

Understanding and Effectively Preventing the ABA Problem in Descriptor-based Lock-free Designs

Damian Dechev

Sandia National Laboratories
Scalable Computing R & D Department
Livermore, CA 94551-0969
ddechev@sandia.gov

Peter Pirkelbauer

Texas A&M University
Department of Computer Science
College Station, TX 77843-3112
peter.pirkelbauer@tamu.edu

Bjarne Stroustrup

Texas A&M University
Department of Computer Science
College Station, TX 77843-3112
bs@cse.tamu.edu

Abstract

An increasing number of modern real-time systems and the nowadays ubiquitous multicore architectures demand the application of programming techniques for reliable and efficient concurrent synchronization. Some recently developed Compare-And-Swap (CAS) based nonblocking techniques hold the promise of delivering practical and safer concurrency. The ABA¹ problem is a fundamental problem to many CAS-based designs. Its significance has increased with the suggested use of CAS as a core atomic primitive for the implementation of portable lock-free algorithms. The ABA problem's occurrence is due to the intricate and complex interactions of the application's concurrent operations and, if not remedied, ABA can significantly corrupt the semantics of a nonblocking algorithm. The current state of the art leaves the elimination of the ABA hazards to the ingenuity of the software designer. In this work we provide the first systematic and detailed analysis of the ABA problem in lock-free Descriptor-based designs. We study the semantics of Descriptor-based lock-free data structures and propose a classification of their operations that helps us better understand the ABA problem and subsequently derive an effective ABA prevention scheme. Our ABA prevention approach outperforms by a large factor the use of the alternative CAS-based ABA prevention schemes. It offers speeds comparable to the use of the architecture-specific CAS2 instruction used for version counting. We demonstrate our ABA prevention scheme by integrating it into an advanced nonblocking data structure, a lock-free dynamically resizable array.

Keywords: ABA problem, nonblocking synchronization, lock-free programming techniques

¹ABA is not an acronym and is a shortcut for stating that a value at a shared location can change from A to B and then back to A

1 Introduction

The modern ubiquitous multi-core architectures demand the design of programming libraries and tools that allow fast and reliable concurrency. In addition, providing safe and efficient concurrent synchronization is of critical importance to the engineering of many modern real-time systems. Lock-free programming techniques [11] have been demonstrated to be effective in delivering performance gains and preventing some hazards, typically associated with the application of mutual exclusion, such as deadlock, livelock, and priority inversion [5], [2]. The ABA problem is a fundamental problem to many CAS-based nonblocking designs. Avoiding the hazards of ABA imposes an extra challenge for a lock-free algorithm's design and implementation. To the best of our knowledge, the literature does not offer an explicit and detailed analysis of the ABA problem, its relation to the most commonly applied nonblocking programming techniques (such as the use of Descriptors) and correctness guarantees, and the possibilities for its avoidance. Thus, at the present moment of time, eliminating the hazards of ABA in a nonblocking algorithm is left to the ingenuity of the software designer. In this work we study in details and define the conditions that lead to ABA in a nonblocking Descriptor-based design. Based on our analysis, we define a generic and practical condition, called the $\lambda\delta$ approach, for ABA avoidance for a lock-free Descriptor-based linearizable design (Section 4). We demonstrate the application of our approach by incorporating it in a complex and advanced nonblocking data structure, a lock-free dynamically resizable array (vector) [2]. The ISO C++ Standard Template Library [17] vector offers a combination of dynamic memory management and constant-time random access. We survey the literature for other known ABA prevention techniques (usually described as a part of a nonblocking algorithm's implementation) and study in detail three known solutions to the ABA problem (Sections 2.1

and 2.3). Our performance evaluation (Section 5) establishes that the single-word CAS-based $\lambda\delta$ approach is fast, efficient, and practical.

2 The ABA Problem

The Compare-And-Swap (CAS) atomic primitive (commonly known as Compare and Exchange, `CMPXCHG`, on the Intel *x86* and *Itanium* architectures [12]) is a CPU instruction that allows a processor to atomically test and modify a single-word memory location. The application of a CAS-controlled speculative manipulation of a shared location (L_i) is a fundamental programming technique in the engineering of nonblocking algorithms [11] (an example is shown in Algorithm 1).

Algorithm 1 CAS-based speculative manipulation of L_i

```

1: repeat
2:   value.type  $A_i = \hat{L}_i$ 
3:   value.type  $B_i = \text{fComputeB}$ 
4: until CAS( $L_i, A_i, B_i$ ) ==  $B_i$ 

```

In our pseudocode we use the symbols $\hat{\cdot}$, $\&$, and \cdot to indicate pointer dereferencing, obtaining an object’s address, and integrated pointer dereferencing and field access. When the value stored at L_i is the target value of a CAS-based speculative manipulation, we call L_i and \hat{L}_i *control location* and *control value*, respectively. We indicate the control value’s type with the string `value.type`. The size of `value.type` must be equal or less than the maximum number of bits that a hardware CAS instruction can exchange atomically (typically the size of a single memory word). In the most common cases, `value.type` is either an integer or a pointer value. In Algorithm 1, the function `fComputeB` yields the new value, B_i , to be stored at L_i .

Definition 1: *The ABA problem is a false positive execution of a CAS-based speculation on a shared location L_i .*

As illustrated in Table 1, ABA can occur if a process P_1 is interrupted at any time after it has read the old value (A_i) and before it attempts to execute the CAS instruction from Algorithm 1. An interrupting process (P_k) might change the value at L_i to B_i . Afterwards, either P_k or any other process $P_j \neq P_1$ can eventually store A_i back to L_i . When P_1 resumes, its CAS loop succeeds (false positive execution) despite the fact that L_i ’s value has been meanwhile manipulated.

Definition 2: *A nonblocking algorithm is ABA-free when its semantics cannot be corrupted by the occurrence of ABA.*

ABA-freedom is achieved when: *a)* occurrence of ABA is harmless to the algorithm’s semantics or *b)* ABA is avoided. The former scenario is uncommon and strictly specific to the algorithm’s semantics. The latter scenario is the general case and in this work we focus on providing details of how to eliminate ABA.

Step	Action
Step 1	P_1 reads A_i from L_i
Step 2	P_k interrupts P_1 ; P_k stores the value B_i into L_i
Step 3	P_j stores the value A_i into L_i
Step 4	P_1 resumes; P_1 executes a false positive CAS

Table 1. ABA at L_i

2.1 Known ABA Avoidance Techniques I

A general strategy for ABA avoidance is based on the fundamental guarantee that no process P_j ($P_j \neq P_1$) can possibly store A_i again at location L_i (Step 3, Table 1). One way to satisfy such a guarantee is to require all values stored in a given control location to be *unique*. To enforce this uniqueness invariant we can place a constraint on the user and request each value stored at L_i to be used only once (*Known Solution 1*). For a large majority of concurrent algorithms, enforcing uniqueness typing would not be a suitable solution since their applications imply the usage of a value or reference more than once.

An alternative approach to satisfying the uniqueness invariant is to apply a *version tag* attached to each value. The usage of version tags is the most commonly cited solution for ABA avoidance [6]. The approach is effective, when it is possible to apply, but suffers from a significant flaw: a single-word CAS is insufficient for the atomic update of a word-sized control value and a word-sized version tag. An effective application of a version tag [3] requires the hardware architecture to support a more complex atomic primitive that allows the atomic update of two memory locations, such as CAS2 (compare-and-swap two co-located words) or DCAS (compare-and-swap two memory locations). The availability of such atomic primitives might lead to much simpler, elegant, and efficient concurrent designs (in contrast to a CAS-based design). It is not desirable to suggest a CAS2/DCAS-based ABA solution for a CAS-based algorithm, unless the implementor explores the optimization possibilities of the algorithm upon the availability of CAS2/DCAS. A proposed hardware implementation (entirely built into a present cache coherence protocol) of an innovative Alert-On-Update (AOU) instruction [16] has been suggested by Spear et al. to eliminate the CAS deficiency of allowing ABA. Some suggested approaches [15] split a version counter into two half-words (*Known Solution 2*): a half-word used to store the control value and a half-word used as a version tag. Such techniques lead to severe limitations on the addressable memory space and the number of possible writes into the shared location. To guarantee the uniqueness invariant of a control value of type pointer in a concurrent system with dynamic memory usage, we face an extra challenge: even if we write a pointer value no

more than once in a given control location, the memory allocator might reuse the address of an already freed object (A_i) and pose an ABA hazard. To prevent this scenario, all control values of pointer type must be guarded by a concurrent nonblocking garbage collection scheme such as Hazard Pointers [14] (that uses a list of hazard pointers per thread) or Herlihy et al.’s Pass The Buck algorithm [10] (that utilizes a dedicated thread to periodically reclaim unguarded objects). While enhancing the safety of a concurrent algorithm (when needed), the application of a complementary garbage collection mechanism might come at a significant performance cost (Section 5).

2.2 The Descriptor Object

Linearizability is an important correctness condition for concurrent objects [11]. A concurrent operation is linearizable if it appears to execute instantaneously in a given point of time τ_{lin} between the time τ_{inv} of its invocation and the time τ_{end} of its completion. The literature often refers to τ_{lin} as a *linearization point*. The implementations of many nonblocking data structures require the update of *two or more memory locations* in a linearizable fashion [2], [5]. The engineering of such operations (e.g. `push_back` and `resize` in a dynamically resizable array) is particularly challenging in a CAS-based design. A common programming technique applied to guarantee the linearizability requirements for such operations is the use of a *Descriptor Object* (δ object) [2], [5]. The pseudocode in Algorithm 2 shows the generalized two-step execution of a Descriptor Object. Our definition of a Descriptor Object requires the Descriptor to store three types of information:

- (1) Global data describing the state of the shared container ($v\delta$), e.g. the `size` of a dynamically resizable array [2].
- (2) A record of a pending operation on a given memory location. We call such a record requesting an update at a shared location L_i from an old value, `old_val`, to a new value, `new_val`, a *Write Descriptor* ($\omega\delta$). The shorthand notation we use is $\omega\delta @ L_i : \text{old_val} \rightarrow \text{new_val}$. The fields in the Write Descriptor Object store the target location as well as the old and the new values.
- (3) A boolean value indicating whether $\omega\delta$ contains a pending write operation that needs to be completed.

The use of a Descriptor allows an interrupting thread to help the interrupted thread complete an operation rather than wait for its completion. As shown in Algorithm 2, the technique is used to implement, using only two CAS instructions, a linearizable update of two memory locations: 1. a reference to a Descriptor Object (data type pointer to δ stored in a location L_δ) and 2. an element of type

`value_type` stored in L_i . In Step 1, Algorithm 2, we perform a CAS-based speculation of a shared location L_δ that contains a reference to a Descriptor Object. Step 1 executes in the following fashion:

1. we read the value of the current δ reference stored in L_δ (line 3)
2. if the current δ object contains a pending operation, we need to help its completion (lines 4-5)
3. we record the current value, A_i , in location L_i (line 6) and compute the new value, B_i , to be stored in L_i (line 7)
4. a new $\omega\delta$ object is allocated on the heap, initialized (by calling $f_{\omega\delta}$), and its fields `Target`, `OldValue`, and `NewValue` are set (lines 8-11)
5. any state carrying data stored in a Descriptor Object must be computed (by calling $f_{v\delta}$). Such data might be a shared element or a container’s size (line 12)
6. a new Descriptor Object is initialized containing the new Write Descriptor and the new descriptor’s data. The new descriptor’s *pending operation* flag (`WDpending`) is set to `true` (lines 13-14)
7. we attempt a swap of the old Descriptor Object with the new one (line 15). Should the CAS fail, we know that there is another process that has interrupted us and meanwhile succeeded to modify L_δ and progress. We need to go back at the beginning of the loop and repeat all the steps. Should the CAS succeed, we proceed with Step 2 and perform the update at L_i .

The size of a Descriptor Object is larger than a memory word. Thus, we need to store and manipulate a Descriptor Object through a reference. Since the control value of Step 1 stores a pointer to a Descriptor Object, to prevent ABA at L_δ , all references to descriptors must be memory managed by a safe nonblocking garbage collection scheme. We use the prefix μ for all variables that require safe memory management. In Step 2 we execute the Write Descriptor, `WD`, in order to update the value at L_i . Any interrupting thread (after the completion of Step 1) detects the pending flag of $\omega\delta$ and, should the flag’s value be still positive, it proceeds to executing the requested update $\omega\delta @ L_i : A_i \rightarrow B_i$. There is no need to execute a CAS-based loop and the call to a single CAS is sufficient for the completion of $\omega\delta$. Should the CAS from Step 2 succeed, we have completed the two-step execution of the Descriptor Object. Should it fail, we know that there is an interrupting thread that has completed it already.

Algorithm 2 Two-step execution of a δ object

```
1: Step 1: place a new descriptor in  $L_\delta$ 
2: repeat
3:    $\delta \mu\text{OldDesc} = \hat{\sim} L_\delta$ 
4:   if  $\mu\text{OldDesc.WDPending} == \text{true}$  then
5:     execute  $\mu\text{OldDesc.WD}$ 
6:   value.type  $A_i = \hat{\sim} L_i$ 
7:   value.type  $B_i = \text{fComputeB}$ 
8:    $\omega\delta \text{WD} = \text{f}\omega\delta()$ 
9:    $\text{WD.Target} = L_i$ 
10:   $\text{WD.OldElement} = A_i$ 
11:   $\text{WD.NewElement} = B_i$ 
12:   $v\delta \text{DescData} = \text{f}v\delta()$ 
13:   $\delta \mu\text{NewDesc} = \text{f}\delta(\text{DescData}, \text{WD})$ 
14:   $\mu\text{NewDesc.WDPending} = \text{true}$ 
15: until  $\text{CAS}(L_\delta, \mu\text{OldDesc}, \mu\text{NewDesc}) == \mu\text{NewDesc}$ 
16:
17: Step 2: execute the write descriptor
18: if  $\mu\text{NewDesc.WDPending}$  then
19:    $\text{CAS}(\text{WD.Target}, \text{WD.OldElement}, \text{WD.NewElement}) == \text{WD.NewElement}$ 
20:    $\mu\text{NewDesc.WDPending} = \text{false}$ 
```

2.3 Known ABA Avoidance Techniques II

A known approach for avoiding a false positive execution of the Write Descriptor from Algorithm 2 is the application of *value semantics* for all values of type `value.type` (*Known Solution 3*). As discussed in [9] and [2], an ABA avoidance scheme based on value semantics relies on:

- a. *Extra level of indirection*: all values are stored in shared memory indirectly through pointers. Each write of a given value v_i to a shared location L_i needs to allocate on the heap a new reference to v_i (η_{v_i}), store η_{v_i} into L_i , and finally safely delete the pointer value removed from L_i .
- b. *Nonblocking garbage collection (GC)*: all references stored in shared memory (such as η_{v_i}) need to be safely managed by a nonblocking garbage collection scheme (e.g. Hazard Pointers, Pass The Buck).

As reflected in our performance test results (Section 5), the usage of both an extra level of indirection as well as the heavy reliance on a nonblocking GC scheme for managing the Descriptor Objects *and* the references to `value.type` objects is very expensive with respect to the space and time complexity of a nonblocking algorithm. However, the use of value semantics is the **only known approach** for ABA avoidance in the execution of a Descriptor Object. In Section 4 we present a 3-step execution approach that helps us eliminate ABA, avoid the need for an extra level of indirection, and reduce the usage of the computationally expensive GC scheme.

3 Descriptor-based Operations Classification

The use of a Descriptor Object provides the programming technique for the implementation of some of the complex nonblocking operations in a shared container, such as

the `push_back`, `pop_back`, and `reserve` operations in a shared vector [2]. The use and execution of a Write Descriptor guarantees the linearizable update of two or more memory locations. Here, to better understand the interactions among these operations and the cause of ABA, we classify the operations in a nonblocking Descriptor-based design.

Definition 3: *An operation whose success depends on the creation and execution of a Write Descriptor is called an $\omega\delta$ -executing operation.*

The operation `push_back` of a shared vector [2] is an example of an $\omega\delta$ -executing operation. Such $\omega\delta$ -executing operations have *lock-free* semantics and the progress of an individual operation is subject to the contention on the shared location L_i . For a shared vector, operations such as `pop_back` do not need to execute a Write Descriptor Object [2]. Their progress is dependent on the state of the global data stored in the Descriptor Object, such as the size of a container.

Definition 4: *An operation whose success depends on the state of the $v\delta$ data stored in the Descriptor Object is a δ -modifying operation.*

A δ -modifying operation, such as `pop_back`, needs only update the shared global data (the `size` of type $v\delta$) in the Descriptor Object. Since an $\omega\delta$ -executing operation by definition always performs an exchange of the entire Descriptor Object, every $\omega\delta$ -executing operation is also δ -modifying. The semantics of a δ -modifying operation are *lock-free* and the progress of an individual operation is determined by the interrupts by other δ -modifying operations. An $\omega\delta$ -executing operation is also δ -modifying but as is the case with `pop_back`, not all δ -modifying operations are $\omega\delta$ -executing. Certain operations, such as the random access `read` and `write` in a vector [2], do not need to access the Descriptor Object and progress regardless of the state of the descriptor. Such operations are non- δ -modifying and have *wait-free* semantics (thus no delay if there is contention at L_δ).

Definition 5: *An operation whose success does not depend on the state of the Descriptor Object is a non- δ -modifying operation.*

3.1 Concurrent Operations

When two δ -modifying operations (O_{δ_1} and O_{δ_2}) are concurrent [11], according to Algorithm 2, O_{δ_1} precedes O_{δ_2} in the linearization history if and only if O_{δ_1} completes Step 1, Algorithm 2 prior to O_{δ_2} .

Definition 6: *We refer to the instant of successful execution of the global Descriptor exchange at L_δ (line 15, Algorithm 2) as τ_δ .*

Definition 7: *A point in the execution of a δ object that determines the order of an $\omega\delta$ -executing operation acting*

on location L_i relative to other writer operations acting on the same location L_i , is referred to as the $\lambda\delta$ -point ($\tau_{\lambda\delta}$) of a Write Descriptor.

The order of execution of the $\lambda\delta$ -points of two concurrent $\omega\delta$ -executing operations determines their order in the linearization history. Let us designate the point of time

Step	Action
Step 1	$O_{\delta_1}: \tau_{read\delta}$
Step 2	$O_{\delta_1}: \tau_{access_i}$
Step 3	$O_{\delta_1}: \tau_{\delta}$
Step 4	$O_{\delta_2}: \tau_{read\delta}$
Step 5	$O_{\delta_1}: \tau_{wd}$
Step 6	$O_i: \tau_{write_i}$
Step 7	$O_{\delta_2}: \tau_{wd}$

Table 2. ABA occurrence in the execution of a Descriptor Object

when a certain δ -modifying operation reads the state of the Descriptor Object by $\tau_{read\delta}$, and the instants when a thread reads a value from and writes a value into a location L_i by τ_{access_i} and τ_{write_i} , respectively. Table 2 demonstrates the occurrence of ABA in the execution of a δ object with two concurrent δ -modifying operations (O_{δ_1} and O_{δ_2}) and a concurrent write, O_i , to L_i . We assume that the δ object’s implementation follows Algorithm 2. The placement of the $\lambda\delta$ -point plays a critical role for achieving ABA safety in the implementation of an $\omega\delta$ -executing operation. As shown in Table 2, at time τ_{wd} when O_{δ_2} executes the write descriptor, O_{δ_2} has no way of knowing whether O_{δ_1} has completed its update at L_i or not. Since O_{δ_1} ’s $\lambda\delta$ -point $\equiv \tau_{\delta}$, the only way to know about the status of O_{δ_1} is to read L_δ . Using a single-word CAS operation prevents O_{δ_2} from atomically checking the status of L_δ and executing the update at L_i .

Definition 8: A concurrent execution of one or more non- $\omega\delta$ -executing δ -modifying operations with one $\omega\delta$ -executing operation, O_{δ_1} , performing an update at location L_i is ABA-free if O_{δ_1} ’s $\lambda\delta$ -point $\equiv \tau_{access_i}$. We refer to an $\omega\delta$ -executing operation where its $\lambda\delta$ -point $\equiv \tau_{access_i}$ as a $\lambda\delta$ -modifying operation.

Assume that in Table 2 the O_{δ_1} ’s $\lambda\delta$ -point $\equiv \tau_{access_i}$. As shown in Table 2, the ABA problem in this scenario occurs when there is a hazard of a spurious execution of O_{δ_1} ’s Write Descriptor. Having a $\lambda\delta$ -modifying implementation of O_{δ_1} allows any non- $\omega\delta$ -executing δ -modifying operation such as O_{δ_2} to check O_{δ_1} ’s progress while attempting the atomic update at L_i requested by O_{δ_1} ’s Write Descriptor. Our 3-step descriptor execution approach, discussed in Section 4, offers a solution based on Definition 8. In an implementation with two or more concurrent $\omega\delta$ -executing oper-

ations, each $\omega\delta$ -executing operation must be $\lambda\delta$ -modifying in order to eliminate the hazard of a spurious execution of an $\omega\delta$ that has been picked up by a collaborating operation.

4 ABA-free Execution of the Descriptor Object

In Algorithm 3 we suggest a design strategy for the implementation of a $\lambda\delta$ -modifying operation. Our approach is based on a 3-step execution of the δ object. While similar to Algorithm 2, the approach shown in Algorithm 3 differs by executing a fundamental additional step: in Step 1 we store a pointer to the new descriptor in L_i prior to the attempt to store it in L_δ in Step 2. Since all δ objects are memory managed, we are guaranteed that no other thread would attempt a write of the value $\mu\text{NewDesc}$ in L_i or any other shared memory location. The operation is $\lambda\delta$ -modifying because, after the new descriptor is placed in L_i , any interrupting writer thread accessing L_i is aware of the Write Descriptor stored at L_i and is required to complete the remaining two steps in the execution of the Write Descriptor. However, should the CAS execution in Step 2 (line 26) fail, we have to unroll the changes at L_i performed in Step 1 by restoring L_i ’s old value preserved in WD.OldElement (line 20) and retry the execution of the routine (line 21). To implement Algorithm 3, we have to be able to distinguish between objects of type `value_type` and δ . A possible solution is to require that all `value_type` variables are pointers and all pointer values stored in L_i are aligned with the two low-order bits cleared during their initialization. That way, we can use the two low-order bits for designating the type of the pointer values. Subsequently, every read must check the type of the pointer obtained from a shared memory location prior to manipulating it. Once an operation succeeds at completing Step 1, Algorithm 3, location L_i contains a pointer to a δ object that includes both: L_i ’s previous value of type `value_type` and a write descriptor WD that provides a record for the steps necessary for the operation’s completion. Any non- δ -modifying operation, such as a random access read in a shared vector, can obtain the value of L_i (of type `value_type`) by accessing WD.OldElement (thus going through a *temporary* indirection) and ignore the Descriptor Object. Upon the success of Step 3, Algorithm 3, the temporary level of indirection is eliminated. Such an approach would preserve the wait-free execution of a non- δ -modifying operation. The $\omega\delta$ data type needs to be amended to include a field `TempElement` (line 9, Algorithm 3) that records the value of the temporary δ pointer stored in L_i . The cost of the $\lambda\delta$ operation is 3 CAS executions to achieve the linearizable update of two shared memory locations (L_i and L_δ). The implementation of our $\lambda\delta$ -modifying operation as shown in Algorithm 3 is similar to the execution of Harris et al.’s *MCAS* algorithm [8]. Just like our $\lambda\delta$ -

modifying approach, for an *MCAS* update of L_δ and L_i , the cost of Harris et al.’s *MCAS* is at least 3 executions of the single-word CAS instruction. Harris et al.’s work on *MCAS* [8] brings forward a significant contribution in the design of lock-free algorithms, however, it lacks any analysis of the hazards of ABA and the way the authors manage to avoid it.

Algorithm 3 Implementing a $\lambda\delta$ -modifying operation through a three-step execution of a δ object

```

1: Step 1: place a new descriptor in  $L_i$ 
2: value.type  $B_i = f_{\text{ComputeB}}$ 
3: value.type  $A_i$ 
4:  $\omega\delta$   $WD = f_{\omega\delta}()$ 
5:  $WD.Target = L_i$ 
6:  $WD.NewElement = B_i$ 
7:  $v\delta$   $DescData = f_{v\delta}()$ 
8:  $\delta$   $\mu\text{NewDesc} = f_\delta(DescData, WD)$ 
9:  $WD.TempElement = \&\mu\text{NewDesc}$ 
10:  $\mu\text{NewDesc}.WDpending = \text{true}$ 
11: repeat
12:    $A_i = \sim L_i$ 
13:    $WD.OldElement = A_i$ 
14: until  $\text{CAS}(L_i, A_i, \mu\text{NewDesc}) == \mu\text{NewDesc}$ 
15:
16: Step 2: place the new descriptor in  $L_\delta$ 
17: bool unroll = false
18: repeat
19:   if unroll then
20:      $\text{CAS}(WD.Target, \mu\text{NewDesc}, WD.OldElement)$ 
21:     goto 3
22:    $\delta$   $\mu\text{OldDesc} = \sim L_\delta$ 
23:   if  $\mu\text{OldDesc}.WDpending == \text{true}$  then
24:     execute  $\mu\text{OldDesc}.WD$ 
25:     unroll = true
26: until  $\text{CAS}(L_\delta, \mu\text{OldDesc}, \mu\text{NewDesc}) == \mu\text{NewDesc}$ 
27:
28: Step 3: execute the Write Descriptor
29: if  $\mu\text{NewDesc}.WDpending$  then
30:    $\text{CAS}(WD.Target, WD.TempElement, WD.NewElement) == WD.NewElement$ 
31:    $\mu\text{NewDesc}.WDpending = \text{false}$ 

```

5 Performance Evaluation

To evaluate the performance of the ABA-free programming techniques discussed in this work, we incorporated the presented ABA elimination approaches in the implementation of a nonblocking dynamically resizable array [2]. Our test results indicate that the $\lambda\delta$ approach offers ABA prevention with performance comparable to the use of the platform-specific CAS2 instruction to implement version counting. This finding is of particular value to the engineering of some embedded real-time systems where the hardware does not support complex atomic primitives such as CAS2 [13]. We ran performance tests on an Intel IA-32 SMP machine with two 1.83GHz processor cores with 512 MB shared memory and 2 MB L2 shared cache running the MAC 10.5.6 operating system. In our performance analysis we compare:

(1) *$\lambda\delta$ approach*: the implementation of a vector with a $\lambda\delta$ -modifying `push_back` and a δ -modifying `pop_back`.

In this scenario the cost of `push_back` is 3 single-word CAS operations and `pop_back`’s cost is one single-word CAS instruction.

(2) *All-GC approach*: the application of Known Solution 3 (Section 2.3), namely the use of an extra level of indirection and memory management for each element. Because of its performance and availability, we have chosen to implement and apply Herlihy et al.’s Pass The Buck algorithm [10]. In addition, we use Pass The Buck to protect the Descriptor Objects for all of the tested approaches.

(3) *CAS2-based approach*: the application of CAS2 for maintaining a reference counter for each element. A CAS2-based version counting implementation is easy to apply to almost any pre-existent CAS-based algorithm. While a CAS2-based solution is not portable, we believe that the approach is applicable for a large number of modern architectures. For this reason, it is included in our performance evaluation. In the performance tests, we apply CAS2 (and version counting) for updates at the shared memory locations at L_i and a single-word CAS to update the Descriptor Object at L_δ .

Similarly to the evaluation of other lock-free algorithms [4], we designed our experiments by generating a workload of the various operations. We varied the number of threads, starting from 1 and exponentially increased their number to 64. Each thread executed 500,000 lock-free operations on the shared container. We measured the execution time (in seconds) that all threads needed to complete. Each iteration of every thread executed an operation with a certain probability (`push_back` (+), `pop_back` (-), random access write (w), random access read (r)). We show the performance graph for a distribution of +:40%, -:40%, w:10%, r:10% on Figure 1. Figure 2 demonstrates the performance results with less contention at the vector’s tail, +:25%, -:25%, w:10%, r:40%. Figure 3 illustrates the test results with a distribution containing predominantly random access read and write operations, +:10%, -:10%, w:40%, r:40%. Figure 4 reflects our performance evaluation on a vector’s use with mostly random access read operations: +:20%, -:0%, w:20%, r:60%, a scenario often referred to as the most common real-world use of a shared container [4]. The number of threads is plotted along the x -axis, while the time needed to complete all operations is shown along the y -axis. According to the performance results, compared to the *All-GC* approach, the $\lambda\delta$ approach delivers consistent performance gains in all possible operation mixes by a large factor, a factor of at least 3.5 in the cases with less contention at the tail and a factor of 10 or more when there is a high concentration of tail operations.

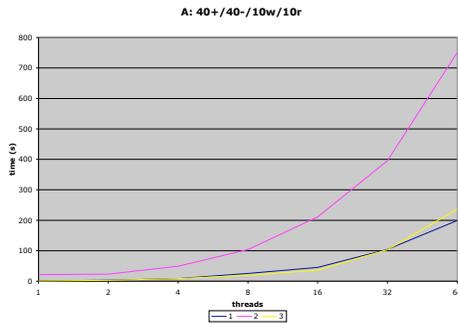


Figure 1. Performance Results A

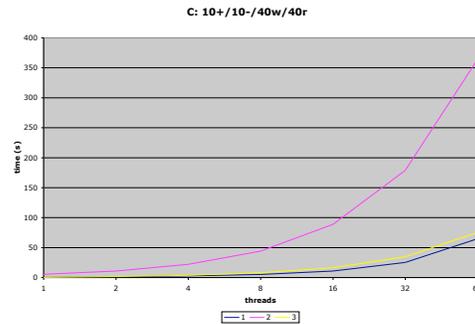


Figure 3. Performance Results C

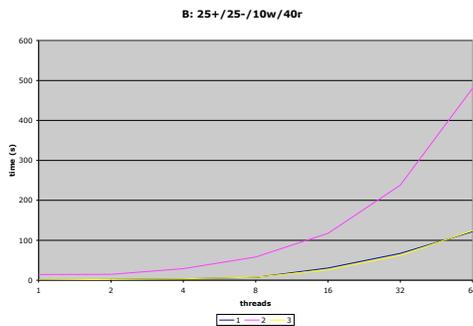


Figure 2. Performance Results B

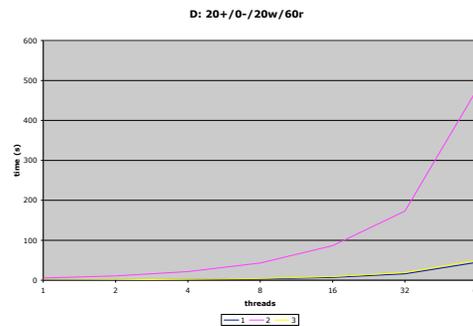


Figure 4. Performance Results D

These observations come as a confirmation to our expectations that introducing an extra level of indirection and the necessity to memory manage each individual element with PTB (or an alternative memory management scheme) to avoid ABA comes with a pricy performance overhead. The $\lambda\delta$ approach offers an alternative by introducing the notion of a $\lambda\delta$ -point and enforces it through a 3-step execution of the δ object. The application of version counting based on the architecture-specific CAS2 operation is the most commonly cited approach for ABA prevention in the literature. Our performance evaluation shows that the $\lambda\delta$ approach delivers performance comparable to the use of CAS2-based version counting. CAS2 is a complex atomic primitive and its application comes with a higher cost when compared to the application of atomic write or a single-word CAS. In the performance tests we executed, we notice that in the scenarios where random access write is invoked more frequently (Figures 3 and 4), the performance of the CAS2 version counting approach suffers a performance penalty and runs slower than the $\lambda\delta$ approach by about 12% to 20%. Ac-

ording to our performance evaluation, the $\lambda\delta$ approach is a systematic, effective, portable, and generic solution for ABA avoidance for Descriptor-based nonblocking designs. The $\lambda\delta$ scheme does not induce a performance penalty when compared to the architecture-specific application of CAS2-based version counting and offers a considerable performance gain when compared to the use of *All-GC*.

6 Conclusion and Future Work

In this work we studied the ABA problem and the conditions leading to its occurrence in a Descriptor-based lock-free linearizable design. We offered a systematic and generic solution, called the $\lambda\delta$ approach, that outperforms by a significant factor the use of garbage collection for the safe management of each shared location and offers speed of execution comparable to the application of the architecture-specific CAS2 instruction used for version counting. Having a practical alternative to the application of the architecture-specific CAS2 is of particular importance

to the design of some modern embedded systems [13]. We defined a condition for ABA-free synchronization that allows us to reason about the ABA safety of a lock-free algorithm. We presented a practical, generic, and portable implementation of the $\lambda\delta$ approach and evaluated it by integrating the $\lambda\delta$ technique into a nonblocking shared vector. The literature *does not offer a detailed analysis of the ABA problem* and the general techniques for its avoidance in a lock-free linearizable design. At the present moment of time, the challenges of eliminating ABA are left to the ingenuity of the software designer. The goal of our work is to deliver a guide for ABA comprehension and prevention in Descriptor-based lock-free linearizable algorithms. For the practical application of Descriptor-based nonblocking techniques in real-time systems, it is important to study the service-time bounds of the operations within the context of the Descriptor's CAS-based retry loop. Anderson et al. [1] present a fundamental approach for such formal timing analysis. In our future work we plan to utilize a model-checker [7] to express the $\lambda\delta$ condition as well as apply Anderson et al.'s [1] approach to derive the timing guarantees for our ABA prevention approach.

7 Acknowledgements

We thank the anonymous referees for their detailed and helpful comments on this work.

References

- [1] James Anderson, Srikanth Ramamurthy, and Kevin Jeffay. Real-time computing with lock-free shared objects. *ACM Trans. Comput. Syst.*, 15(2):134–165, 1997.
- [2] Damian Dechev, Peter Pirkelbauer, and Bjarne Stroustrup. Lock-Free Dynamically Resizable Arrays. In Alexander A. Shvartsman, editor, *OPODIS*, volume 4305 of *Lecture Notes in Computer Science*, pages 142–156. Springer, 2006.
- [3] David L. Detlefs, Paul A. Martin, Mark Moir, and Guy L. Steele Jr. Lock-free reference counting. *Distrib. Comput.*, 15(4):255–271, 2002.
- [4] Keir Fraser. Practical lock-freedom. Technical Report UCAM-CL-TR-579, University of Cambridge, Computer Laboratory, February 2004.
- [5] Keir Fraser and Tim Harris. Concurrent programming without locks. *ACM Trans. Comput. Syst.*, 25(2):5, 2007.
- [6] David Gifford and Alfred Spector. Case study: IBM's system/360-370 architecture. *Commun. ACM*, 30(4):291–307, 1987.
- [7] Peter Gluck and Gerard Holzmann. Using SPIN Model Checker for Flight Software Verification. In *Proceedings of the 2002 IEEE Aerospace Conference*, 2002.
- [8] Timothy L. Harris, Keir Fraser, and Ian A Pratt. A practical multi-word compare-and-swap operation. In *Proceedings of the 16th International Symposium on Distributed Computing*, 2002.
- [9] Danny Hendler, Nir Shavit, and Lena Yerushalmi. A scalable lock-free stack algorithm. In *SPAA 2004*, pages 206–215, New York, NY, USA, 2004. ACM Press.
- [10] Maurice Herlihy, Victor Luchangco, Paul Martin, and Mark Moir. Nonblocking memory management support for dynamic-sized data structures. *ACM Trans. Comput. Syst.*, 23(2):146–196, 2005.
- [11] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming*. Morgan Kaufmann, March 2008.
- [12] Intel. IA-32 Intel Architecture Software Developer's Manual, Volume 3: System Programming Guide, 2004.
- [13] Michael R. Lowry. Software Construction and Analysis Tools for Future Space Missions. In Joost-Pieter Katoen and Perdita Stevens, editors, *TACAS*, volume 2280 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2002.
- [14] Maged M. Michael. Hazard Pointers: Safe Memory Reclamation for Lock-Free Objects. *IEEE Trans. Parallel Distrib. Syst.*, 15(6):491–504, 2004.
- [15] Kirk Reinholtz. Atomic Reference Counting Pointers, C++ User Journal. December 2008.
- [16] Michael F. Spear, Arrvindh Shriraman, Hemayet Hossain, Sandhya Dwarkadas, and Michael L. Scott. Alert-on-update: a communication aid for shared memory multiprocessors. In *PPoPP '07: Proceedings of the 12th ACM SIGPLAN symposium on Principles and practice of parallel programming*, pages 132–133, New York, NY, USA, 2007. ACM.
- [17] Bjarne Stroustrup. *The C++ Programming Language*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2000.