

Doc. No.: N3712
Date: 2013-08-30
Author: Arch D. Robison <arch.robison@intel.com>

Policy-Based Design for Safe Destruction in Concurrent Containers

Contents

1	Introduction	1
1.1	Weak Guarantee	2
1.2	Marked Pointers	3
2	Interface	3
2.1	Possible Concrete Implementations	5
3	Alternatives	5
4	Concrete Interface	6
4.1	enable_concurrent_ptr	6
4.2	concurrent_ptr	6
4.3	marked_ptr	8
4.4	guard_ptr	9
5	Examples	10
5.1	Classic Lock-Free Stack	10
5.2	Lock-Free List-Based Set	12
6	Acknowledgments	15

1 Introduction

Writing containers that support concurrently access/modification by multiple threads incurs a special problem: when is it safe to delete an object? The problem is surprisingly tricky when lock-free operations are involved. Numerous general means have been proposed, such as Hazard Pointers [1], Pass the

Buck [2], Epoch Based Reclamation [3], Read Copy Update (RCU) [4], Differential Reference Counting (DRC) [5], Reference Counting [6], Garbage Collection (GC), and Transactional Memory.¹ Each means has its own strengths and weaknesses, and no one approach is ideal in all contexts.

This paper sketches a way of abstracting these means so that concurrent container implementations could be parameterized over these means, thus allowing users to select the means best for their context. Though few users will want to change a container at whim, say from Hazard Pointers to RCU, different applications will have different requirements. For example, though RCU is popular for kernel writing, Hazard Pointers may be preferable for user-space code. Hence abstracting over the particular means could enable wider use of common concurrent container code.

What would be standardized is a conceptual *reclaimer* interface along with a concrete implementation. Concurrent containers (see N3732) would be parameterized over a *reclaimer* argument. This arrangement would be analogous to current practice with allocators: they have a conceptual *allocator* interface, a concrete implementation `std::allocator`, and are used as parameters to container templates.

This paper presents a sketch, not a complete proposal. It is intended to elicit discussion from experts as to whether the abstraction is viable. As such, this paper presents a fairly minimal interface that omits knobs for weak memory consistency and extensive overloading of operators, so that experts can focus on a core question: Is the abstraction sufficient to support various means for deferred object deletion?

1.1 Weak Guarantee

The basic idea is to define some abstractions that support the weakest common guarantee that the various means can support. There are two fundamental abstractions:

- A *concurrent_ptr* acts like a pointer that supports atomic operations.
- A *guard_ptr* is an object that can atomically take a snapshot of the value of a *concurrent_ptr* and **if** the target has not yet been deleted, guarantee that the target will not be deleted while the *guard_ptr* holds a pointer to it.

¹ The recent “Drop the Anchor” [7] approach is omitted from the list since it seems difficult to apply generically.

It is important to note that this guarantee is much weaker than classic garbage collection: It is only *guard_ptr* references that protect against deletion.

In effect, a *concurrent_ptr* is a “weak” pointer and a *guard_ptr* is a “shared ownership” pointer. The key semantic differences from `std::weak_ptr` and `std::shared_ptr` are that

- *concurrent_ptr* and *guard_ptr* are abstract interfaces (a.k.a. “concepts”), not concrete interfaces.
- They support a snapshot operation that is an *atomic* analog of method `std::weak_ptr::lock()`.
- A `std::weak_ptr` can indicate whether it has “expired” — that is its target was deleted. A *concurrent_ptr* gives no such indication even if, as it can in some implementations, point to freed memory.

1.2 Marked Pointers

Many lock-free algorithms rely on marking pointers with one or more low-order bits borrowed from the pointer value itself. See references [8, 1, 9] for examples. Because of the way the bit borrowing interacts with some of the deferred deletion mechanisms, the notion of a marked pointer is part of the interface.

The abstractions allow more than one mark bit per pointer, as opposed to Java’s `AtomicMarkableReference` that supports only a single bit. If the number of mark bits exceeds practicality on the target machine, the implementation can fall back on using locks, much the way existing `std::atomic<T>` can. Like the latter, a marked pointer has a method `is_lock_free()` that indicates whether operations on it are lock free.

2 Interface

Each concurrent container would take a *reclaimer* argument in addition to an allocator argument. Similar to the allocator argument, there would be a default reclaimer, for example `std::default_reclaimer`.

A reclaimer type `R` would define two abstractions necessary for safe destruction and deletion.

- `R::concurrent_ptr<T,D>` : Acts like an atomic markable pointer to objects of type `T` with deleter of type `D`. It supports atomic operations such as `load`, `store`, and `compare_exchange_weak`. Class `T` must be derived from `enable_concurrent_ptr<T,N>`.
- `R::enable_concurrent_ptr<T,N>` : Mandatory base class for targets of `concurrent_ptr`. `T` is the derived class. `N` is the number of mark bits supported, which defaults to zero.

The intent of `enable_concurrent_ptr<T,N>` is to provide implementers of reclaimers with two things:

- A way to force the alignment of targets, which is a common way to provide mark bits in the pointers.
- A place to embed reclaimer state, such as reference counts, in the user's objects.

Class `concurrent_ptr<T,D>` provides two auxiliary types for working with `concurrent_ptr` objects.

- `concurrent_ptr<T,D>::marked_ptr` : Acts like a pointer, but has `N` mark bits, where `N` is specified by the base class `enable_concurrent_ptr<T,N>` of `T`.
- `concurrent_ptr<T,D>::guard_ptr` : Similar to a `marked_ptr`, but has shared ownership of its target *if* the target has not been deleted.

The key operation is `acquire`, which takes a snapshot with the weak guarantee. Here is an example of the syntax:

```
extern concurrent_ptr p;
extern guard_ptr g;
g.acquire(p);
```

In wait-free algorithms, `acquire` may be problematic when implemented with hazard pointers or pass the buck, because it may have to loop indefinitely to get a good snapshot. Thus a fancier form is also provided that enables quitting early if the value in `p` does not match a provided value `m`.

```
extern concurrent_ptr p;
extern guard_ptr g;
extern marked_ptr m;
bool b = g.acquire_if_equal(p,m);
```

If the pointer in `p` does not match `m`, `b==false`. Otherwise `b==true` and the net effect is the same as for a plain `acquire`, and afterwards `g.get()==m`. Listing 9 in Section 5.2 will show it in use.

Releasing a guard follows the standard smart pointer interface. The operation `g.reset` releases ownership and sets `g` to `nullptr`. The destructor of `guard_ptr` implicitly calls `reset`.

2.1 Possible Concrete Implementations

Some possible concrete implementations:

- Hazard Pointers: `guard_ptr` is a hazard pointer. The operation `acquire(p)` repeatedly copies `p` into a hazard pointer until the hazard pointer and `p` match. The operation `g.reset()` sets the hazard pointer to null.
- Pass the Buck: `guard_ptr` is a “guard”. Operation of `acquire(p)` is similar to hazard pointers.
- RCU: The operation `acquire` indicates start of a read critical section followed by a `rcu_dereference` operation. The operation `g.reset()` indicates the end of a read critical section.
- DRC: A `concurrent_ptr` is a “strong pointer” and a `guard_ptr` is a “basic pointer”.
- GC: `guard_ptr` is just a traceable pointer.

Note that RCU, DRC, and GC provide much stronger guarantees than our “weak guarantee”. Programmers should take care to not accidentally rely on the strong guarantees.

3 Alternatives

In theory, `shared_ptr` and `weak_ptr` could be used, but their higher degree of safety adds significant overhead. In particular, atomic operations on these are quite expensive, requiring expensive locking protocols and solve the problem poorly. (Michael and Scott [6] discuss why reference counting is problematic.) In contrast, this proposal enables lightweight mechanisms such as Hazard Pointers or RCU to be employed.

The sketched proposal separates *allocator* and *reclaimer* policy arguments. An alternative is for concurrent containers to take a single policy argument that combines both into a “concurrent allocator”. The combination would eliminate the need for the type `enable_concurrent_ptr` since the concurrent allocator could do the equivalent alignment adjustments when allocating an object.

4 Concrete Interface

This section shows a basic concrete interface, sufficient for the examples in Section 5. Class templates `enable_concurrent_ptr` and `concurrent_ptr` are assumed to be members of a concrete reclaimer class, for example `rcu_reclaimer` or `hazard_pointer_reclaimer`.

4.1 `enable_concurrent_ptr`

```

1  template<class T, size_t MarkSize=0>
2  class enable_concurrent_ptr {
3  protected:
4      enable_concurrent_ptr() noexcept;
5      enable_concurrent_ptr(const enable_concurrent_ptr&) noexcept;
6      enable_concurrent_ptr& operator=(const enable_concurrent_ptr&)
          noexcept;
7      ~enable_concurrent_ptr() noexcept;
8  };

```

Listing 1: Concrete interface for `enable_concurrent_ptr`

4.2 `concurrent_ptr`

```

1  //! T must be derived from enable_concurrent_ptr<T>. D is a deleter.
2  template<class T, class D=std::default_delete<T> >
3  class concurrent_ptr {
4  public:
5
6      // See Section 4.3
7      typedef implementation-defined marked_ptr;
8
9      // See Section 4.4

```

```

10     typedef implementation-defined guard_ptr;
11
12     // Construct concurrent_ptr
13     concurrent_ptr(const marked_ptr& p=marked_ptr());
14
15     // Atomic load that does not guard target from being reclaimed.
16     marked_ptr load() const;
17
18     // Atomic store.
19     void store(const marked_ptr& src);
20
21     // Shorthand for store(src.get())
22     void store(const guard_ptr& src);
23
24     // Atomic exchange.
25     void swap(guard_ptr& q);
26
27     // Compare-and-swap.
28     bool compare_exchange_weak(marked_ptr& expected, marked_ptr
        desired);
29
30     // Compare-and-swap.
31     bool compare_exchange_weak(guard_ptr& expected, marked_ptr
        desired);
32
33     // Shorthand for compare_exchange_weak(src.get())
34     bool compare_exchange_weak(marked_ptr& expected, guard_ptr&
        desired);
35
36     // Shorthand for compare_exchange_weak(src.get())
37     bool compare_exchange_weak(guard_ptr& expected, guard_ptr&
        desired);
38 };

```

Listing 2: Concrete interface for concurrent_ptr

The interface for atomic exchange departs somewhat from the rest of the standard library to avoid introducing the need for a third `guard_ptr` object, whose construction might incur unnecessary overhead.

The several flavors of `compare_exchange_weak` exist to support various cases where one or both arguments do not need guard protection.

4.3 `marked_ptr`

The implementation-defined type `concurrent_ptr<T,D>::marked_ptr` should behave as if it were defined as follows.

```

1  template<class T, class D>
2  class concurrent_ptr<T,D>::marked_ptr {
3  public:
4      // Construct a marked_ptr
5      marked_ptr(T* p=nullptr, uintptr_t mark=0) noexcept;
6
7      // Copy constructor
8      marked_ptr(const marked_ptr&) noexcept;
9
10     // Copy underlying marked pointer of a guard_ptr.
11     marked_ptr(const guard_ptr& g) noexcept;
12
13     // Destructor
14     ~marked_ptr() noexcept;
15
16     // Assignment
17     marked_ptr& operator=(const marked_ptr&) noexcept;
18
19     // Set to nullptr
20     void reset() noexcept;
21
22     // Get mark bits
23     uintptr_t mark() const noexcept;
24
25     // Get underlying pointer (with mark bits stripped off).
26     T* get() const noexcept;
27
28     // True iff get()==nullptr && mark()==0
29     explicit operator bool() const noexcept;
30
31     // Get pointer with mark bits stripped off.
32     T* operator->() const noexcept;
33
34     // Get reference to target of pointer.

```



```

35     T& operator*() const noexcept;
36
37     // True iff operations are lock-free
38     static bool is_lock_free() noexcept;
39 };

```

Listing 3: Concrete interface for `marked_ptr`

Additionally, there is a requirement that `operator==` and `operator!=` operate on marked pointers, and include mark bits in their comparisons.

4.4 `guard_ptr`

The implementation-defined type `concurrent_ptr<T,D>::guard_ptr` should behave as if it were defined as follows.

```

1  template<class T, class D>
2  class concurrent_ptr<T,D>::guard_ptr {
3  public:
4      // Guard a marked_ptr.
5      guard_ptr(const marked_ptr& p=marked_ptr());
6
7      // Copy constructor
8      explicit guard_ptr(const guard_ptr& p);
9
10     // Copy assignment
11     guard_ptr& operator=(const guard_ptr& p);
12
13     // Destructor
14     ~guard_ptr();
15
16     // Atomically take snapshot of p, and *if* it points to unreclaimed
17         object, acquire shared ownership of it.
18     void acquire(concurrent_ptr& p);
19
20     // Like acquire, but quit early if a snapshot != expected.
21     bool acquire_if_equal(concurrent_ptr& p, const marked_ptr&
22         expected);
23
24     // Release ownership. Postcondition: get()==nullptr.
25     void reset() noexcept;

```

```

24
25     // Get underlying pointer
26     T* get() const noexcept;
27
28     // Get mark bits
29     uintptr_t mark() const noexcept;
30
31     // True if snapshot is nullptr with mark bits set to zero.
32     explicit operator bool() const noexcept;
33
34     // Get pointer with mark bits stripped off. Undefined if target has
35     been reclaimed.
36     T* operator->() const noexcept;
37
38     // Get reference to target of pointer. Undefined if target has been
39     reclaimed.
40     T& operator*() const noexcept;
41
42     // Swap two guards
43     void swap(guard_ptr& g) noexcept;
44
45     // Reset. Deleter d will be applied some time after all owners release
46     their ownership.
47     void reclaim(D d=D()) noexcept;
48
49     // True iff operations are lock-free
50     static bool is_lock_free() noexcept;
51 };

```

Listing 4: Concrete interface for guard_ptr

5 Examples

5.1 Classic Lock-Free Stack

The classic lock-free stack is a linked list. Listing 5 declares the data structures.

```

1 struct node;

```

```

2 typedef typename Reclaimer::template concurrent_ptr<node>
   concurrent_ptr;
3
4 // A node in a linked list
5 struct node: Reclaimer::template enable_concurrent_ptr<node> {
6     const Value value;
7     node* next;
8     node(Value k) : value(k) {}
9 };
10
11 // Root of the linked list
12 concurrent_ptr root;
13
14 typedef typename concurrent_ptr::marked_ptr marked_ptr;
15 typedef typename concurrent_ptr::guard_ptr guard_ptr;

```

Listing 5: Data declarations for a lock-free stack

A subtle point is that `next` is declared as a plain pointer, not a `concurrent_ptr` or a `atomic<node*>`, though it could be. The reason is that its value is written only before it is published to other threads or after reclamation has determined it is accessible to only one thread.

Listing 6 shows how an item is pushed onto the list. Each `mark_ptr` here has zero mark bits. No `guard_ptr` is required since no other thread can possibly delete the newly created node until it is published in the list.

```

1 void push(Value value) {
2     marked_ptr n = new node(value);
3     marked_ptr m;
4     do {
5         m = root.load();
6         n->next = m.get();
7     } while(!root.compare_exchange_weak(m,n));
8 }

```

Listing 6: Pushing onto a lock-free stack

Listing 7 shows how to pop an item from the stack. The `guard_ptr` is essential for preventing an ABA problem that could corrupt the list representation.

```

1 bool try_pop(Value& v) {
2     guard_ptr g;

```

```

3      do {
4          // Get guarded snapshot of root.
5          g.acquire(root);
6          if(!g)
7              // Stack is empty.
8              return false;
9      } while(!root.compare_exchange_weak(g, g->next));
10     // Successfully popped item from stack.
11     v = g->value;
12     g.reclaim();
13     return true;
14 }

```

Listing 7: Popping from a lock-free stack

5.2 Lock-Free List-Based Set

This example is adapted from Figure 9 of reference[1]. Listing 8 declares the data structures.

```

1      struct node;
2      typedef typename Reclaimer::template concurrent_ptr<node>
           concurrent_ptr;
3      typedef Key key;
4      struct node: Reclaimer::template enable_concurrent_ptr<node,1> {
5          const Key key;
6          concurrent_ptr next;
7          node(Key k) : key(k) {}
8      };
9      typedef typename concurrent_ptr::guard_ptr guard_ptr;
10     typedef typename concurrent_ptr::marked_ptr marked_ptr;
11     concurrent_ptr head;

```

Listing 8: Data declarations for a lock-Free list-based set

Listing 9 shows routine `find`, which is the workhorse for the lock-free list. It searches the list for a given key or suitable insertion point for the key.

```

1 bool find(Key key, concurrent_ptr*& prev, guard_ptr& cur,
           marked_ptr& next, guard_ptr& save) {
2     retry:
3         prev = &head;

```

```

4     next = prev->load();
5     for(;;) {
6         if(!cur.acquire_if_equal(*prev, next))
7             goto retry;
8         if(!cur)
9             return false;
10        next = cur->next.load();
11        if(next.mark() != 0) {
12            // Node *cur is marked for deletion.
13            next = marked_ptr(next.get(), 0);
14            // Try to splice out node
15            if(!prev->compare_exchange_weak(cur, next))
16                goto retry;
17            cur.reclaim();
18        } else {
19            Key ckey = cur->key;
20            if(prev->load() != marked_ptr(cur))
21                // *cur might be cut from list.
22                goto retry;
23            if(ckey >= key)
24                return ckey == key;
25            prev = &cur->next;
26            cur.swap(save);
27        }
28    }
29 }

```

Listing 9: Find operation on a lock-free List

If it returns true, then the following properties hold:

- `prev` points to a live (not in reclaimed memory) `concurrent_ptr`.
- `cur` points to a live node with a key greater or equal to the search key.
- `save` is guarding the node containing `*prev`, unless `prev == &head`, in which case no guard is necessary.
- While `cur->next` and `next` match, the node pointed to by `cur` is still part of the list.
- While `prev->load()` and `cur` match, it is safe to insert a new node there. See Listing 10 for details.

Listing 10 uses find to insert an item.

```

1 bool insert(Key key) {
2     node* n = new node(key);
3     concurrent_ptr* prev;
4     guard_ptr cur, save;
5     marked_ptr next;
6     do {
7         if(find(key,prev,cur,next,save)) {
8             delete n;
9             return false;
10        }
11        // Try to install new node
12        n->next.store(cur.get());
13    } while(!prev->compare_exchange_weak(cur,n));
14    return true;
15 }
```

Listing 10: Insert operation on a lock-Free list

Listing 11 uses find to erase an item.

```

1 bool erase(Key key) {
2     concurrent_ptr* prev;
3     guard_ptr cur, save;
4     marked_ptr next;
5     // Find node in list with matching key and mark it for erasure.
6     do {
7         if(!find(key,prev,cur,next,save))
8             // No such node in the list
9             return false;
10    } while(!cur->next.compare_exchange_weak(next,marked_ptr(next.
11        get(),1)));
12    // Try to splice out node
13    if(prev->compare_exchange_weak(cur,next))
14        cur.reclaim();
15    else
16        // Another thread interfered. Rewalk list to ensure reclamation of
17        marked node before returning.
18        find(key,prev,cur,next,save);
19    return true;
20 }
```

18 }

Listing 11: Erase operation on a lock-Free list

The `find` on line 16 assures progress guarantees [8].

6 Acknowledgments

Paul McKenney raised the issue of flexibly supporting different mechanisms for deferred deletion, and provided feedback on an earlier draft. Maged Michael and Paul raised issues of lock-free versus wait-free, which motivated extending `acquire` to `acquire_if_equal`. Discussions with them also reminded me of the ubiquity of mark bits in lock-free algorithms. Dmitry Vyukov’s “strong pointer” and “basic pointer” interface inspired the basic interface. Artur Laksberg critiqued an earlier draft, and Anton Potapov corrected a later draft.

References

- [1] Maged M. Michael. Hazard pointers: Safe memory reclamation for lock-free objects. *IEEE Trans. Parallel Distrib. Syst.*, 15(6):491–504, June 2004.
- [2] Maurice Herlihy, Victor Luchangco, and Mark Moir. The repeat offender problem: A mechanism for supporting dynamic-sized, lock-free data structures. In *Proceedings of the 16th International Conference on Distributed Computing*, DISC ’02, pages 339–353, London, UK, UK, 2002. Springer-Verlag.
- [3] Keir Fraser. Practical lock-freedom. Technical Report UCAM-CL-TR-579, University of Cambridge, Computer Laboratory, February 2004.
- [4] Mathieu Desnoyers, Paul E. McKenney, Alan S. Stern, Michel R. Dagenais, and Jonathan Walpole. User-level implementations of read-copy update. *IEEE Trans. Parallel Distrib. Syst.*, 23(2):375–382, February 2012.
- [5] Differential reference counting. <http://www.1024cores.net/home/lock-free-algorithms/object-life-time-management/differential-reference-counting>. Accessed 2013-08-07.

- [6] Maged M. Michael and Michael L. Scott. Correction of a memory management method for lock-free data structures. Technical report, Rochester, NY, USA, 1995.
- [7] Anastasia Braginsky, Alex Kogan, and Erez Petrank. Drop the anchor: lightweight memory management for non-blocking data structures. In *Proceedings of the 25th ACM symposium on Parallelism in algorithms and architectures*, SPAA '13, pages 33–42, New York, NY, USA, 2013. ACM.
- [8] Timothy L. Harris. A pragmatic implementation of non-blocking linked-lists. In *Proceedings of the 15th International Conference on Distributed Computing*, DISC '01, pages 300–314, London, UK, UK, 2001. Springer-Verlag.
- [9] Hans-J. Boehm. An almost non-blocking stack. In *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, PODC '04, pages 40–49, New York, NY, USA, 2004. ACM.