the `mmapbarrier` syscall, and it is considered a reasonable tradeoff in practice. Similarly to the HP implementation, this version is applied to the Harris-Michael linked list because it cannot be applied to the Harris-Herlihy-Shavit linked list; it is denoted HPMB.

5.1 Methodology

For the linked-list data structure, we consider two configurations, LL5K and LL128; they differ in the number of keys in the list. In the LL5K configuration, the range is set to 10K so that the average size of the list is 5K. This configuration frequently reads from the shared memory and exhibits low contention. In the LL128 configuration, the range is set to 256 so that the average size of the list is 128. The contention is significantly higher and frequency of memory reclamation increases. For the hash set, the average table-size is 10K and the load factor is 1. The operations are extremely fast, hence the frequency of memory reclamation increases further. However, contention is relatively low as different threads operate on different buckets.

We ran the experiments on a machine featuring 4 AMD Opteron (TM) 6376 2.3GHz processors, each with 16 threads (overall 64 hardware threads). We varied the number of threads between 1 and 64, in power-of-two jumps to check the behavior of the proposed schemes under different scalability settings. The operation distribution was fixed to 50% contains and 50% modifications, equally distributed between inserts and removes. This models a write-heavy distribution that stresses memory reclamation. Each execution runs for 1 second and each test was repeated 10 times. We depicted the ratio between the average throughput and the throughput of the NR implementation and the 95% confidence intervals. The code was compiled using LLVM version 5.0.0 with the `-O3` optimization flag, running on an Ubuntu 16.04 OS (kernel version 4.4.0-72).

For the Free Access and the AOA schemes, before the measurements began we initialized the allocation pool with 50,000 nodes; a new collection phase is triggered when the allocation pool is exhausted. For HP and HPMB implementations, reclamation is triggered locally by a thread once its local list of retired nodes (recall Section 2.2.1) reaches a threshold; this is different than with the FA and AOA schemes, where reclamation is triggered globally. To ensure a fair comparison, we set the local reclamation threshold to 100,000 nodes divided by the number of threads, which approximates a slightly lower reclamation rate than the FA scheme⁷.

For the HPMB implementation, the `mmapbarrier` syscall is invoked with the `MEMBARRIER_CMD_PRIVATE_EXPEDITED` parameter that was recently introduced to Linux. This parameter provides a significant performance boost to HPMB compared to the older variant of the `mmapbarrier` syscall.

⁷Setting reclamation to 50,000 nodes divided by the number of threads ensures similar maximum memory size but a higher average reclamation rate. We verified that a lower reclamation rate improves performance.

Fig. 3. Comparing the NR, AOA, FA, RC, HP, and HPMB memory reclamation schemes for a linked list with 5,000 elements. For each scheme, we depict the throughput ratio between the throughput of the depicted scheme and the leaky implementation (NR).

Fig. 4. Comparing the NR, AOA, FA, RC, HP, and HPMB schemes for a linked list with 128 elements.

5.2 Results

The results for the linked list with 5,000 elements are presented in Figure 3. Each operation is relatively long (around $20\mu s$) and most of the time is spent on reads. Compared to the leaky implementation, FA has an overhead of 0.1% – 1.4%. Compared to the AOA scheme, FA is faster by 1.3% – 2.5%. We attribute the difference to implementing the Free Access scheme at the LLVM IR level (Section 4.6) that generates better code than the manually modified version. In this case, simpler is strictly better. FA is significantly faster than RC: over 5x for a single thread and more than 10x for 2 threads and above. It is well known that the overhead of RC is extremely high; still, this was the only available memory-reclamation scheme that could be applied to any data structure (without garbage cycles). The Free Access scheme is clearly superior to RC: it can be applied automatically (including garbage cycles), easily (using our compiler plug-in), and provides significantly better performance. Due to this huge overhead, we ignore RC throughout the rest of this section.

The HP implementation performs slower than FA by 56% – 72%. This is mainly the cost of the memory fences required during navigation. Finally, the HPMB is slightly slower than FA by

Fig. 5. Comparing the NR, AOA, FA, RC, HP, and HPMB schemes for a hash table with 10,000 elements.

around 1% up to 16 threads. For 32 threads and above, the overhead of HPMB increases to 3% – 13%. The main reason for this overhead is the baseline data structure: HPMB is applied to Harris-Michael linked list data structure, which disables the wait-free navigation optimization discussed in Section 2.2.1. Note that without this optimization, contain operations are more heavyweight as they need to synchronize with removal operations and might need to restart.

The results for the linked list with 128 are presented in Figure 4. The operations are faster, but contention is significantly higher: a single operation takes around 230ns for the single threaded case and around 7.5μs for the 64-thread case. The overhead of the FA scheme is 12% for the single threaded case and less than 4% for 2 threads and above. Recall that memory reclamation efforts are mostly thread-local. Consequently, when contention on the list grows, the overhead of these local operations decreases. FA is faster than the AOA scheme by 2% – 7%. The HP implementation is slower than FA by 6x for the single threaded case, but this overhead reduces when contention on the list increases and reaches only 11% for 64 threads. Comparing FA and HPMB, we see that FA performs better for 1 – 8 threads by 3% – 16%, but slower when the number of threads is 16 or higher. Interestingly, for 32 and 64 threads, HPMB performs even faster than NR by up to 8%. Profiling shows that part of this benefit is due to quickly reusing reclaimed nodes, while they are still in the processor cache. The reader should recall that HPMB calls the `membarrier` method, which waits until all threads execute a memory fence. As a result, HPMB does not guarantee lock-freedom progress. If the programmer is willing to compromise for the progress guarantee, Free Access trades off some performance for the sake of simplicity.

The results for the hash table are presented in Figure 5. Hash operations are extremely short: on average, an operation visits only a single node; still, contention is low as each operation uses a different bucket. The overhead of the FA scheme (compared to NR) is 16% – 44% and increases for a high-thread count. Profiling shows that the main overhead comes from the tracing algorithm; because operations are extremely short, the overhead of tracing cannot be hidden. Furthermore, tracing is not as scalable as hash-map operations, which explains the increased overhead when the number of threads increases. Comparing FA to AOA, we find the FA scheme is faster by 3% – 6%. Most overhead comes from the tracing algorithm that is shared by both implementations. The overhead of the HP implementation is higher than FA by 22% for a single-threaded case, similar to FA for 32 threads, and lower than FA by 9% for 64 threads. The higher throughput at the 64-threaded case again shows the benefits of manually reclaiming memory, also for the strictly lock-free version of the hazard-pointers scheme. There is no need to trace the data structure at runtime to gather

reclaimed nodes; and reclamation decisions can be done in a local manner. By avoiding frequent memory fences, the HPMB version provides further performance improvements over HP. Up to 32 threads, FA performs similar to HPMB: between 8% faster and 3% slower. But for 64 threads, HPMB outperforms FA by 15%. Again, the comparison against the hazard-pointers scheme represents a tradeoff. The Free Access scheme is the simplest to apply and provides good performance for a low-thread count; but when operations are extremely short it suffers from reduced scalability. The HP implementation is hard to apply, but provides some performance benefits for 64 threads. The HPMB implementation compromises the lock-freedom guarantee, in addition to being hard to apply; but it provides best performance for 64 threads.

Overall, the overhead of simplicity is quite low. The Free Access scheme consistently outperforms the AOA scheme, even though FA is significantly simpler to apply and does not require the data structure to be presented in a normalized form. Compared to the hazard-pointers scheme, the Free Access scheme incurs some overhead for the 64 threads case. In this case, the Free Access scheme makes a tradeoff similar to garbage collected systems: it pays some cost at runtime, for the sake of simplicity for the programmer. In other cases, where the number of threads is lower or operations are longer, the Free Access scheme provides better performance by avoiding expensive memory fences and supporting the faster Harris-Herlihy-Shavit data structure. Finally, HPMB compromises the lock-freedom guarantee by avoiding frequent memory-fences. For a low-thread count, it significantly outperforms HP, but FA still provides better performance. For the 64-threads case, HPMB outperforms FA by a wider margin (15% for the hash configuration) but remains hard to apply as HP.

6 ASSUMPTION: SEPARATING WRITE-ONLY AND READ-ONLY PERIODS

A central assumption of the Free Access algorithm is the possibility to partition the execution of a program into write-only and read-only periods. Read-only periods must not at all modify the shared memory. But write-only periods restrict only the reads of new pointers that have a special meaning for root gathering. Write-only periods are permitted to read any non-pointer field from the shared memory. Furthermore, many lock-free data structures use the CAS instruction that writes a new value to the shared memory, but only if the current value matches the expected value. Typically, these data structures use the boolean version of CAS, which returns a boolean indicator whether the expected value matches and the new value is written⁸. This instruction is still considered write-only as it does not load a new pointer to the local state of the thread (hence can be ignored by root gathering). All the lock-free data structures we are aware of in the literature modify the shared memory, either by a direct store, the boolean version of CAS, or by fetch-and-add (FAA) on an integer field. Thus, the Free Access scheme is applicable to all these cases.

Current architectures, however, also support instructions that can modify the shared state and simultaneously load a pointer. Consider the SWAP instruction (x86 xchg instruction) that swaps the content of a memory location and a register. If SWAP is applied to an integer field, it is considered as a write-only instruction. But if it is applied to a pointer field, then this instruction both writes a memory location and reads a pointer to a register. This usage cannot be partitioned into either write-only or read-only period and is not supported by the Free Access scheme. This instruction cannot be restarted because the SWAP can be visible twice to other threads. But if the thread is stuck immediately after executing the SWAP, it is not possible to discover the read pointer that resides in a register that is inaccessible by other threads.

Although the lack of support for SWAP is a limitation of the Free Access scheme, it seems to be a general limitation for lock-free data structures in general. The problem is similar: If a lock-free data

⁸Java only supports the boolean version of CAS, which can be implemented using LoadLinked/StoreConditional primitives.

structure uses a SWAP and the thread executing the SWAP is stuck, it is not possible to recover the original node. We are also not aware of any lock-free reclamation schemes that handle the SWAP instruction: the hazard pointers and similar schemes cannot be applied as they need to restart reading a pointer. Recall that restarting a SWAP instruction that reads a pointer might foil the correctness of the data structure. The lock-free reference counting cannot be applied as it cannot SWAP and increment the reference count of the target node in a single operation. The AOA scheme cannot be applied as SWAP cannot be represented in the normalized form.

Finally, the Free Access scheme does not support a data structure that uses the SWAP instruction, whereas, it is easy to modify the data structure — in a completely mechanical way — so that Free Access can be applied. This can be done by replacing the SWAP instruction with a loop of (boolean) CASes, as illustrated on Listing 8.

Listing 8. Emulate SWAP

```

1 Node *SWAP(atomic<Node*> *addr, Node *reg){
2     do{
3         Node *cur = addr->load(memory_order_relaxed);
4     }while(!addr->compare_exchange_strong(cur, reg));
5     return cur;
6 }
```

We believe this solution is not ideal as it can reduce the progress guarantee of the modified data structure from wait-freedom to lock-freedom (a thread can be delayed an arbitrary number of steps under a worst-case scheduling) and it might reduce performance. Still, as no such lock-free data structure is currently available and no better solution is known, we consider this to be an acceptable solution. If SWAP becomes common in the future, it would be interesting to study how it could overcome the difficulty of a stuck thread and whether Free Access could be modified to handle this usage, without foiling the wait-freedom guarantee of the data structure.

Additional instructions that do not satisfy the Free Access assumption are (i) the CAS version that returns the old value of the field and (ii) the FAA instruction on a field that smartly encodes a counter and a pointer in a single 64-bit word. CASes are typically used to modify the shared memory in a synchronous manner, hence the old value of the CAS is typically not important to the data-structure algorithm. Still, if a data structure uses the old value returned by a CAS or uses FAA on a pointer field, Free Access can be applied if the unsupported instruction is replaced by a CAS loop.

7 RELATED WORK

In this section, we discuss existing memory-reclamation schemes for non-blocking data structures. We consider three aspects of each of these schemes: performance, non-blocking guarantees, and the simplicity of applying them to a given data structure. A table summarizing these schemes appears in Table 1.

The first memory reclamation scheme for lock-free data structures was based on reference counting [Michael and Scott 1995; Valois 1995]. It is the only scheme that can be applied to general code (under the assumption of Section 6). It offers full lock-freedom support. The main problem of the reference-counting scheme is performance: it adds significant overhead and prevents scalability [Hart et al. 2007]. The reclamation of cyclic garbage is another limitation.

Harris [2001] designed the epoch-based reclamation (EBR) method based on the RCU paradigm. The EBR scheme permits threads to have different epoch numbers, but the maximum difference between all threads in the system is at most one. Threads reclaim nodes that were retired at least two epochs ago, thus ensuring that all threads finish executing all operations on that epoch. EBR is

Table 1. Existing lock-free memory reclamation schemes

Scheme	Based on	Performance	Simplicity	Lock Freedom Dependences if exist
RC [V95,MS95]		Very slow	Automatic (§6)	Yes
EBR [H01]	RCU	Fast	Simple	No
HP [M02a,M04]	RC	Slow	Complex (§2.2.1)	Yes
DTA [BKP13]	HP	Fast (depends)	Very complex	Yes
StackTrace [ALMS15]	HP	Fast	Complex (§2.2.1)	Partial. Wait for threads to respond. Depends on OS signals.
OA [CP15b]	HP	Fast	Complex (§2.2.1)	Yes
DEBRA+ [B15]	EBR	Fast	Complex	Partial. Depends on OS signals and stack unwinding (longjmp).
AOA [CP15a]	OA M&S GC	Fast	Simple Normalized form	Yes
QSense [BGHZ16]	HP, EBR	Fast	Complex (§2.2.1)	Partial. Assumes threads interact in fixed time. Depends on OS scheduler.
FenceElision [DHK16]	HP	Fast	Complex (§2.2.1)	Partial. Wait for threads to respond. Depends on OS interrupts.
HazardEras [RC17]	HP, EBR	Fast	Complex (§2.2.1)	Yes
IBR [WICBS18]	HP, EBR	Fast	Complex (§2.2.1)	Yes
ForkScan [ALMS17]	StackTrace OnTheFly GC	Fast	Simple Programmer hints	Partial. Wait for threads to respond. Depends on OS signals. Copy-on-write.
FreeAccess (this paper)	AOA	Fast	Automatic (§6)	Yes

not lock-free, because if a thread does not advance its epoch number, other threads in the system cannot advance the global epoch number, thus preventing reclamation. Eventually, the system will run out of memory. EBR is easy to apply, but still requires the programmer to insert retire calls, which make the node a candidate for reclamation.

The hazard-pointers (HP) scheme [Michael 2002a, 2004] was presented in Section 2.2.1. HP provides full lock-freedom support: threads never wait for other threads to respond. However, it often suffers from a high overhead (Section 5). In addition, applying HP is complex; it cannot be applied to standard lock-free data structure such as Harris’s linked list [Michael 2002a]. Instead, it requires changing the data structure, reducing its scalability and performance, and changing its properties (Section 2.2.1).

The drop the anchor (DTA) scheme [Braginsky et al. 2013] was designed to improve the performance of the hazard-pointers scheme for linked-list-based data structures. DTA publishes a single anchor node that implicitly publishes a set of k nodes reachable from the anchor. A freezing mechanism is used to recover a delayed thread. DTA offers full lock-freedom support. It is faster than the HP scheme on long linked lists but suffers from scalability issues on smaller lists [Cohen and Petrank 2015b]. It is required to re-design a drop-the-anchor scheme for each data structure, and the authors presented such a design only for the linked-list data structure.

ThreadScan [Alistarh et al. 2015] reduces the overhead of the HP scheme by using OS signals. During a reclamation, a signal is sent to all threads in the system; each thread traces its own roots and reports them to the reclamation process. The restrictions posed by ThreadScan on the data structures are the same as these posed by the HP scheme. ThreadScan cannot be considered fully lock-free as it must wait for all threads to respond before proceeding with the reclamation.

DEBRA+ [Brown 2015] extends the EBR scheme to tolerate threads failures. If a thread is delayed and fails to progress its epoch number, other threads naturalize the delayed thread by sending it an OS signal. When a naturalized thread wakes up, it must restart the current operation. The algorithm never waits for a thread: once a thread is naturalized, other threads can continue executing their operations. Nevertheless, DEBRA+ uses the OS signaling mechanism, which acquires locks in the Linux implementation. Applying a DEBRA+ scheme to a data structure requires to re-design the data structure and develop a helping mechanism for naturalized threads.

QSense [Balmau et al. 2016] combines the HP and EBR schemes: When all threads respond, it uses EBR; when a thread fails to respond, it falls back to using HP. To avoid the cost of memory barriers at the fast path, it uses rooster processes that the OS scheduler is required to run at fixed intervals. There is one rooster process per worker thread and it runs on the same processor core as the thread. Its task is to flush, at these intervals, the write buffer of the processor core. QSense does not support strict lock freedom as it assumes that threads respond within a fixed time interval and relies on the lock-freedom of the OS scheduler. It poses the same restrictions on the data structure as the HP scheme does.

Fence eliding techniques for the HP scheme were considered by Dice et al. [2016]. The basic idea is to avoid publishing the hazard pointers during the data-structure operation. Only during reclamation is a global synchronization barrier executed, which implies a memory fence of each of the participating threads. They offer three such mechanisms. First, they use OS services to execute a global-synchronization barrier. Such an OS service does not purely adhere to the lock-freedom model as it assumes that a thread can be interrupted in a fixed amount of time. The Linux implementation takes locks to execute this operation. A second mechanism relies on a x86-specific hardware global-synchronization barrier; this barrier remains only to support legacy code. Such hardware support is not available on other architectures and it might not be supported on future x86 processors. A third mechanism is a hardware-assisted hazard-pointers mechanism. All these mechanisms pose the same restrictions on the data structure as the HP scheme does.

The Optimistic Access scheme [Cohen and Petrank 2015b] improves the hazard-pointers scheme by replacing publishing hazard pointers, which requires a memory fence, by checking a dirty flag. The handling of read instructions in this paper is based on this mechanism. If the dirty flag is found to be set, the thread restarts the operation. They also concisely define restarting an operation by using the normalized form of Timnat and Petrank [2014]. The Optimistic Access scheme poses the same restrictions on the data structure as the HP scheme does.

Hazard Eras [Ramalheite and Correia 2017] and Interval-Based Reclamation [Wen et al. 2018] use a combination of epoch-based reclamation and hazard pointers schemes. Each thread publishes the current epoch it observes. These schemes do not reclaim any node that was alive during a published epoch, but can reclaim nodes removed before such an epoch or allocated after such an epoch. Both these schemes pose the same restrictions on the data structure as the HP scheme does.

Alistarh et al. [Alistarh et al. 2017] presented ForkScan, a garbage collection scheme that uses programmer hints to achieve scalability. These hints are checked by the garbage-collection system, hence the programmer does not have to restrict itself in the data structure design. Similar to on-the-fly garbage-collection algorithms, ForkScan relies heavily on OS services, both to interrupt threads in order to gather roots and to provide a snapshot of the heap, hence it cannot be considered fully lock-free.

The Free Access scheme relies on the ability to restart read-only periods. This technique is used by Java's StampLock [Lea and JSR-166 2013] for synchronizing optimistic readers and was later used by Layout Lock [Cohen et al. 2017] for synchronizing structural changes to a data structure.

8 CONCLUSION

In this paper, we presents Free Access, a lock-free memory-reclamation scheme that is applicable to general data-structure code. The Free Access scheme does not require the data structure to be presented in a special normalized-form or support a special restart mechanism. Furthermore, applying Free Access is easy because code modifications are done automatically by a compiler plug-in. This enables a separation of concerns: Memory reclamation does not prevent designs of novel lock free data structures and does not prevent a practitioner from applying these data structures in real applications. The overhead of the Free Access scheme is surprisingly low: it always outperforms the automatic optimistic access scheme and in many cases also the hazard-pointers scheme (with and without fence elision).

A HARDNESS OF APPLYING THE HAZARD-POINTERS SCHEME

The Harris's linked list is incompatible with the hazard-pointers scheme (as reported by the original hazard pointers paper [Michael 2002a]). The reason is similar to the one described in Section 2.2.1, but the details are somewhat different and provided here for completeness. Consider a snippet of a linked list with four nodes $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 4$ and suppose that both 2 and 3 were marked as removed from the list, but were not yet disconnected from the list (so 1 's next pointer still points to 2). In such a case, it clearly make sense to remove both nodes (2 and 3) by using a single CAS instruction that modifies 1 's next pointer to 4 . This optimization is incorporated in Harris's linked-list algorithm: when a marked node is found, the algorithm attempts to find additional sequentially marked nodes and to disconnect all these nodes together.

Next, suppose that no memory reclamation is used, and that a thread T currently traverses node 2 . As 2 is marked, T would read 2 's next pointer (3) and find that it is also marked. Then it would read 3 's next pointer (presumably 4) and find (by reading 4 's content) it is not marked. Finally, T would attempt to CAS 1 's next pointer from 2 to 4 .

This technique does not work when the hazard-pointer scheme is used. When T traverses 2 , it only protects 2 via a hazard pointer. A concurrent thread might have disconnected both 2 and 3 from the list and retire both. Since only 2 is guarded by a hazard pointer, 3 would be reclaimed and its content becomes arbitrary. When T resumes its operation, it would read 3 's next pointer, but obtain an arbitrary address (instead of 4). Then, it would crash when checking, which is a read from an arbitrary address, whether the node is marked.

REFERENCES

- Dan Alistarh, William Leiserson, Alexander Matveev, and Nir Shavit. 2017. Conservative Memory Reclamation for Modern Operating Systems. In *Proc. Twelfth Eur. Conf. Comput. Syst. - EuroSys '17*. ACM Press, New York, New York, USA, 483–498. <https://doi.org/10.1145/3064176.3064214>
- Dan Alistarh, William M. Leiserson, Alexander Matveev, and Nir Shavit. 2015. ThreadScan: Automatic and Scalable Memory Reclamation. In *Proc. 27th ACM Symp. Parallelism Algorithms Archit. - SPAA '15*. ACM Press, New York, New York, USA, 123–132. <https://doi.org/10.1145/2755573.2755600>
- Joshua Auerbach, David F. Bacon, Perry Cheng, David Grove, Ben Biron, Charlie Gracie, Bill McCloskey, Aleksandar Micic, and Ryan Sciampanone. 2008. Tax-and-Spend: Democratic Scheduling for Real-Time Garbage Collection. In *EMSOFT*. 245–254.
- Oana Balmau, Rachid Guerraoui, Maurice Herlihy, and Igor Zablotchi. 2016. Fast and Robust Memory Reclamation for Concurrent Data Structures. In *Proc. 28th ACM Symp. Parallelism Algorithms Archit. - SPAA '16*. ACM Press, New York, New York, USA, 349–359. <https://doi.org/10.1145/2935764.2935790>
- Dmitry Basin, Edward Bortnikov, Anastasia Braginsky, Guy Golan-Gueta, Eshcar Hillel, Idit Keidar, and Moshe Sulamy. 2017. KiWi: A Key-Value Map for Scalable Real-Time Analytics. In *Proc. 22nd ACM SIGPLAN Symp. Princ. Pract. Parallel Program. - PPOPP '17*. ACM Press, New York, New York, USA, 357–369. <https://doi.org/10.1145/3018743.3018761>

- Hans Boehm. 2012. Can seqlocks get along with programming language memory models?. In *Proc. 2012 ACM SIGPLAN Work. Mem. Syst. Perform. Correctness - MSPC '12*. ACM Press, New York, New York, USA, 12. <https://doi.org/10.1145/2247684.2247688>
- Anastasia Braginsky, Nachshon Cohen, and Erez Petrank. 2016. CBPQ: High Performance Lock-Free Priority Queue. In *22nd Int. Eur. Conf. Parallel Distrib. Comput.* Springer, Cham, 460–474. https://doi.org/10.1007/978-3-319-43659-3_34
- Anastasia Braginsky, Alex Kogan, and Erez Petrank. 2013. Drop the anchor: lightweight memory management for non-blocking data structures. In *Proc. twenty-fifth Annu. ACM Symp. Parallelism algorithms Archit.* ACM, 33–42.
- Trevor Alexander Brown. 2015. Reclaiming Memory for Lock-Free Data Structures. In *Proc. 2015 ACM Symp. Princ. Distrib. Comput. - Pod. '15*. ACM Press, New York, New York, USA, 261–270. <https://doi.org/10.1145/2767386.2767436>
- Nachshon Cohen and Erez Petrank. 2015a. Automatic memory reclamation for lock-free data structures. In *Proc. 2015 ACM SIGPLAN Int. Conf. Object-Oriented Program. Syst. Lang. Appl. - OOPSLA 2015*, Vol. 50. ACM Press, New York, New York, USA, 260–279. <https://doi.org/10.1145/2814270.2814298>
- Nachshon Cohen and Erez Petrank. 2015b. Efficient Memory Management for Lock-Free Data Structures with Optimistic Access. In *Proc. 27th ACM Symp. Parallelism Algorithms Archit. - SPAA '15*. ACM Press, New York, New York, USA, 254–263. <https://doi.org/10.1145/2755573.2755579>
- Nachshon Cohen, Arie Tal, Erez Petrank, Nachshon Cohen, Arie Tal, and Erez Petrank. 2017. Layout Lock: A Scalable Locking Paradigm for Concurrent Data Layout Modifications. In *Proc. 22nd ACM SIGPLAN Symp. Princ. Pract. Parallel Program. - PPOPP '17*, Vol. 52. ACM Press, New York, New York, USA, 17–29. <https://doi.org/10.1145/3018743.3018753>
- Dave Dice, Timothy L Harris, Alex Kogan, Yossi Lev, and Mark Moir. 2014. Pitfalls of lazy subscription. *6th Work. Theory Trans. Mem. - WTTM'14* (2014).
- Dave Dice, Maurice Herlihy, and Alex Kogan. 2016. Fast non-intrusive memory reclamation for highly-concurrent data structures. In *Proc. 2016 ACM SIGPLAN Int. Symp. Mem. Manag. - ISMM 2016*. ACM Press, New York, New York, USA, 36–45. <https://doi.org/10.1145/2926697.2926699>
- Panagiota Fatourou and Nikolaos D. Kallimanis. 2011. A highly-efficient wait-free universal construction. In *Proc. 23rd ACM Symp. Parallelism algorithms Archit. - SPAA '11*. ACM Press, New York, New York, USA, 325. <https://doi.org/10.1145/1989493.1989549>
- Timothy L Harris. 2001. A pragmatic implementation of non-blocking linked-lists. In *Distrib. Comput.* Springer, 300–314. https://doi.org/10.1007/3-540-45414-4_21
- Thomas E Hart, Paul E McKenney, Angela Demke Brown, and Jonathan Walpole. 2007. Performance of memory reclamation for lockless synchronization. *JPDC* (2007), 1270–1285.
- Maurice Herlihy. 1991. Wait-free synchronization. *ACM Trans. Program. Lang. Syst. - TOPLAS* 13, 1 (1991), 124–149.
- M.P. Herlihy and J.E.B. Moss. 1992. Lock-free garbage collection for multiprocessors. *IEEE Trans. Parallel Distrib. Syst.* 3, 3 (1992). <https://doi.org/10.1109/71.139204>
- Maurice Herlihy and Nir Shavit. 2012. *The Art of Multiprocessor Programming*. Elsevier.
- Richard L. Hudson and J. Eliot B. Moss. 2001. Sapphire: Copying GC Without Stopping The World. In *Joint ACM-ISCOPE Conference on Java Grande*. 48–57.
- Alex Kogan and Erez Petrank. 2011. Wait-free queues with multiple enqueueers and dequeuers. In *Proc. 16th ACM Symp. Princ. Pract. parallel Program. - PPOPP '11*, Vol. 46. ACM Press, New York, New York, USA, 223. <https://doi.org/10.1145/1941553.1941585>
- Alex Kogan and Erez Petrank. 2012. A methodology for creating fast wait-free data structures. In *Proc. 17th ACM SIGPLAN Symp. Princ. Pract. Parallel Program. - PPOPP '12*, Vol. 47. ACM Press, New York, New York, USA, 141. <https://doi.org/10.1145/2145816.2145835>
- Daug Lea and JSR-166. 2013. <http://greppcode.com/file/repository.greppcode.com/java/root/jdk/openjdk/8u40-b25/java/util/concurrent/locks/StampedLock.java>.
- Maged M Michael. 2002a. *Safe Memory Reclamation for Dynamic Lock-Free Objects Using Atomic Reads and Writes*. Research Report RC22317. {IBM} Corp., Thomas J Watson Research Center, Yorktown Heights, NY. 21 pages. <https://doi.org/10.1145/571825.571829>
- Maged M. Michael. 2004. Hazard pointers: Safe memory reclamation for lock-free objects. *IEEE Trans. Parallel Distrib. Syst.* 15, 6 (2004), 491–504. <https://doi.org/10.1109/TPDS.2004.8>
- Maged M Michael and Michael L Scott. 1995. *Correction of a Memory Management Method for Lock-Free Data Structures*. Technical Report 8930. University of Rochester. 1–7 pages. <https://doi.org/1802/503>
- Mm Maged M Michael. 2002b. High performance dynamic lock-free hash tables and list-based sets. In *Proc. fourteenth Annu. ACM Symp. ACM*, 73–82. <https://doi.org/10.1145/564870.564881>
- Erez Petrank. 2012. Can parallel data structures rely on automatic memory managers?. In *MSPC*. ACM, 1.
- Filip Pizlo, Daniel Frampton, Erez Petrank, and Bjarne Steensgard. 2007. Stopless: A Real-Time Garbage Collector for Multiprocessors. In *ISMM*. 159–172.

- Filip Pizlo, Erez Petrank, and Bjarne Steensgaard. 2008. A Study of Concurrent Real-Time Garbage Collectors. In *PLDI*. 33–44.
- Filip Pizlo, Lukasz Ziarek, Petr Maj, Antony L. Hosking, Ethan Blanton, and Jan Vitek. 2010. Schism: Fragmentation-Tolerant Real-Time Garbage collection. In *PLDI*. 146–159.
- Pedro Ramalhete and Andreia Correia. 2017. Brief announcement: Hazard eras-non-blocking memory reclamation. In *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures*. ACM, 367–369.
- Shahar Timnat, Anastasia Braginsky, Alex Kogan, and Erez Petrank. 2012. Wait-free linked-lists. In *Proc. 17th ACM SIGPLAN Symp. Princ. Pract. Parallel Program. - PPoPP '12*, Vol. 47. ACM Press, New York, New York, USA, 309. <https://doi.org/10.1145/2145816.2145869>
- Shahar Timnat and Erez Petrank. 2014. A practical wait-free simulation for lock-free data structures. In *Proc. 19th ACM SIGPLAN Symp. Princ. Pract. parallel Program. - PPoPP '14*. ACM, 357–368.
- John D Valois. 1995. Lock-free linked lists using compare-and-swap. In *Proc. fourteenth Annu. ACM Symp. Princ. Distrib. Comput. - Pod. ACM*, 214–222.
- Haosen Wen, Joseph Izraelevitz, Wentao Cai, H Alan Beadle, and Michael L Scott. 2018. Interval-based memory reclamation. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. ACM, 1–13.
- Chaoran Yang and John Mellor-Crummey. 2016. A wait-free queue as fast as fetch-and-add. In *Proc. 21st ACM SIGPLAN Symp. Princ. Pract. Parallel Program. - PPoPP '16*, Vol. 51. ACM Press, New York, New York, USA, 1–13. <https://doi.org/10.1145/2851141.2851168>