



# The ERA Theorem for Safe Memory Reclamation

Gali Sheffi\*

Pliops  
Ramat Gan, Israel  
sheffiga@gmail.com

Erez Petrank

Department of Computer Science, Technion  
Haifa, Israel  
erez@cs.technion.ac.il

## ABSTRACT

Safe memory reclamation (SMR) schemes for concurrent data structures offer trade-offs between three desirable properties: ease of integration, robustness, and applicability. In this paper we rigorously define SMR and these three properties, and we present the ERA theorem, asserting that any SMR scheme can only provide at most two of the three properties.

## CCS CONCEPTS

• **Software and its engineering** → *Garbage collection*; **Scheduling**; **Deadlocks**; **Multithreading**.

## KEYWORDS

concurrency, safe memory reclamation, lock-freedom, robustness

### ACM Reference Format:

Gali Sheffi and Erez Petrank. 2023. The ERA Theorem for Safe Memory Reclamation. In *ACM Symposium on Principles of Distributed Computing (PODC '23)*, June 19–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3583668.3594564>

## 1 INTRODUCTION

Managing memory for concurrent data structures is known to be non-trivial. The main problem is that a node  $n$  that is detached from the data structure and is headed for reclamation, may still be accessed by concurrent threads that gained access to  $n$  prior to its detachment. Once a detached node is reclaimed, the executing threads are in danger of accessing freed memory, potentially causing a system crash, a segmentation fault, or correctness failure [36, 37]. This problem can be prevented using a rigid access discipline with locking, but such locking is not typically used because it is often detrimental to performance, and because locking foils the progress guarantee of non-blocking data structures.

To deal with this problem, SMR schemes were presented. The task of an SMR scheme is to either prevent hazardous accesses, or delay the reclamation of nodes that are still accessible by concurrent threads. With an SMR scheme, nodes are first manually *retired*, announcing that they are candidates for reclamation (and re-allocation) and the SMR scheme is responsible for determining

when a retired node can be safely reclaimed and reused. Retired nodes are typically held in per-thread retire lists [24, 37, 41, 42, 46, 52] until they are eligible for reclamation.

Each SMR comes with some benefits over existing schemes, but also with some disadvantages for the concurrent system. In this paper we formulate three desirable properties that were obtained (never simultaneously) in prior work: Ease of integration, Robustness, and wide Applicability (ERA, in short). We define these properties formally, and present a theorem, asserting that it is not possible for any SMR scheme to deliver all three qualities.

Concurrent data-structure implementations usually adhere to some correctness criteria, typically, linearizability [30], and provide some progress guarantee, typically lock-freedom or wait-freedom [26, 28]. Interestingly, there is no analogue widely accepted correctness criteria for SMR implementations. Various works have phrased ad-hoc correctness conditions [5, 17, 32–34, 37]. For example, [37, 41, 52] introduce the terms *protection*, or *hazardous access*, which are specific to the SMR scheme they present, but are not adequate for other schemes. Some previous work prove ad-hoc properties that relate to correctness [42, 46, 52], but there is no correctness condition that is applicable to all SMR techniques in the literature. Clearly, a reclamation scheme should preserve the original implementation's correctness (e.g., linearizability) and it is desirable to also keep its progress guarantee. The notion of safety is often used, but with no rigorous acceptable definition available. For example, safety may be interpreted as preventing access to reclaimed memory, but then, optimistic methods [13, 14, 42], that allow access to reclaimed data, may erroneously be considered as unsafe.

In addition to the lack of an adequate definition for safe memory reclamation, there are also some properties of SMR schemes that have not been previously rigorously defined. Previous works [32, 40, 46] listed various desirable SMR properties, which are often used to evaluate the overall quality of an SMR scheme. The most natural properties that we care about include performance (and scalability), progress guarantee, applicability to a wide set of concurrent data-structure implementations, the ease of integration with a given concurrent data structure, and the memory footprint, a.k.a. robustness. While some of these properties seem intuitively clear, most of them were not formally defined. Moreover, robustness was only recently introduced, only in the scope of memory reclamation [5, 17], and is sometimes used to mean different things. Let us briefly discuss the three notions that are at the focus of this paper.

Defining wide *applicability* is not straightforward. It is not always easy to tell if a given SMR scheme is applicable to a given data structure. For example, it is not trivial to see that the Hazard Pointers (HP) method [37] and some of its extensions [39, 41, 52] are not applicable to Harris's linked-list [24], and it is not clear how to test HP's applicability to a new data structure. While the

\*This work was done while the author was at the Technion - Israel.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

PODC '23, June 19–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0121-4/23/06...\$15.00

<https://doi.org/10.1145/3583668.3594564>

applicability notion has been considered previously [13, 32, 42, 46], it has not been formally defined.

*Robustness* has two different meanings in the literature. According to Dice et al. [17], an SMR scheme is considered as robust if there is a bound on the number of retired objects that cannot be reclaimed. According to Balmau et al. [5], the number of reclamation-related computational steps should be bounded. While the first definition has subsequently been more widely accepted [42, 46, 52], both were stated informally. First, the bound is not defined. It is unclear whether it should be a constant, or whether it may depend on the size of the data structure, or even the execution length, and in what way. If the data structure requires space, exponential in the number of insert operations, is it robust? If a scheme exhausts the heap even with small data structures, is it robust? Furthermore, what does "cannot be reclaimed" mean? Various SMR schemes are driven by different reclamation triggers, and a reclamation of a certain set of objects may be postponed, regardless of safety. Moreover, the terminology in existing definitions is not obvious. What is a reclamation-related step? Some SMR schemes change the original implementation's layout fundamentally, making the difference between the original implementation steps and the newly inserted ones ambiguous. We feel that a formal definition of robustness, which rigorously clarifies the concept of memory overhead for SMR schemes, is missing.

Finally, we look at how difficult it is to integrate an SMR scheme into a given data structure, i.e., at *ease of integration*. An SMR scheme typically provides a set of API operations that should be inserted into the given code (e.g., *alloc()*, *retire()*, *beginOperation()* [20, 24, 52]). Sometimes, the integration also involves changing the original program. E.g., the Automatic Optimistic Access (AOA) scheme [13] requires the data structure to be in a normalized form [49] for the integration. Free Access (FA) [12] and Neutralization-Based Reclamation (NBR) [46] require that the code is divided into separate read and write phases, and Version-Based Reclamation (VBR) [42] provides a designated mechanism for adding code checkpoints. Note that wide applicability and easy integration are independent properties. A scheme may be easily integrated to a given data-structure code, but, still, not be correctly applicable to it, or it may be applicable but hard to integrate.

In this paper, we formally define safe memory reclamation, along with the three desirable properties: robustness, wide applicability and ease of integration. In addition, we present and prove the ERA theorem, asserting that the three properties (Ease of integration, Robustness, and wide Applicability) cannot co-exist. I.e., any SMR scheme can have at most two out of the three desirable properties. This paper is organized as follows: Related work is surveyed in Section 2. In Section 3 we specify the shared-memory model. We formally define safe memory reclamation in Section 4, and the three desirable properties in Section 5. We present the ERA Theorem in Section 6, and conclude in Section 7.

## 2 RELATED WORK

Detailed surveys of SMR schemes appear in the literature [10, 42, 46, 52]. We focus on previous efforts to address and define suitable correctness conditions for SMR schemes and desirable reclamation properties. To the best of our knowledge, while there exist

formal methods for verifying safety in manually reclaimed environments [23, 33, 34], and various reclamation schemes have been shown to possess highly desirable properties (e.g., robustness, easy integration and wide applicability), these notions have never been formally defined.

Michael [37] and Herlihy et al. [27] were the first to address the need to bound the amount of memory occupied by removed elements. In both suggested schemes, each thread has a pool of global pointers (called hazard pointers in [37] or guards in [27]), used to postpone the reclamation of certain nodes that might still be in use. Both schemes were designed to provide an upper bound guarantee on the total number of deleted nodes that are not yet eligible for reuse at any given moment. Braginsky et al. [8] suggested a flexible trade-off between space and runtime overhead incurred by these two reclamation schemes. The term *Robustness* was first introduced by Dice et al. [17] and Balmau et al. [5]. While the first definition puts a limit on the number of deleted objects that could not be reclaimed, the latter bounded the number of steps during any action related to memory reclamation. Subsequently, the first definition (bounding the space overhead) was widely adopted [32, 39–42, 46, 52]. However, the bound on the space overhead was never specified, and so reclamation schemes with liberal bounds (that can potentially exhaust the available memory) [39, 41, 52] are considered robust. Reference counting-based schemes [16, 21, 27] are usually not robust, mainly due to the existence of cyclic structures of retired objects (these cycles can be broken [15, 48], at the cost of higher performance overheads).

Although some reclamation schemes are designed for a limited set of data-structures [8], most methods are designed for a general, wide set of data-structure implementations. A common issue which arises in this context is that many allegedly general schemes [37, 39, 41, 47, 52] require restricted access to deleted nodes. I.e., traversing a deleted node is forbidden. This requirement is problematic, as many state-of-the-art lock-free data-structures allow such traversals [9, 18, 24, 25, 38] to obtain fast searches. Indeed, this issue was discussed in various works [10, 13, 20, 32]. An extensive study by Singh et al. [46] lists many popular lock-free data-structures for which such schemes are not applicable. Michael [35] suggested a modification of Harris's lock-free linked-list [24], which makes it suitable for these reclamation schemes. However, this modification reduces the performance of the linked-list (for more details, see [13]). Furthermore, it is not a universal construction and it cannot be used to modify other data-structures. Gidenstam et al. [21] suggested a more general solution, but their construction adds even higher performance overheads, and involves a complicated manual integration into the given code. Wide applicability was also referred to as *generality* in [39].

Applicability does not guarantee an easy integration. First, the programmer must issue *retire()* statements at a location where the node is already detached from the data structure. Also, reclamation schemes often require additional integration. The simplest integration is provided by the seminal EBR scheme [10, 20, 24], which solely requires inserting calls to external methods in the beginning and end of each data-structure operation. Other schemes [32, 37, 41, 52] provide a slightly more complicated interface, which is still relatively easy for integration. Such interfaces may include an explicit node protection method (preventing the reclamation of retired

nodes while they are potentially still in use), and read and write barriers (i.e., code to be executed with any read or write of a data structure field). Such methods can be automatically integrated into an existing implementation, and do not require a significant familiarity with the original code. However, some reclamation schemes (e.g., [10, 13, 14, 42, 46]) require a more complicated integration procedure. Such schemes may cause the original algorithm to fail in accessing data structure fields and retry. This implies control flow changes to determine where to branch when such a failure occurs. Such integration is non-trivial, and it requires a deep understanding of the underlying data structure. Automatic integration is no longer possible. However, as the ERA Theorem asserts (see Section 6), such harder integration efforts are required to obtain general applicability and a low space overhead.

In addition to performance, robustness, wide applicability, and easy integration, other reclamation properties have been considered in previous work. Singh et al. [46], considered *consistency*, which requires the overall performance to not be affected by workload changes or under a system over-subscription. Nikolaev and Ravindran [40] claimed that a reclamation scheme should be *transparent*. Namely, threads can be created and deleted dynamically throughout the execution, and without affecting the scheme's safety. Cohen and Petrank introduced an *optimistic* reclamation method [13, 14], which allows threads to read reclaimed memory, while taking care to preserve correctness. Such schemes provide strong robustness and high throughput. Sheffi et al. [42] extended optimistic access to writes, providing a fully optimistic solution.

In addition to ensuring desirable properties, avoiding common unwanted side effects is also important. Many reclamation schemes introduce significant memory overheads, as they add extra fields to the data-structure nodes' layout. E.g., epoch-related data [41, 42, 52], or reference counters [16, 21, 27]. Optimistic methods rely on type-preservation [13, 14, 42], which means that nodes should be re-allocated to the same type. This is adequate particularly when there is a small number of major data structures in the code. Some reclamation algorithms require special compiler or hardware support (and therefore, are not considered as *self-contained* [32]). E.g., DEBRA+ [10] and NBR [46] rely on lock-free OS signals to preserve lock-freedom, VBR [42, 43] and Hyaline [39] rely on a hardware-provided wide CAS instruction, StackTrack [2] and ThreadScan [3] rely on transactional memory, Dice et al. [17] and PEBR [32] rely on the existence of process-wide memory fences, and QSense [5] requires control over the OS scheduler.

### 3 PRELIMINARIES

We follow previous work and use the basic asynchronous shared memory model from [26], and related definitions from [30]. We consider a fixed set of  $N$  executing threads, communicating by applying operations on *shared objects*. An object is an instance of an *abstract data type*, which specifies a set of possible values, and a set of operations that provide the only means to access it. For example, we define the *set* data type as a set of integer keys (denoted as *set keys*), initially empty. Its associated operations are *insert(key)*, *delete(key)* and *contains(key)*, where *key* is an integer. The *insert(key)* operation inserts *key* into the set and returns *true* if the set does not already contain *key*, and returns *false* otherwise.

The *delete(key)* operation removes *key* from the set and returns *true* if the set indeed contains *key*, and returns *false* otherwise. The *contains(key)* operation returns *true* if the set contains *key*, and returns *false* otherwise.

*Implementations and Data-Structures.* An *implementation* of an object (or several objects) provides a data-representation by applying primitive memory access operations (e.g., reads, writes, atomic read-modify-write instructions [26]) on a set of base objects (i.e., shared memory locations). Specifically, *set objects* are represented by *shared data-structures*. Each set key is represented by a *node*, and each data-structure has a fixed set of entry points (e.g., a linked-list head [24] or a tree root [38]), which are node pointers. Nodes may contain node pointer fields. We say that a node  $m$  is a successor of a node  $n$  (or that  $n$  is a predecessor of  $m$ ) if at least one of  $n$ 's node pointer fields points to  $m$ . Accordingly, we say that a node  $m$  is reachable from a node  $n$  if there exist nodes  $n_0, n_1, \dots, n_k$  such that (1)  $n = n_0$ , (2)  $m = n_k$ , and (3) for every  $0 \leq i < k$ ,  $n_i$  is a predecessor of  $n_{i+1}$ . We say that a certain node is *reachable* if it is reachable from an entry point.

*Executions.* A *step* is either a shared-memory access, a local variable access, an operation invocation, or the return from an operation. In all cases, the step includes the executing thread id, the accessed object (when exists), and the respective access input and output values. Steps are considered to be atomic. A *configuration* specifies the value of each shared memory address and the state of each thread (including the content of its local variables and program counter). The initial configuration  $C_0$  is the configuration in which all memory addresses have their initial values and all threads are in their initial states. In particular, all data-structures are initialized, and represent empty sets. An *execution*  $E = C_0 \cdot s_1 \cdot C_1 \cdot \dots$  is an alternating sequence of configurations and steps, starting from the initial configuration. Specifically, an *execution of an implementation* is an execution where, starting from the initial configuration, each step is issued according to the given implementation, each memory read matches the preceding configuration, and each memory write is reflected in the following configuration. Given a sub-sequence  $E' = C_k \cdot s_{k+1} \cdot \dots \cdot s_m \cdot C_m$ , we say that  $E'$  is a *solo-run* if the steps  $s_{k+1}, s_{k+2}, \dots, s_m$  are executed by the same thread.

*Histories.* An execution  $E$  induces a *history*  $H$ , which is its sub-sequence of operation invocation and response steps. Given an implementation, its set of *derived histories* is the set of all histories that model executions of that implementation. Given a history  $H$  and a thread  $T$ , we denote with  $H|T$  the sub-history of  $H$ , consisting of exactly all the steps executed by  $T$  in  $H$ . Similarly, given a history  $H$  and a shared object  $O$  (may be a memory word or a data-structure as described above), we denote with  $H|O$  the sub-history of  $H$ , consisting of exactly all the steps executed on  $O$  in  $H$ . Accordingly, given a history  $H$ , a thread  $T$  and a shared object  $O$ , we denote with  $H|(T, O)$  the sub-history of  $H$ , consisting of exactly all the steps executed by  $T$  on  $O$  in  $H$ . Two histories  $H, H'$  are *equivalent* if for every thread  $T$ , it holds that  $H|T = H'|T$ .

Given a history  $H$  and an object  $O$ , we say that  $H|O$  is sequential if it begins with an invocation step, and each invocation step (except for possibly the last one) is immediately followed by its matching response. We say that a history  $H$  is a *sequential history* if for

every object  $O$ ,  $H|O$  is sequential. An object is associated with a *sequential specification*, which is a prefix-closed set of all of its possible sequential histories.

Given a history  $H$  that contains an operation invocation, we say that this operation is *complete* in  $H$  if  $H$  also contains its matching response. Otherwise, we say that this operation is *pending* in  $H$ . A history is *complete* if all of its contained operations are complete.

**Well-Formed Histories.** The standard definition of well-formed histories [30] assumes that, given a history  $H$  and a thread  $T$ ,  $H|T$  is a sequence of operation invocations and their immediate matching responses. However, describing the integration of a safe memory reclamation scheme into a given data-structure implementation requires nesting operations. I.e., the reclamation scheme's operations (e.g., *retire()*, *alloc()*) are called in the scope of the data-structure operations. Therefore, we cannot use the standard definition of well-formed histories from [30]. Instead, we follow the extended definition from [4]. Given a history  $H$  and an object  $O$ , we say that  $H|O$  is *well-formed* if for every thread  $T$ ,  $H|T, O$  starts with an invocation step, and is an alternating sequence of invocation steps and their immediate matching response. We say that a history  $H$  is *well-formed* if (1) for every object  $O$ ,  $H|O$  is well-formed, and (2) for every thread  $T$ , two of its invocation steps  $s_{inv_1}, s_{inv_2}$  and their respective matching response steps  $s_{res_1}, s_{res_2}$ , if  $s_{inv_1}$  precedes  $s_{inv_2}$  and  $s_{inv_2}$  precedes  $s_{res_1}$  in  $H$ , then  $s_{res_2}$  precedes  $s_{res_1}$  in  $H$ . A *well-formed implementation* is an implementation for which all derived histories are well-formed.

**Linearizability.** A complete history  $H$  is linearizable if it is well-formed, and for every object  $O$ , its sequential specification contains a sequential history  $S$  such that (1)  $H|O$  and  $S$  are equivalent, and (2) if a response step precedes an invocation step in  $H|O$ , then it also precedes it in  $S$ . A history  $H$  is linearizable if it can be completed (by adding matching response steps to a subset of pending operations in  $H$ , and removing the rest of  $H$ 's pending operations) to a linearizable complete history. A *linearizable implementation* is an implementation for which all derived histories are linearizable.

**Lock-Freedom.** We follow Herlihy and Shavit's definition of lock-freedom [28]. Given an execution  $E = C_0 \cdot s_1 \cdot \dots$  and an executing thread  $T$ , we say that  $T$  is *effective* in a configuration  $C_i$  if  $T$  performs the step  $s_j$  for some  $j > i$  (informally,  $T$  is not starved by the scheduler). Now, let  $s_k$  be an operation invocation by a thread  $T$  during an execution  $E$ . We say that  $s_k$  is *effective* if either  $s_k$  has a matching response step in  $E$ , or  $T$  is effective in  $C_m$  for every  $m \geq k$ .

A history  $H$  provides *minimal progress* if in every suffix of  $H$ , some pending effective invocation has a matching response.  $H$  provides *maximal progress* if in every suffix of  $H$ , every pending effective invocation has a matching response. An implementation is *lock-free* if every derived history provides minimal progress, and some derived history provides maximal progress.

## 4 DEFINING SAFE MEMORY RECLAMATION

In this section we present a formal definition of safe memory reclamation. For ease of presentation, we focus on data-structures that implement set objects. This allows defining a life cycle of a node in the data structure. Extensions to other object types are not difficult.

### 4.1 Nodes' Life-Cycles

Following Meyer and Wolff [33, 34] we assume that (for a set implementation) each node goes through stages in a life-cycle, and can be in one of four possible states: unallocated, local, shared, or retired. Initially, a node is *unallocated*. I.e., its memory is not available for use by the executing threads. After being allocated by a certain thread, the node becomes *local*. While being local, no thread but the allocating thread has access to this node. In particular, the node cannot be reachable (from an entry point of the data structure), and it cannot represent a set item at this stage. Next, the node may or may not become *shared* (e.g., by making it reachable). While being shared, the node may become alternately reachable and unreachable, and may also represent a set item. When a node is either local or shared, we also say that it is *active*. At some point, an executing thread may retire the node, announcing that this node is about to become garbage. Once a thread retires the node, it becomes *retired* (and cannot be retired again). Note that some nodes never become shared, and therefore may become retired after being local. In the lifetime of a node, we assume that a node always becomes unreachable before it becomes retired (generally, nodes can only be reachable while they are shared<sup>1</sup>). Finally, a retired node may be reclaimed, meaning that its memory may now be used for re-allocation and its state becomes *unallocated* again. We consider nodes as logical entities. I.e., after a node returns to being unallocated, a new allocation from the same address is considered as an allocation of a different node (though both nodes are allocated on the same address).

### 4.2 Safe and Unsafe Memory Accesses

In this section we define a safe memory access. We think of the memory as segregated into two separate spaces – the *system space* and the *program space*. The program space is the area in which the executing threads keep their local and shared nodes and variables. In particular, new allocations are always within the program space, and all nodes reside in the program space, until they are reclaimed. At the time of reclamation, the memory reclamation scheme decides whether to keep the reclaimed nodes for potential subsequent re-allocation in the program space, or to return the node space to the system, in which case, the node moves to system space. If the program attempts to access a node in system space, the result is undefined, and may include a segmentation fault.

From now on, we refer to a given set implementation (with no memory reclamation) as the *plain* implementation. Note that although plain implementations do not include memory reclamation, they do follow the life cycle defined in Section 4.1. Specifically, they include adequate *retire()* instructions. Upon integrating a given memory reclamation scheme into a plain implementation, we refer to the derived implementation as the *integrated* implementation. We further use the *plain* and *integrated* terms to describe the respective executions and histories.

**Dereferencing** a pointer  $p$  means reading or updating a value in a node whose address is stored in the pointer  $p$ .<sup>2</sup> After memory is

<sup>1</sup>This assumption is necessary for most safe memory reclamation schemes. However, there exist scenarios in which it is not needed [51].

<sup>2</sup>Note that the pointer content is not necessarily equal to the stored address (e.g., marked pointers [24] contain addressees, possibly along with a marked bit).

reclaimed, dereferencing a pointer to it might cause a segmentation fault [36]. Given an execution  $E = C_0 \cdot s_1 \cdot \dots$ , a pointer variable  $p$ , and any configuration  $C_m$  in the execution ( $m \geq 1$ ), let  $s_i$  be the last update of  $p$  in the sub-execution  $C_0 \cdot s_1 \cdot \dots \cdot C_m$ . Namely, for  $i \leq m$ ,  $p$  is updated in  $s_i$ , and for every  $i < j \leq m$ ,  $s_j$  does not update  $p$ . The update of  $p$  in  $s_i$  may be an allocation of a new node to  $p$  or a pointer assignment to  $p$  from another pointer  $q$ .<sup>3</sup>

Let  $n$  be the node that is referenced by  $p$  in  $C_m$ . We separate into two cases according to whether  $s_i$  is an allocation or an assignment. If the last update of  $p$  is an allocation in  $s_i$ , and the node is not in an *unallocated* state in any of the intermediate configurations  $C_j$  (for  $i \leq j \leq m$ ), then we say that  $p$  is *valid* in  $C_m$ . Otherwise, we say that  $p$  is *invalid* in  $C_m$ . For example, by definition,  $p$  is always valid in  $C_i$ . The second case is that  $s_i$  is a pointer assignment, and let  $q$  be the pointer whose content is assigned into  $p$  in  $s_i$ . Similarly, we say that  $p$  is *valid* in  $C_m$  if  $q$  is valid in  $C_i$ , and for all intermediate configurations  $C_j$  for  $i \leq j \leq m$ ,  $n$  is not in an *unallocated* state in  $C_j$ . If  $p$  is not valid in  $C_m$ , we say that  $p$  is *invalid* in  $C_m$ . We can now formally define a safe memory accesses.

**Definition 4.1.** A memory access is unsafe if it dereferences an invalid pointer. It is safe otherwise.

### 4.3 Defining SMR in the Presence of Unsafe Memory Accesses

When designing a memory reclamation scheme, one must either provide an adequate solution for coping with unsafe memory accesses, or make sure that all memory accesses are safe. Some SMR schemes allow unsafe memory accesses, while taking care to preserve correctness. E.g., AOA [13] and VBR [42] allow reading via invalid pointers, as it is ensured that stale values are always ignored. VBR further allows trying to update the shared memory via invalid pointers, as it is guaranteed that the update fails. Definition 4.2 below encapsulates the conditions for a memory reclamation scheme to be considered as a safe one. Loosely speaking, an SMR should either only use safe memory accesses, or be very careful in how it executes unsafe memory accesses. In particular, it should not access system space (that might end up in a segmentation fault), it should not modify the data on the node (which might have been reclaimed), and it should not use a value read during an unsafe memory access.

**Definition 4.2.** A memory reclamation scheme is a safe memory reclamation (SMR) scheme with respect to a given plain implementation, if for each respective integrated execution  $E$ , all memory accesses in all steps are safe, or if all unsafe memory accesses satisfy the following conditions. Let  $s_i$  be a step with an unsafe memory access to a memory node  $n$  via a pointer  $p$ , then the following three conditions must hold:

- (1) In configuration  $C_{i-1}$ ,  $n$ 's occupied memory belongs to the program space.
- (2)  $s_i$  does not update  $n$ 's content. Namely,  $n$ 's content in  $C_{i-1}$  equals  $n$ 's content in  $C_i$ .

<sup>3</sup>In the following clean formalization, we assume no address arithmetic, and that a pointer points to the head of the object. One can easily extend the definitions below to programs that use pointer arithmetic, as long as there is a clear mapping from pointers to nodes.

- (3) If data from the dereferenced node  $n$  is read into a variable or field  $v$  (local or shared), then the value in  $v$  is never used.

The term "used" in Condition 3 refers to the standard program analysis terminology. In particular, if the modified variable (or field) is read in a step  $s_j$  (for some  $j > i$ ), then there exists  $i < k < j$  such that  $s_k$  overwrites the content of this variable (or field) before it is read.

Note that we only define safety with respect to a given plain implementation, as most schemes are not necessarily safe when integrated with all existing plain implementations. E.g., the HP [37] scheme is safe with respect to Michael's linked-list [35], but is not safe with respect to Harris's linked-list [24]. This implies that HP is not applicable to Harris's linked-list plain implementation. We further discuss applicability (and HP's applicability in particular) in Section 5.3 and in the full version of this paper [45].

## 5 DESIRABLE SMR PROPERTIES

In this section we present formal definitions for three of the most desirable reclamation scheme properties. Robustness is defined in Section 5.1, easy integration is defined in Section 5.2, and wide applicability is defined in Section 5.3. One property may affect another in a design. For example, robustness may affect progress, and preserving progress guarantees may affect applicability. But in the definitions, we separate the notions and define each of them independently of the others.

### 5.1 Robustness (Memory Footprint)

Similarly to [42, 46, 52], we adopt the definition of robustness from Dice et al. [17], which relates to the space overhead or memory footprint of a reclamation scheme. According to this definition, a reclamation scheme is considered robust if a failed or delayed thread cannot totally prevent memory reclamation. This is formalized by a bound on the amount of retired nodes that exist at any point in the execution. As in Section 4.1, retired nodes are nodes that have already been retired, but are not in the state of *unallocated*. This includes nodes that cannot be reclaimed (due to some reclamation condition that the nodes do not satisfy), together with the retired nodes that have simply not yet been reclaimed, typically because some periodic process that reclaims objects has not yet processed them, typically, because they are still "waiting" in a fixed-size retire list [10, 20, 46].

To the best of our knowledge, previous work does not specify any general definition for the bound on the number of such objects. It is not clear whether this bound should be a pre-defined constant, may depend on the specific execution, or may depend on the data-structure size. In definitions 5.1-5.2 below we classify the different levels of robustness, according to the bound that an SMR scheme satisfies.

In the following definitions we consider a bound on the number of retired objects that depends on the size of the data structure. The size of the data structure is dynamic and is bounded by the number of active nodes, i.e., the nodes that have been allocated and have not yet been retired (see Section 4.1). This choice follows the tradition in the memory management community, where the space overhead of a program execution is compared to the minimum space that is always required to execute the program. For us, this minimum

space is the space that the program must use even if any node that is disconnected from the data structure is immediately reclaimed and immediately becomes available for re-allocation.

We first define the class of robust reclamation schemes. Robustness bounds the number of retired nodes at any time by a function that is asymptotically smaller than the maximum size of the data structure so far in the execution, multiplied by the number of threads. Formally, given an execution  $E$ , we denote the number of active nodes in  $C_i$  by  $active_E(i)$ . In addition, we set the function  $max\_active_E(i)$  to be  $\max\{active_E(0), \dots, active_E(i)\}$ .

**Definition 5.1. (Robustness)** We say that a reclamation scheme is *robust* if for every integrated execution  $E$ , there exists a function  $f_E : \mathbb{N} \rightarrow \mathbb{N}$ , such that (1)  $f_E = o(max\_active_E)$ , and (2) for every configuration  $C_i$ , the number of retired nodes in  $C_i$  is bounded by  $f_E(i) \cdot N$ .

VBR [42] is robust, with  $f_E(i)$  being a constant function, bounded by the local retire list size (which does not depend on the execution). This scheme presents the strongest robustness available by an SMR today. Its bound does not depend on the size of the data structure.

HP [37], AOA [13], and NBR [46] all use hazard pointers [37] for write protection<sup>4</sup>. For all three schemes,  $f_E$  depends on the pre-defined local retire list size (similarly to VBR) plus the number of hazard pointers. The number of hazard pointers is typically a small constant (e.g., 3 for linked-lists [24, 35]), but may also depend on the number of active nodes (e.g., for skip lists with a dynamic number of levels [1]). While the number of hazard pointers is not guaranteed to be asymptotically smaller than the number of active nodes, in all known data structures it is. It is an open question for the study of hazard pointers to bound their number (in all applicable data structures). If the number of hazard pointers is always asymptotically smaller than the number of active nodes, then all three schemes are robust.

The robustness property is very useful in practice. The asymptotic space overhead of robust schemes is smaller than the space consumed by active nodes in the data structure at any point in time. Some reclamation schemes do not provide robustness, but a bound still exists. Accordingly, we define a relaxed term of robustness, denoted *weak robustness*, in Definition 5.2 below. Note that robust schemes are also considered as weakly robust schemes, but not vice-versa.

**Definition 5.2. (Weak Robustness)** We say that a reclamation scheme is *weakly robust* if for every integrated execution  $E$ , there exists a function  $f_E : \mathbb{N} \rightarrow \mathbb{N}$ , such that (1)  $f_E$  is polynomial in  $max\_active_E$ , and (2) for every configuration  $C_i$ , the number of retired nodes in  $C_i$  is bounded by  $f_E(i) \cdot N$ .

A weakly robust scheme might incur a larger space overhead. Usually, this happens only at worst-case scenarios, but a large space overhead is theoretically possible. In some hybrids of the epoch-based [20, 24] and pointer-based [37] reclamation approaches, the execution is divided into epochs, and the number of retired nodes is bounded by the number of active nodes during a certain set of epochs. Given an integrated execution  $E$  and an epoch  $e$ , that starts

in a configuration  $C_i$ , the total number of active nodes during  $e$  is  $active_E(i)$  plus the number of allocations during  $e$ . In IBR [52], each thread might prevent the reclamation of nodes that were active during a small set of *reserved* epochs<sup>5</sup>. As IBR allows only a constant number of allocations per epoch, the number of retired nodes in a configuration  $C_i$  is linear in  $max\_active_E(i) \cdot N$  (which is not asymptotically smaller than  $max\_active_E$ ). Therefore, IBR is weakly robust. In HE [41], the number of allocations per epoch is also a pre-defined constant, but the number of reserved epochs may depend on the number of active nodes (similarly to the number of hazard pointers in HP). However, the number of retired nodes is still polynomial in  $max\_active_E(i) \cdot N$ , and HE is considered as weakly robust as well.

EBR [10, 20, 24] is not even weakly robust. Once a thread is halted, all subsequently allocated nodes can never be reclaimed. There are two different avenues by which previous work chose to look at such schemes. Some papers consider such schemes as not providing lock-freedom, because when the memory is exhausted, the program cannot allocate anymore and hence, cannot guarantee progress. Other papers think of this scheme as guaranteeing progress assuming that the memory is unbounded. We adopt the second school for this paper, but augmenting the definitions to fit the first school is not hard.

## 5.2 Easy Integration

We assume that the plain implementation already contains proper *retire()* invocations. Namely, we do not consider *retire()* calls installations as part of the reclamation scheme integration. The obvious easiest integrate-able SMR is the EBR scheme. EBR provides a *retire()* implementation, along with two API operations, *beginOp()* and *endOp()*, to be respectively inserted in the beginning and end of every data-structure operation. Namely, any plain implementation can be easily integrated with EBR, and the integration process does not require an understanding of the plain implementation. A definition of easily integrated scheme should obviously include EBR.

Easy integration includes schemes that provide designated *alloc()* and *retire()* implementations, along with code to replace reading and writing from shared memory. Some other schemes provide *beginOp()* and *endOp()* implementations, to be inserted at the beginning and end of each operation, respectively. Although the integration of such schemes is slightly more complicated than EBR's, they are still considered easily integrated, as their integration does not require any familiarity with the original code.

Some schemes cannot be easily integrated with many plain implementations. An important example of integration obstacle is the requirement of an SMR to insert roll-back instructions for handling unsafe memory accesses. Namely, when some validation test fails, one needs to return program control to a point in the code from which it is safe to re-execute. Roll-backs are something that we rule out for easy integration. Another undesirable property of an SMR is that it modifies fields of the data structure. An SMR may add fields to a data structure node to be used for its own activity, but expect the SMR to not modify fields that the plain implementation

<sup>4</sup>The HP scheme uses hazard pointers for read protection as well, whereas for NBR the number of required hazard pointers is the number of addresses that will be written in a write phase.

<sup>5</sup>There is a trade-off between IBR's easy integration and the bound on the number of reserved epochs. For more details, see Section 5.2.

uses. This *modularity*, *encapsulation of information*, and *separation of concerns* between the data structure and the SMR activity are standard principles in software engineering. We stress that marking a pointer to signify a deleted object as in Harris' linked-list is not a problem, because this is a modification by the data structure to support its own deletion activity, that the SMR is not involved in. In contrast, if the SMR modifies a data structure field, then the programmer that integrates the SMR into the data structure must have an intimate acquaintance with the fields of the data structure to make sure that the SMR does not foil the data structure operations correctness or performance. Such an intimate knowledge, if required during integration, makes the integration difficult.

Some schemes deal with difficult roll-backs by adhering to specific code shapes, which allow easier placement of roll-back mechanisms. AOA [13] requires that the plain implementation is first transformed into a normalized form [50]. NBR [46] requires that the code is first divided into separate read and write phases (to be further discussed in Section 5.3), which enable easier rollback call installations. VBR [42] relies on linearizability [30] when installing checkpoints and roll-back instructions. We define the easy integration property more rigorously in Definition 5.3 below. This definition excludes reclamation schemes that alter the plain implementation layout and disallows rolling back into wisely chosen code locations. Consequently, it classifies AOA, FA, NBR and VBR as reclamation schemes that cannot be easily integrated. We state the definition and follow up with more explanations.

*Definition 5.3.* A reclamation scheme is considered an easily integrated scheme if the following conditions hold:

- (1) The reclamation scheme is provided as an object.
- (2) The reclamation scheme's API operations may only be inserted in the following code locations: (1) upon the invocation or before the termination of any operation of the plain implementation, (2) as a replacement to *alloc()* and *retire()* calls, or (3) as a replacement to primitive memory access operations.
- (3) An API operation that replaces a primitive memory access operation, should be a linearizable implementation of that primitive.
- (4) The integrated implementation should be well-formed.
- (5) The reclamation scheme may add new fields to the node's layout and it may access these fields, but it cannot access any other node fields.

According to Condition 1, the reclamation scheme should be provided as an object that can be used with all implementations. Namely, as defined in Section 3, it should provide a uniform set of API operations, which are the only way to use its functionality (for more details, see Section 3). This condition also ensures that a reclamation scheme is not adjusted in order to fit specific plain implementations. Condition 2 ensures that the integration procedure is indeed relatively easy (as in EBR and IBR). Condition 3 treats shared memory addresses as objects. Namely, each memory address is an object that provides a set of operations (e.g., read, write, read-modify-write), usually implemented via atomic primitives. According to Condition 3, if a reclamation scheme operation replaces such a primitive, then it should implement it in a linearizable manner. By the locality property of linearizability [30], this

maintains some level of equivalency between the plain implementation and the integrated one. In particular, requiring manual changes to the control flow of the (plain) data structure operations (beyond replacing memory accesses with linearizable implementations of such accesses) implies that a reclamation scheme is not easy to integrate. Condition 4 builds on Condition 1, and treats the plain implementation and the reclamation scheme implementation as implementations of two separate objects. In particular, the meaning of *well-formed* (as in Section 3) in this sense is that the integration cannot move control from within a reclamation-related operation to a point in the plain implementation. Namely, rollbacks from a reclamation API code back into code of the plain implementation are not allowed. Finally, Condition 5 ensures that the reclamation scheme does not assume anything regarding the node's layout, and does not access any of its original fields. It may only access new fields that it adds to the node.

AOA, NBR and VBR do not satisfy Definition 5.3, as they do not provide a uniform set of API operations. In particular, inserting roll-back instructions foils Condition 4, as it moves control to an external point in the data-structure code before terminating the current reclamation-related operation execution. Note that the reclamation API should not include an explicit reclamation operation, as reclamation is expected to occur in the scope of the reclamation scheme code.

### 5.3 Wide Applicability

In this section we define the applicability of a reclamation scheme to a given plain implementation, and define the wide-applicability property accordingly. Our applicability definition includes a proper safety engagement, and a reference to the plain implementation's correctness and progress guarantees. We use linearizability [30] as our correctness condition, but the definition can be easily adapted to fit other correctness conditions.

*Definition 5.4.* We say that a reclamation scheme is applicable to a plain implementation if the following hold:

- (1) **Memory safety:** The reclamation scheme is safe with respect to the plain implementation according to Definition 4.2.
- (2) **Correctness:** The integrated implementation is linearizable<sup>6</sup>.
- (3) **Progress:** The integrated implementation provides the same progress guarantee as the plain implementation.

The progress guarantee of a scheme is determined according to [28], and set to the weakest guarantee of any of its operations. For example, Herlihy and Shavit's linked-list is lock-free although its *contains()* operation is wait-free [29].

Given the definition of applicability of a reclamation scheme to a plain implementation in Definition 5.4, we now define strong applicability and wide applicability of a given SMR. A good property of a reclamation scheme is applicability to as many data structures as possible.

The seminal EBR scheme [10, 20, 24] is *strongly applicable*, as defined in Definition 5.5 below (the full proof appears in the full

<sup>6</sup>Note that linearizability (as defined in Section 3) refers to the object implemented by the plain implementation, regardless of integration. Namely, the integrated implementation is linearizable with respect to this object even when the reclamation scheme cannot be treated as a separate object (see Section 5.2).



version of this paper [45]). It is the strongest scheme in terms of applicability. The only assumption it makes with respect to the plain implementation is that *retire()* instructions are properly installed (for more details, see Section 4.1).

**Definition 5.5. (Strong Applicability)** We say that a reclamation scheme is *strongly applicable* if it is applicable to every plain implementation.

While strong applicability is highly desirable, to the best of our knowledge, EBR is the only scheme that satisfies it. There are still different extents of applicability for the various reclamation schemes in the literature. We are interested in reclamation schemes that, while not applicable to any imaginable data structure, are still widely applicable to many known data structures. To accurately define *widely applicable* reclamation schemes, we adopt the definition from previous work [46], that defines a large class of well-known and widely-used concurrent data-structure implementations (e.g., [18, 24, 25, 31, 38]), all applicable to the NBR reclamation scheme. This class is described in [46] as containing all data-structure implementations that can be divided into separate interleaving read and write phases. We provide the formal definition of this class of data-structure implementations in the full version of this paper [45]. We denote such implementations as *access-aware data-structure implementations*, and define wide applicability accordingly:

**Definition 5.6. (Wide Applicability)** We say that a reclamation scheme is *widely applicable* if it is applicable to all access-aware data-structure implementations.

Not all SMRs are widely applicable. We show in the full version of this paper [45] that HP, IBR and HE are not widely applicable.

## 6 THE ERA THEOREM

In Section 5 we showed that there exists a reclamation scheme which is both widely applicable and easily integrated (EBR). Assuming that the number of hazard pointers [37] is asymptotically smaller than the data-structure size (see Section 5.1), there also exists a reclamation scheme which is both robust and widely applicable (NBR). We do not know of a scheme that is both robust and easily integrated. Building such a scheme is an open problem<sup>7</sup>. In this section we prove the main theorem of this paper.

**THEOREM 6.1.** *Any memory reclamation scheme can provide at most two of the following three guarantees: robustness, easy integration and wide applicability.*

In fact, we prove a stronger result, namely that even weak robustness (see Definition 5.2) cannot be achieved when easy integration and wide applicability are provided. This stronger result immediately implies Theorem 6.1.

Standard proofs of impossibility for concurrent and distributed computing typically employ indistinguishability arguments, showing that a thread that needs to arrive at two different decisions for two different executions, cannot distinguish the two executions based on the data that it can access (e.g., [4, 7, 19, 22, 26, 44]). Our

<sup>7</sup>One can build such a scheme that is trivial in the sense that it is not applicable to anything, or it is applicable to data structures that do not reuse objects, but this would not be very interesting.

proof techniques are different. We consider the three desirable properties and build an execution that cannot satisfy all three properties while providing correctness. Similar proofs are common in the literature of complexity theory and foundations of cryptology (e.g., impossibility proof for short black-box concurrent zero-knowledge or strong program obfuscation [6, 11]).

In the full version of this paper [45] we show that Harris' linked-list is access-aware, and therefore, a widely applicable reclamation scheme must be also applicable to Harris's linked-list implementation [24]. The main idea in the proof is to assume in a way of contradiction that an SMR does satisfy all three properties. In this case, it must be applicable to Harris' linked-list, and we then build a specific execution that is not safe for this concurrent linked-list implementation. Thus, no SMR can satisfy all three properties.


As described in Section 3, the list API provides the *insert()*, *delete()* and *contains()* operations, and the nodes comprise of two fields – an immutable key and a *next* pointer to the node's successor in the list. The list maintains two sentinel nodes, *head* and *tail*, with the respective  $-\infty$  and  $\infty$  keys, that are never removed from the list. Nodes are logically inserted into the list by physically linking them into the list and making them reachable, and are logically deleted from the list by marking their *next* pointer (for more details, see [24, 29]). Note that after a node is marked for deletion, it is not necessarily unlinked by the thread that had previously marked it, as it might be unlinked during a concurrent operation. However, the marked node is guaranteed to be unlinked and retired before the *delete()* operation returns.

All three API operations use the *search()* auxiliary method, which is in charge of (1) locating a given key in the list, and (2) unlinking logically deleted (i.e., marked) nodes from the list. This method traverses the list (by following *next* pointers) until it finds the first unmarked node with a key greater than or equal to the searched key. After locating such a node, the method might try to physically unlink a sequence of marked nodes from the list. The crucial point here is that marked nodes are not unlinked during the traversal. As opposed to Michael's implementation [35] (that was originally designated to fit HP [37]), when the *search()* method encounters a marked node, it just continues its traversal.

We prove Theorem 6.1 by constructing a specific execution. Let  $N \geq 2$  be some fixed constant, and assume  $N$  threads are executing Harris's linked list plain implementation, integrated with a widely applicable memory reclamation scheme. By Definition 5.6, the given reclamation scheme is applicable to this plain implementation. In particular, by Definition 5.4, the integrated implementation must provide the same progress guarantee as Harris's algorithm, namely lock-freedom. Now, assume by contradiction that the scheme is both weakly robust and easily integrated.

Initially (stage a in Figure 1), there are two reachable nodes in the list (besides the *head* and *tail* sentinels). Assume that  $T_1$  starts executing a *delete(3)* operation. It calls *search(3)*, and starts its traversal by reading *head's next* pointer, which is currently referencing node 1. At this stage, the scheduler moves control to  $T_2$ , which executes a *delete(1)* operation.  $T_2$  marks node 1 for deletion (stage b) and physically unlinks it from the list (stage c). Next,  $T_2$  executes an *insert(3)* operation (stage d) and a *delete(2)* operation (stages e-f). In a similar way,  $T_2$  continues calling an alternating sequence of *insert(n+1)* and *delete(n)* (starting from  $n = 3$ ).



Figure 1: Lower bound illustration ( $n = 2$ ).

For every  $n \geq 1$ , let  $i_n$  be the integer such that  $C_{i_n}$  is the configuration after  $T_2$  returns from the  $delete(n)$  execution. Note that for every  $n \geq 2$ , after  $T_2$  executes  $insert(n+1)$ , there are four active nodes in the system ( $head$ ,  $n$ ,  $n+1$  and  $tail$ ), as the rest of the nodes are already retired before their respective  $delete()$  operations return. Finally, after  $T_2$  executes  $delete(n)$ , there are three active nodes in the system ( $head$ ,  $n+1$  and  $tail$ ), and the nodes  $1, \dots, n$  are already retired. Therefore, for every  $n \geq 1$ ,  $max\_active_E(i_n) = 4$ . As the integrated reclamation scheme is weakly robust, there exists a function  $f_E$ , which is polynomial in  $max\_active_E$ , such that the number of retired nodes in  $C_{i_n}$  is bounded by  $f_E(i_n) \cdot N$ .

Let  $n > f_E(i_n) \cdot N$  ( $n$  must exist as the number of threads  $N$  and  $max\_active_E$  are constants). In  $C_{i_n}$ , for every  $1 \leq i \leq n$ , node  $i$  has already been retired, and as  $n > f_E(i_n) \cdot N$ , by Definition 5.2, at least one of the nodes  $1, 2, \dots, n$  must already be reclaimed. In  $C_{i_n}$ , if a node  $i$ ,  $1 \leq i \leq n$ , is already retired but not yet reclaimed, then it must be marked and pointing to node  $i+1$  via its  $next$  pointer. To see that the latter is true, recall that, as the scheme is easily integrated, according to Condition 5 from Definition 5.3, the reclamation scheme does not update  $next$  pointers.

Starting from  $C_{i_n}$ , let the scheduler apply a solo-run by  $T_1$ . I.e.,  $T_1$  is the only effective thread in  $C_{i_n}$  and on (for more details, see Section 3).  $T_1$  continues its traversal from node 1. As all nodes along its path are marked, the traversal's stopping condition does not hold for all nodes along its path (which have keys smaller than  $n+1$ , and that have not been reclaimed yet), and  $T_1$  should continue its traversal. By Condition 4 from Definition 5.3 (forcing well-formedness),  $T_1$  must return from this read operation (as implemented by the reclamation scheme) before it continues its execution. As  $T_1$  is the only effective thread, and as lock-freedom is guaranteed (see Section 3), every such read operation by  $T_1$  indeed terminates. In addition,

according to Condition 3 from Definition 5.3, as long as a node  $x$  (for any  $1 \leq x \leq n$ ) is not reclaimed, a read of its  $next$  pointer by  $T_1$  must return a marked reference to the memory (either previously or currently) occupied by node  $x+1$ . Therefore, as long as  $T_1$  does not encounter a reclaimed node, every read of a  $next$  pointer must terminate, returning a (marked) reference to the next node.

Recall that there exists an already reclaimed node  $m$  (for some  $1 \leq m \leq n$ ), such that for every  $i < m-1$ , node  $i$ 's next pointer is marked, and is pointing to node  $i+1$ . Eventually,  $T_1$  must dereference a pointer stored in the memory formerly occupied by  $m$ . It must assign the content of an invalid  $next$  pointer (as  $m$  has already been reclaimed) to a local pointer variable,  $p$ , performing an unsafe memory access by Definition 4.1. By Definition 5.4, the given reclamation scheme is safe with respect to this plain implementation, and therefore, by Definition 4.2,  $T_1$ 's local pointer  $p$  must be overridden before  $T_1$  dereferences it. By Condition 2 and 3 from Definition 5.3,  $T_1$  does not perform any updates before dereferencing  $p$ , as it contains a marked reference – a contradiction to the scheme's applicability to the given plain implementation. Therefore, a memory reclamation scheme cannot provide robustness, easy integration and wide applicability. ■

*Discussion.* We prove this result using a specific data structure (Harris's linked-list). This is enough to prove the impossibility in Theorem 6.1, but it actually provides a stronger result. Even if one tries to achieve robustness and ease of integration only for Harris's linked-list, then this attempt must fail. Therefore, other weaker interpretations of applicability that only require applicability to smaller sets of implementations must still respect the impossibility of Theorem 6.1, as long as the notion of applicability for an SMR requires it to be applicable to Harris's list (among other implementations).

An interesting open question is to characterize implementations that are similar to Harris' linked-list. Namely, that memory reclamation with (weak) robustness and ease of integration cannot be obtained for them. This may shed light on which data structures require special care.

Finally, we stress the practical importance of this theorem. In order to apply the HP scheme to Harris's linked-list, Maged [35] modified Harris's implementation to disallow the simultaneous removal of multiple consecutive nodes from the list. While this enabled applying HP to the list, Michael's implementation was slower than the original implementation, see for example the evaluation in [13]. While there may be another modifications of Harris's list that is both efficient and for which HP is applicable, such a modification is not known today. Thus, avoiding specific data-structure implementations may yield a noticeable performance reduction in practice.

## 7 CONCLUSION

In this paper we proposed some theoretical foundation for safe memory reclamation for concurrent data structures. We provided definitions for safe memory reclamation and for three fundamental desirable properties: robustness, easy integration and wide applicability. We then proved that no reclamation scheme can provide all three desirable properties. I.e., robust reclamation schemes (with limited space overhead) should either be designed for specific implementations, or come with a relatively complicated manual for proper integration. Open questions include the formalization of additional interesting properties of SMRs and formal safety proofs for existing SMR schemes.

## ACKNOWLEDGMENTS

This work was supported by the Israel Science Foundation Grant No. 1102/21. We thank Jeehoon Kang, Jaehwang Jung, Michael Scott, and Trevor Brown for helpful discussions.

## REFERENCES

- [1] Vitaly Aksenov, Dan Alistarh, Alexandra Drozdova, and Amirkeivan Mohtashami. 2020. The splay-list: A distribution-adaptive concurrent skip-list. *arXiv preprint arXiv:2008.01009* (2020).
- [2] Dan Alistarh, Patrick Eugster, Maurice Herlihy, Alexander Matveev, and Nir Shavit. 2014. Stacktrack: An automated transactional approach to concurrent memory reclamation. In *Proceedings of the Ninth European Conference on Computer Systems*. 1–14.
- [3] Dan Alistarh, William Leiserson, Alexander Matveev, and Nir Shavit. 2018. Threadscan: Automatic and scalable memory reclamation. *ACM Transactions on Parallel Computing (TOPC)* 4, 4 (2018), 1–18.
- [4] Hagit Attiya, Ohad Ben-Baruch, and Danny Hendler. 2018. Nesting-safe recoverable linearizability: Modular constructions for non-volatile memory. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*. 7–16.
- [5] Oana Balmau, Rachid Guerraoui, Maurice Herlihy, and Igor Zablotchi. 2016. Fast and robust memory reclamation for concurrent data structures. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*. 349–359.
- [6] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. 2001. On the (im) possibility of obfuscating programs. In *Annual international cryptography conference*. Springer, 1–18.
- [7] Ohad Ben-Baruch, Danny Hendler, and Matan Rusanovsky. 2020. Upper and lower bounds on the space complexity of detectable objects. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*. 11–20.
- [8] Anastasia Braginsky, Alex Kogan, and Erez Petrank. 2013. Drop the anchor: lightweight memory management for non-blocking data structures. In *Proceedings of the twenty-fifth annual ACM symposium on Parallelism in algorithms and architectures*. 33–42.
- [9] Trevor Brown, Faith Ellen, and Eric Ruppert. 2014. A general technique for non-blocking trees. In *Proceedings of the 19th ACM SIGPLAN symposium on Principles and practice of parallel programming*. 329–342.
- [10] Trevor Alexander Brown. 2015. Reclaiming memory for lock-free data structures: There has to be a better way. In *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*. 261–270.
- [11] Ran Canetti, Joe Kilian, Erez Petrank, and Alon Rosen. 2002. Black-box concurrent zero-knowledge requires (almost) logarithmically many rounds. *SIAM J. Comput.* 32, 1 (2002), 1–47.
- [12] Nachshon Cohen. 2018. Every data structure deserves lock-free memory reclamation. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 143:1–143:24. <https://doi.org/10.1145/3276513>
- [13] Nachshon Cohen and Erez Petrank. 2015. Automatic memory reclamation for lock-free data structures. *ACM SIGPLAN Notices* 50, 10 (2015), 260–279.
- [14] Nachshon Cohen and Erez Petrank. 2015. Efficient memory management for lock-free data structures with optimistic access. In *Proceedings of the 27th ACM symposium on Parallelism in Algorithms and Architectures*. 254–263.
- [15] Andreia Correia, Pedro Ramalhete, and Pascal Felber. 2021. OrcGC: automatic lock-free memory reclamation. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 205–218.
- [16] David L Detlefs, Paul A Martin, Mark Moir, and Guy L Steele Jr. 2002. Lock-free reference counting. *Distributed Computing* 15, 4 (2002), 255–271.
- [17] Dave Dice, Maurice Herlihy, and Alex Kogan. 2016. Fast non-intrusive memory reclamation for highly-concurrent data structures. In *Proceedings of the 2016 ACM SIGPLAN International Symposium on Memory Management*. 36–45.
- [18] Faith Ellen, Panagioti Fatourou, Eric Ruppert, and Franck van Breugel. 2010. Non-blocking binary search trees. In *Proceedings of the 29th ACM SIGACT-SIGOPS symposium on Principles of distributed computing*. 131–140.
- [19] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. 1985. Impossibility of Distributed Consensus with One Faulty Process. *J. ACM* 32, 2 (apr 1985), 374–382. <https://doi.org/10.1145/3149.214121>
- [20] Keir Fraser. 2004. *Practical lock-freedom*. Technical Report. University of Cambridge, Computer Laboratory.
- [21] Anders Gidenstam, Marina Papatriantafyllou, Håkan Sundell, and Philippas Tsigas. 2008. Efficient and reliable lock-free memory reclamation based on reference counting. *IEEE Transactions on Parallel and Distributed Systems* 20, 8 (2008), 1173–1187.
- [22] Wojciech Golab. 2018. Recoverable consensus in shared memory. *arXiv preprint arXiv:1804.10597* (2018).
- [23] Alexey Gotsman, Noam Rinetzy, and Hongseok Yang. 2013. Verifying concurrent memory reclamation algorithms with grace. In *European Symposium on Programming*. Springer, 249–269.
- [24] Timothy L Harris. 2001. A pragmatic implementation of non-blocking linked-lists. In *International Symposium on Distributed Computing*. Springer, 300–314.
- [25] Steve Heller, Maurice Herlihy, Victor Luchangco, Mark Moir, William N Scherer, and Nir Shavit. 2005. A lazy concurrent list-based set algorithm. In *International Conference On Principles Of Distributed Systems*. Springer, 3–16.
- [26] Maurice Herlihy. 1991. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 13, 1 (1991), 124–149.
- [27] Maurice Herlihy, Victor Luchangco, Paul Martin, and Mark Moir. 2005. Non-blocking memory management support for dynamic-sized data structures. *ACM Transactions on Computer Systems (TOCS)* 23, 2 (2005), 146–196.
- [28] Maurice Herlihy and Nir Shavit. 2011. On the nature of progress. In *International Conference On Principles Of Distributed Systems*. Springer, 313–328.
- [29] Maurice Herlihy, Nir Shavit, Victor Luchangco, and Michael Spear. 2020. *The art of multiprocessor programming*. Newnes.
- [30] Maurice P Herlihy and Jeannette M Wing. 1990. Linearizability: A correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 12, 3 (1990), 463–492.
- [31] Shane V Howley and Jeremy Jones. 2012. A non-blocking internal binary search tree. In *Proceedings of the twenty-fourth annual ACM symposium on Parallelism in algorithms and architectures*. 161–171.
- [32] Jeehoon Kang and Jaehwang Jung. 2020. A marriage of pointer and epoch-based reclamation. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. 314–328.
- [33] Roland Meyer and Sebastian Wolff. 2019. Decoupling lock-free data structures from memory reclamation for static analysis. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 1–31.
- [34] Roland Meyer and Sebastian Wolff. 2019. Pointer life cycle types for lock-free data structures with memory reclamation. *Proceedings of the ACM on Programming Languages* 4, POPL (2019), 1–36.
- [35] Maged M Michael. 2002. High performance dynamic lock-free hash tables and list-based sets. In *Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures*. 73–82.
- [36] Maged M Michael. 2004. ABA prevention using single-word instructions. *IBM Research Division, RC23089 (W0401-136), Tech. Rep* (2004).
- [37] Maged M Michael. 2004. Hazard pointers: Safe memory reclamation for lock-free objects. *IEEE Transactions on Parallel and Distributed Systems* 15, 6 (2004),

- 491–504.
- [38] Aravind Natarajan and Neeraj Mittal. 2014. Fast concurrent lock-free binary search trees. In *Proceedings of the 19th ACM SIGPLAN symposium on Principles and practice of parallel programming*. 317–328.
  - [39] Ruslan Nikolaev and Binoy Ravindran. 2020. Universal wait-free memory reclamation. In *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 130–143.
  - [40] Ruslan Nikolaev and Binoy Ravindran. 2021. Snapshot-free, transparent, and robust memory reclamation for lock-free data structures. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*. 987–1002.
  - [41] 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*. 367–369.
  - [42] Gali Sheffi, Maurice Herlihy, and Erez Petrank. 2021. VBR: Version Based Reclamation. In *35th International Symposium on Distributed Computing, DISC 2021, October 4–8, 2021, Freiburg, Germany (Virtual Conference) (LIPIcs, Vol. 209)*, Seth Gilbert (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 35:1–35:18. <https://doi.org/10.4230/LIPIcs.DISC.2021.35>
  - [43] Gali Sheffi, Maurice Herlihy, and Erez Petrank. 2021. Vbr: Version based reclamation. In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures*. 443–445.
  - [44] Gali Sheffi and Erez Petrank. 2020. Functional faults. In *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures*. 453–463.
  - [45] Gali Sheffi and Erez Petrank. 2022. *The ERA Theorem for Safe Memory Reclamation*. Technical Report. Full version of this paper.
  - [46] Ajay Singh, Trevor Brown, and Ali Mashtizadeh. 2021. NBR: neutralization based reclamation. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 175–190.
  - [47] Daniel Solomon and Adam Morrison. 2021. Efficiently reclaiming memory in concurrent search data structures while bounding wasted memory. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 191–204.
  - [48] Håkan Sundell. 2005. Wait-free reference counting and memory management. In *19th IEEE International Parallel and Distributed Processing Symposium*. IEEE, 10–pp.
  - [49] Shahar Timnat, Anastasia Braginsky, Alex Kogan, and Erez Petrank. 2012. Wait-free linked-lists. In *International Conference On Principles Of Distributed Systems*. Springer, 330–344.
  - [50] Shahar Timnat and Erez Petrank. 2014. A practical wait-free simulation for lock-free data structures. *ACM SIGPLAN Notices* 49, 8 (2014), 357–368.
  - [51] Yuanhao Wei, Naama Ben-David, Guy E Blelloch, Panagiota Fatourou, Eric Ruppert, and Yihan Sun. 2021. Constant-time snapshots with applications to concurrent data structures. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 31–46.
  - [52] Haosen Wen, Joseph Izraelevitz, Wentao Cai, H Alan Beadle, and Michael I Scott. 2018. Interval-based memory reclamation. *ACM SIGPLAN Notices* 53, 1 (2018), 1–13.