

# P2055R0: A Relaxed Guide to `memory_order_relaxed`

Hans Boehm                      Paul E. McKenney  
hboehm@google.com              paulmck@kernel.org

The Indefatigible TBD

January 12, 2020 (Pre-Prague)

Audience: SG1 (Informational)

## Abstract

The out-of-thin-air (OOA) and read-from-untaken-branch (RFUB) properties of the specification of `memory_order_relaxed` have resulted in considerable consternation over the years. Although there are no known instances of full-blown OOTA behavior, and no known RFUB-induced failures of production code, the theoretical possibility of these properties severely complicates automated analysis of large C and C++ code bases. Thus far, attempts to eliminate OOTA and RFUB properties from the memory model have resulted in otherwise needless added overheads on weakly ordered systems on the one hand or excessive implementation complexity on the other. However, `memory_order_relaxed` never was intended to be used in arbitrary code, but rather as a part of deliberate application of specific concurrency designs. This paper forms an initial catalog of patterns underlying such designs.

## 1 OOTA and RFUB Background

There has been considerable work done on OOTA and RFUB [5, 14, 11, 20, 4, 3]. This work has taken place over many years, and builds on prior work in the Java community.<sup>1</sup>

There has long been hope that additional research effort will identify a model of OOTA that all can live with, for example, on the part of Paul, and that everyone would come to appreciate the relative simplicity of strengthening `memory_order_relaxed` to forbid prior reads to be reordered with later writes, for example, on the part of Hans. In addition, although there is general agreement that OOTA behaviors must be forbidden, there is some debate on the need to forbid RFUB behaviors. Perhaps agreement on these points will be reached, but in the meantime, `memory_order_relaxed` use is increasing, and thus an increasing need to identify known-safe usage

---

<sup>1</sup> See for example the infamous “Causality Test Cases” (<http://www.cs.umd.edu/~pugh/java/memoryModel/unifiedProposal/testcases.html>).

---

```
1 int x = 0;
2 int y = 0;
3
4 void thread1()
5 {
6     if (x)
7         y = 1;
8 }
9
10 void thread2()
11 {
12     if (y)
13         x = 1;
14 }
```

Listing 1: Non-Atomic Accesses Sometimes Respect Control Dependencies

patterns. In the best case, these usage patterns might be automatically recognized in existing code, but at a minimum we hope that this list will be useful to code reviewers.

This paper is a first step toward such a set of patterns.

The term *full C++* refers to the C++20 memory model as stated in the current draft. The term *strict C++* refers to the subset of full C++ obtained by dropping the following normative encouragement from the C++20 memory model: “Implementations should ensure that no “out-of-thin-air” values are computed that circularly depend on their own computation.” Some (but not all) of the proto-patterns in this document are safe in strict C++, but all of them are safe in full C++.

## 2 Relaxed Design Patterns

Many of these patterns are taken from Hans’s `memory-model-design` posting on September 4, 2018.<sup>2</sup>

### 2.1 Non-Racing Accesses

Any non-racing access to an atomic object can be a relaxed access. Because the access is not concurrent with a conflicting access (store against either store or load), further ordering is unnecessary.<sup>3</sup> In fact, such accesses can in theory be non-atomic. In environments where atomicity is controlled by the access rather than the object definition, such accesses are often non-atomic in practice [1].

For example, given concurrent execution of `thread1()` and `thread2()` in Listing 1, the only permitted outcome results in both `x` and `y` being equal to zero in both

---

<sup>2</sup> Message-ID: <CAMOCf+jchGw6DeE2NyCJA3wfFbNH-WFn59JruZPSWt9\_jPW9NQ@mail.gmail.com>.

<sup>3</sup> This covers case #8 in Hans’s September 4, 2019 email.

## 2.1 Non-Racing Accesses

---

```
1 std::atomic<int> x = 0;
2 std::atomic<int> y = 0;
3
4 void thread1()
5 {
6     if (x.load(memory_order_relaxed))
7         y.store(1, memory_order_relaxed);
8 }
9
10 void thread2()
11 {
12     if (y.load(memory_order_relaxed))
13         x.store(1, memory_order_relaxed);
14 }
```

Listing 2: Strict C++ Does Not Require Atomics to Respect Control Dependencies

full C++ and strict C++. Any other outcome would violate the “sequential consistency for data race free programs” principle, and must effectively be due to a compiler-created data race, which is forbidden.

In contrast, in the analogous program using C++ atomics (see Listing 2), additional behaviors are permitted by strict C++, including the one resulting in the final values of both `x` and `y` being 1. The restriction to “strict C++” is important because this code fragment is considered to be an example of the OOTA behavior that is forbidden by the normative encouragement in that same standard.

In short, although any non-racing access to an atomic object may be relaxed, strict C++ counter-intuitively classifies many access patterns as racy.

### 2.1.1 Initialization and Cleanup

Important special cases of this pattern are the single-threaded initialization and cleanup phases of an otherwise concurrent program. These use cases are one motivation for the strong ordering guarantees of thread creation and destruction. These guarantees permit the single-threaded initialization and cleanup code to run race free, with no need to consider interference from the intervening code that runs multithreaded.

### 2.1.2 Lock-Based Critical Sections

Exclusive locks provide mutual exclusion, so that objects accessed only while holding a given lock may be accessed using `memory_order_relaxed` accesses, or, for that matter, using non-atomic accesses.

Reader-writer locks provide a weaker form of mutual exclusion. However, objects that are updated only while a given reader-writer lock is write-held and read only when that same lock is either read-held or write-held may also be accessed using `memory_order_relaxed` accesses, or, again, using non-atomic accesses.

Of course, non-atomic accesses are almost always used with pure locking. However, `memory_order_relaxed` accesses are sometimes quite useful, for example, in cases where objects pass through a software pipeline, where one stage uses pure locking and another stage relies on atomic operations.

### 2.2 Single-Location Data Structures

Relaxed atomic operations provide sequentially consistent access to a single object. This means that data structures that fit into a single object can be accessed with relaxed atomics with no possibility of OOTA or RFUB behavior.

Note well that a group of single-location data structures might well interact in a way that could raise the spectre of OOTA or RFUB. As before, design review should therefore pay careful attention to information flow.

### 2.3 Shared Fences

The `atomic_thread_fence()` function can be used to order multiple sets of accesses, for example, by replacing a series of acquire loads with relaxed loads followed by an `atomic_thread_fence(memory_order_acquire)` [19, Section 4.1] or a series of release stores with an `atomic_thread_fence(memory_order_release)` followed by a series of relaxed stores [19, Section 4.2]. In many cases, other ordered atomic operations may be substituted for the fence operations.

In this design pattern, OOTA and RFUB behaviors are ruled out by the semantics of `atomic_thread_fence()`.

### 2.4 Atomic Reference-Count Updates

In certain reference-count use cases, the ordering of the increments and decrements is irrelevant. One common case is where it is only legal to increment the reference count when the incrementing thread already holds a reference, in which case the count cannot possibly decrease to zero in the meantime. Because only the one-to-zero transition requires ordering, reference-count increments can be relaxed in cases where another reference is guaranteed to be held throughout.

Similarly, reference-count decrements can also be relaxed, but only if the thread will still hold at least one reference after the decrement. In other words, a thread releasing its last reference is forbidden from using a relaxed operation to do so, because in that case there is no guarantee that another reference is guaranteed to be held throughout.<sup>4</sup>

We suspect that this is an example of a more general class of patterns, but other examples of such a class do not immediately come to mind. One can of course imagine things like preprocessed sensor values where these values are irrelevant except in their relation to cutoff values. We would welcome examples used in actual code.

---

<sup>4</sup> More elaborate variants of this pattern allow these rules to be relaxed. For example, if a parent thread is guaranteed not to release its last reference until after joining with its child threads, then those child threads may use relaxed decrements to release their final reference.

### 2.5 Untrusted Loads

In many cases, it is acceptable for a load from an atomic shared variable to occasionally return random bits because the value is checked by some later operation. In such cases, the load can be a relaxed load.

#### 2.5.1 Pre-Load for Compare and Swap

Perhaps the most well-known later checking operation is a non-relaxed compare-and-swap (CAS). The `atomic_compare_exchange_*()` family of read-modify-write CAS operations are typically used in a loop, and often require an initial load prior to the first pass through that loop. For non-relaxed CAS operations, this initial load can typically be a relaxed load, with the CAS operation's ordering preventing OOTA and RFUB behaviors. Relaxed CAS operations need to be part of some other design pattern (for example, the shared fences pattern called out in Section 2.3) if cycles containing them are to be guaranteed to be OOTA/RFUB-free in conjunction with an initial relaxed load. One common design pattern is the single-location data structure discussed in Section 2.2.

Additional examples are presented by Sinclair et al. [20].

#### 2.5.2 Sequence Locking

The accesses within a sequence-locking read-side critical section can use relaxed loads because any concurrency with the corresponding update will result in a retry, thus discarding any loaded values. Assuming that sequence-locking readers never store to shared memory, this not only prevents the surfacing of any OOTA or RFUB cycles, but also of any other non-SC behaviors.

Note that a proposal<sup>5</sup> provides an `atomic_load_per_byte_memcpy()` that allows safe non-atomic access to data for sequence-lock readers, as well as an `atomic_store_per_byte_memcpy()` to update that same data by sequence-lock updaters. It is nevertheless quite possible that some sequence-lock readers might continue to use relaxed atomics in order to permit reliable computations within readers in the presence of data objects having trap representations.

Furthermore, sophisticated sequence-locking use cases may need to use relaxed accesses for other reasons. For example, the Linux kernel's lockless `path-to-inode` traversal uses the closely related sequence counters to detect large-scale changes to the filesystem tree that would otherwise confuse this traversal [7, 8]. Such confusion could result in a number of anomalies, including successful lookup of paths that never actually existed.

### 2.6 Unidirectional Data Flow

If data flows only in one direction, then OOTA cycles cannot form. The following sections give several examples of this general design pattern.

---

<sup>5</sup> <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2019/p1478r0.html>

### 2.6.1 Independent Input Data

Input data consisting of independent objects may be read using relaxed accesses because these objects are not affected by downstream computations. Here input data is defined broadly, including:

1. Measurements of outside environmental conditions.
2. Device configuration data.
3. Software configuration data.
4. Security policies.
5. Network routing information.

The key point is that the concurrent-computation portion of application references but does not modify this data, and that there are no object-to-object consistency constraints.

### 2.6.2 Independent Output Data

Similarly, output data consisting of independent objects may be written using relaxed accesses because these objects do not affect upstream computations. As before, output data is defined broadly, including:

1. Control of objects external to the computer.
2. Many classes of debug output.
3. Some use cases involving video frame buffers.
4. Some use cases involving communication to a later stage of a software pipeline.

Similar to the independent input data discussed in the preceding section, the key point is that the concurrent-computation portion of application modifies but does not reference this data, and that there are no object-to-object consistency constraints.

### 2.6.3 Statistical Counters

The canonical instance of a unidirectional data-flow pattern is the statistical counter, in which each thread (or CPU, as the case may be) updates its own counter, and the aggregate value of the counter is read out by summing all threads' counters [13, Section 5.2].

Statistical counters do have concurrent updates and reads, and thus must use atomics. However, the concurrent reads can be modeled as returning approximate results (for example, for monitoring or debugging), and can in fact be modeled as sequentially consistent approximate operations. But more to the point, data flow in real use cases is always unidirectional, proceeding from the updater responding to an event and flowing

```
1 StatCounter<int> a;  
2 StatCounter<int> b;  
3  
4 void thread1()  
5 {  
6     int r1 = a.readout();  
7     b.increase(r1);  
8 }  
9  
10 void thread2()  
11 {  
12     int r2 = b.readout();  
13     a.increase(r2);  
14 }
```

Listing 3: Statistical-Counter Abuse and OOTA

through the counter to some reader displaying or logging statistics. This unidirectional data flow precludes the cycles required for OOTA or RFUB behavior to manifest.

An example abuse is shown in Listing 3. Lines 1 and 2 define a pair of statistical counters `a` and `b`. The `thread1()` and `thread2()` functions form a classic data-dependent OOTA cycle. Assuming both statistical counters start out with all counters zero, we could in theory see the following OOTA sequence of events:

1. Line 6 sums `a`'s counters, obtaining the sum 42.
2. Line 7 increases the current component of `b`'s counter by 42.
3. Line 12 sums `b`'s counters, obtaining the sum 42 due to the increase from line 7.
4. Line 13 increases the current component of `a`'s counter by 42, thus justifying the sum of 42 obtained by line 6.

Of course, the code in Listing 3 is complete nonsense: Counters should count events, not each others's cumulative values. The code as written is about as useful as the proverbial screen door in a submarine. Problems of this sort should be located in a code review, or better yet during the preceding design review.<sup>6</sup>

Exact values are sometimes obtained from statistical counters in stop-the-world situations, such as checking for consistent results at the end of a stress test or benchmarking run [13, Sections 5.3 and 5.4]. Alternatively, counter updates might be carried out while read-holding a given reader-writer lock and counter reads while write-holding that same lock. In all of these cases, OOTA and RFUB behaviors are additionally avoided due to the fully synchronized nature of the readout.

---

<sup>6</sup> Yes, this could be considered analogous to a difference-equation control system. But in that case, the system being controlled is part of the loop, and proper synchronization must be used when communicating with that system. In addition, the actual difference-equation computation will normally be single-threaded. More importantly, if the system being controlled might pose a threat to life and limb, the design review had jolly well better be sufficiently well-informed and thorough as to avoid this sort of problem.

#### 2.6.4 Software Pipelines

Software pipelines break computation up into stages that might proceed concurrently. If the interface between a consecutive pair of stages is simple enough, relaxed accesses might be used for the corresponding communication of data. Pipelines are not necessarily strictly linear, in fact it can be quite advantageous to have concurrent stages feeding into a single subsequent stage via a reduction step. If the output of the concurrent stages is sufficiently simple, the reduction step might be a simple relaxed atomic fetch-and-op operation to a single scalar object. An example of a sufficiently simple output is event counts emanating from concurrent stream processing feeding into later sequential logic.

Note that the independent input and output data patterns discussed in Sections 2.6.1 and 2.6.2 might be endpoints of a software pipeline.

#### 2.6.5 Owner Field for Re-Entrant Mutex

This pattern is first analyzed for full C++, and then for strict C++. Spoiler warning: There is reason to believe that this pattern works in full C++, but not in strict C++.

A re-entrant exclusive mutex must track its owner in order to avoid self-deadlock when the owner re-acquires a mutex that it already holds. This owner field is updated only while the mutex is held, and its value is used only to compare for equality to the current thread's ID. Before releasing the mutex, the owner writes a special ID to the owner field that is guaranteed not to match the ID of any thread. Other threads can access the owner concurrently with the owner's update, so the owner field must be atomic in order to avoid data races. Of course, a nesting counter must also be used in order to identify the outermost lock-release operation, however this counter is accessed only by the thread currently holding the lock. Therefore, if the lock works correctly, exclusive access will be provided to the nesting counter, as is required. Those wishing to produce a proof of correctness are encouraged to try induction.

However, the only time that the owner field can be equal to the thread ID is when that thread carried out the last update to the owner field and still holds the mutex:

1. Each thread writes only its ID or the special ID.
2. Because `memory_order_relaxed` loads are single-variable SC, and because each thread sets the owner field to the special ID before releasing the mutex, a given thread cannot see its own ID unless it still holds the mutex.
3. Because atomic accesses forbid load tearing, each load from the owner field will return either the special ID or the thread ID corresponding to some thread that recently held the mutex.
4. Therefore, when a thread is not holding the mutex, it is guaranteed not to load its own ID from the owner field.

No other thread is allowed to write to the owner field while the mutex is held, so it is impossible to form the cycles required for OOTA or RFUB behavior to manifest. Therefore, both the reads from and the writes to the owner field may use `memory_order_relaxed`.



This is a special case of unidirectional data flow, with the data flowing from the mutex holder to threads not holding the mutex. The mutual exclusion provided by the mutex prevents any OOTA or RFUB cycles from forming.

However, things might well be more difficult in strict C++.

These potential difficulties stem from the possibility of undefined behavior (UB) back-propagating through a cycle so as to justify the OOTA behavior [9]. To see the rationale for this back-propagating self-justifying UB (BPSJUB?), consider a pair of threads each concurrently attempting to acquire a mutex, but where (incorrectly and inconsistently) each see that the owner field matches their respective thread IDs. Both threads would then simultaneously execute within their respective critical sections, which could result in UB. UB can back-propagate in time, which could quite possibly result in the threads each seeing their own values in the owner field, which is what instigated the UB in the first place.<sup>7</sup>

Therefore, developers and reviewers should assume that owner fields for re-entrant mutexes require full C++ in order to operate correctly.

### 2.6.6 One-Way Memory Allocation

One-way memory allocation provides fresh memory that is never deallocated, or that is deallocated using a heavy weight one-sided mechanism, for example, a stop-the-world deallocation phase. Such an allocator might use a relaxed atomic fetch-and-add operation on a shared pointer to allocate memory from a contiguous buffer. This pointer would be initialized to reference the beginning of the buffer, and each fetch-and-add operation would add the desired allocation size (perhaps rounded up to meet alignment constraints), returning the initial value of the pointer and leaving the pointer referencing the portion of the buffer following the just-completed allocation.

The semantics of the C++ fetch-and-add operation guarantees that data flows from one runtime operation to the next, acyclicly. Therefore, OOTA cycles cannot be formed on this type of allocator's pointer manipulation alone.<sup>8</sup> However, OOTA cycles can form based on the pointer values returned from such an allocator and from dereferences of these pointers. Adventurous readers can find an early drafty draft analysis of this situation (but on a more complicated allocator) in Appendix A.

In the meantime, code reviewers should view relaxed stores of pointers to newly allocated objects with great suspicion.

### 2.6.7 Relaxed Consumption

In cases where a full-speed `memory_order_consume` is needed on a weak-memory system and where the developers are willing to live within strict coding standards [12], `memory_order_relaxed` may be used to head dependency chains. In many (but not all!) use cases, the data flow is also unidirectional, proceeding from the thread installing the new object to the threads consuming it.

Note well that this design pattern is outside of the current standard.

---

<sup>7</sup> Full disclosure: Paul wrote this paragraph, and he is not completely sold on back-propagating self-justifying UB. The critical reader should therefore review David Goldblatt's working paper [9].

<sup>8</sup> This pattern can also be considered to be a single-location data structure, as discussed in Section 2.2.

## 2.7 Java-Style Hashcode Access

Given the “Java” in the title, it is only natural to ask why this applies to C or C++. The answer is simple: It applies because this portion of Java is written in C++.<sup>9</sup>

Java-style hashcode access can be used to track values that are expensive to compute and not known until runtime on the one hand, but immutable and deterministic on the other. A natural way to handle this situation is to have each access check to see if the desired value has already been computed, and, if not, compute the value and store it for later use. Of course, it is possible that two threads might concurrently load the initial not-yet-computed value, in which case, both threads will compute the value and store it. This does waste CPU time, but this waste is often greatly outweighed by reduced synchronization cost. This reduced synchronization cost is due to the `memory_order_relaxed` loads and stores used to access the value. After all, the fact that the values are deterministic means that the two threads will be storing the same value, so strongly ordered stores provide no benefit. Furthermore, any other thread is guaranteed to see either the before-computation initial value or the exact same computed value.

Given that each value is deterministic, there cannot be a cyclic chain of interdependent values, so this pattern works in both full C++ and in strict C++.

This pattern must be used with caution in cases where the value is a pointer to allocated memory, especially if that memory is allocated and initialized at runtime. First, racing initializations will result in one-time memory leaks. Second, a store of a pointer to recently initialized memory should be a `memory_order_release` store or stronger, and the corresponding loads should be `memory_order_consume` loads or stronger. Therefore, this pattern might not be helpful when the immutable and deterministic values are linked data structures, but it is often used for scalar values.

## 2.8 Chaotic Relaxation

There are a number of iterative numerical algorithms for which unsynchronized access does not slow convergence as much as waits for barrier synchronization. These algorithms can use relaxed loads and stores to update the numerical data [2]. The idea is that the iterative convergence tests correct any small errors due to accessing data from some other iteration.

In theory, these algorithms are subject to OOTA and RFUB behaviors, however, in practice, current implementations avoid such behaviors.

## 2.9 Garbage Collection

By definition, concurrent garbage collectors read pointers in the user’s heap while the application is running. On weakly ordered machines, such accesses *must* use `memory_order_relaxed` accesses, since anything else would require all pointer accesses by the user program to be ordered, which is usually far too expensive. Since such collectors often both read and write heap pointers, it is currently difficult or impossible to strictly

---

<sup>9</sup> Documented here: <https://docs.oracle.com/javase/7/docs/api/java/lang/Object.html>.

	Multiple Threads	Concurrent WW	Concurrent RW	But Checked	But Discarded	But Fungible	Unordered Cycle	Strict C++ Unsafe
Non-Racing Accesses (Section 2.1)	Y	■	■	■	■	■	■	■
Single-Location Data Structures (Section 2.2)	Y	Y	Y	■	■	■	■	■
Shared Fences (Section 2.3)	Y	Y	Y	■	■	■	■	■
Atomic Reference-Count Updates (Section 2.4)	Y	Y	Y	■	■	Y	■	■
Untrusted Loads (Section 2.5)	Y	■	Y	y	y	y	■	■
Unidirectional Data Flow (Section 2.6)	Y	Y	Y	■	■	■	■	y
Java-Style Hashcode Access (Section 2.7)	Y	Y	Y	N	N	Y	■	■
Chaotic Relaxation (Section 2.8)	Y	Y	Y	Y	■	■	Y	?
Garbage Collection (Section 2.9)	Y	Y	Y	N	N	N	?	?

Table 1: Attributes of Categories of Relaxed Design Patterns

preclude OOTA in strict C++. However, neither OOTA nor RFUN behavior has been observed in practice.

Note that although such garbage collection for C++ is rare, such garbage collectors, e.g. for Java, are often implemented in C++.

### 3 Attributes of Relaxed Design Patterns

Table 1 shows attributes of design patterns. The attributes are as follows:

1. *Multiple Threads*: The design pattern uses multiple threads in and of itself. Note that ostensibly single-threaded patterns often interact with other patterns extending across multiple threads. For example, the allocator caches discussed in Section A operate within a single thread, but the resulting memory blocks and associated pointers might be passed to other threads using some other pattern such as release-acquire or release-consume.
2. *Concurrent WW*: The design pattern involves concurrent relaxed writes to a given object.
3. *Concurrent RW*: The design pattern involves concurrent relaxed reads and writes to a given object, but not necessarily concurrent relaxed writes.
4. *But Checked*: The values from the concurrent reads are checked if there might have been a concurrent write. See for example Section 2.5.1.
5. *But Discarded*: The values from the concurrent reads are discarded if there might have been a concurrent write. See for example Section 2.5.2.

- 
6. *But Fungible*: A set of writers are fungible if a reader running concurrently with those writers will exhibit the same behavior regardless of which of those writes' values that read returns. An important special case is when all the writers are storing the same value, as discussed in Section 2.7. Another important special case is where readers take one action if the value is (say) zero and another for any non-zero value, and all concurrent writers will write a non-zero value, as discussed in Section 2.4.

In all of these “But” columns, lower case “y” indicates that some examples in the category might possess the corresponding attribute.

7. *Unordered Cycle*: The design pattern can produce an unordered cycle in and of itself. Of course, a combination of design patterns that individually exclude the possibility of an unordered cycle might nevertheless produce an unordered cycle when used in combination. Design and code reviews should therefore carefully consider ordering at the intersection of multiple design patterns.
8. *Strict C++ Unsafe*: The design pattern is expected to work correctly in full C++, but not necessarily in strict C++. A lower case “y” indicates that only some examples in the category are agreed to be safe for strict C++, in which case the examples that are believe to be unsafe for strict C++ are called out in their respective sections.

## 4 Marking of Relaxed Design Patterns

It is currently believed that these design patterns will need to be explicitly marked in order for code reviewers and automatic verifiers to recognize them and validate their usage, although the authors would love to be proven wrong on this point. Here are some candidate marking strategies that have been discussed within the C++ standards committee:

1. Create new `memory_order` enum members for each new design pattern. This has the benefit of calling out the pattern in an unmistakable way that is visible to the compiler, but requires that each new design pattern be standardized. It also does not support the case where a given access plays a role in multiple overlapping design patterns.
2. Use structured comments to mark each design pattern. This avoids the time delays and administrative overhead inherent in standardization, and could potentially allow multiple comments to handle a given access that plays a role in multiple overlapping design patterns.
3. Use structured comments with a per-instance identifier for a given use of a pattern. The idea here is to enable tools to more easily spot unintended interactions between different design patterns being applied to a given group of objects. On the other hand, this raises the issue of namespace management.

- 
4. Define C++ `template` types for each design pattern. This is an excellent idea where it applies, as it might well for the statistical counters discussed in Section 2.6. However, we have reason to doubt that `template` types can be reasonably created for all possible relaxed-access design patterns.

More ideation and discussion is needed on this topic.

## 5 Concluding Remarks

This paper starts the work of classifying known-safe design patterns involving `memory_order_relaxed`. It is hoped that this work will be of use in design and code reviews, and that it might eventually lead to improved theoretical models of `memory_order_relaxed` accesses.

---

## A Allocator Caches

This appendix expands on the one-way memory allocator discussed in Section 2.6.6 by way of a multithreaded allocator with caches. Once agreement is reached on the simpler one-way case, this appendix might be promoted to the main paper.

Allocator caches provide per-CPU or per-thread pools of free memory in order to provide high-performance scalable memory allocation in the common case [15, 6]. Accesses to these pools are normally single-threaded by design for reasons of performance and scalability. However, objects are often allocated for concurrent algorithms, It may be helpful to list phases of a dynamically allocated object’s typical lifetime in a concurrent context:

1. Allocation.
2. Initialization, including construction.
3. Use. This might include subphases, but given that any such subphases are defined by the user, safely transitioning between them is the user’s responsibility. This is usually the only phase that permits concurrent access.
4. Cleanup, including destruction.
5. Deallocation.

Note the possibility of memory reuse means that this is a cycle rather than a sequence.

The key point is that there must be a happens-before edge for each phase transition. In the case of the cleanup to deallocation to allocation to initialization transitions, this happens-before edge is frequently supplied by sequenced-before, courtesy of the fact that allocator caches cause all of those transitions to occur within a single thread in the common case. However, some sort of happens-before edge is required for each phase transition regardless of which thread is executing any given phase.

In the common case, the transitions requiring other-thread-visible ordering are those to and from the “Use” phase. In particular, the complexities of transitioning from the “Use” phase to the “Cleanup” phase has inspired safe memory reclamation schemes, including reference counting, hazard pointers [17, 10, 18], and RCU [16].

In less-common cases where inter-thread transitions occur between other phases, proper synchronization must be provided. For example, the earlier phase might use a release store and the later phase might use an acquire or consume load.

Proper phase-transition synchronization rules out the infamous RFUB cycle shown in Listing 4.<sup>10</sup> This is because the allocation phase on line 15 is required to happen before any later phase, a requirement that is violated by the relaxed accesses on lines 5, 12, and 17.

Again, code reviewers should view relaxed stores of pointers to newly allocated objects with great suspicion.

---

<sup>10</sup> Adapted from Boehm and Demsky[5, Figure 5].

## REFERENCES

---

```
1 void *heap;
2
3 void thread1()
4 {
5     r1 = x.load(memory_order_relaxed);
6     y.store(r1, memory_order_relaxed);
7 }
8
9 void thread2()
10 {
11     bool allocated(false);
12     r1 = y.load(memory_order_relaxed);
13     if (r1 != heap) {
14         allocated = true;
15         r1 = heap;
16     }
17     x.store(r1, memory_order_relaxed);
18     assert_not(allocated);
19 }
```

Listing 4: RFUB Allocator-Like Example

## References

- [1] Jade Alglave, Will Deacon, Boqun Feng, David Howells, Daniel Lustig, Luc Maranget, Paul E. McKenney, Andrea Parri, Nicholas Piggin, Alan Stern, Akira Yokosawa, and Peter Zijlstra. Who’s afraid of a big bad optimizing compiler? Linux Weekly News, July 2019.
- [2] Gregory R. Andrews. *Concurrent Programming, Principles, and Practices*. Benjamin Cummins, 1991.
- [3] Mark Batty, Simon Cooksey, Scott Owens, Anouk Paradis, Marco Paviotti, and Daniel Wright. D1780R0: Modular relaxed dependencies: A new approach to the out-of-thin-air problem. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2019/p1780r0.html>, June 2019.
- [4] Hans-J. Boehm. P1217R2: Out-of-thin-air, revisited, again. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2019/p1217r2.html>, June 2019.
- [5] Hans-J. Boehm and Brian Demsky. Outlawing ghosts: Avoiding out-of-thin-air results. In *Proceedings of the Workshop on Memory Systems Performance and Correctness*, MSPC ’14, pages 7:1–7:6, New York, NY, USA, 2014. ACM.
- [6] Jeff Bonwick. The slab allocator: An object-caching kernel memory allocator. In *USENIX Summer Technical Conference*, pages 87–98, 1994.

## REFERENCES

---

- [7] Neil Brown. Pathname lookup in Linux. <https://lwn.net/Articles/649115/>, June 2015.
- [8] Neil Brown. RCU-walk: faster pathname lookup in Linux. <https://lwn.net/Articles/649729/>, July 2015.
- [9] David Goldblatt. P1916R0: There might not be an elegant OOTA fix. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2019/p1916r0.pdf>, October 2019.
- [10] Maurice Herlihy, Victor Luchangco, and Mark Moir. The repeat offender problem: A mechanism for supporting dynamic-sized, lock-free data structures. In *Proceedings of 16<sup>th</sup> International Symposium on Distributed Computing*, pages 339–353, October 2002.
- [11] Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. Repairing sequential consistency in C/C++11. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017*, pages 618–632, New York, NY, USA, 2017. ACM.
- [12] Paul E. McKenney. Proper care and feeding of return values from `rcu_dereference()`. [https://www.kernel.org/doc/Documentation/RCU/rcu\\_dereference.txt](https://www.kernel.org/doc/Documentation/RCU/rcu_dereference.txt), February 2014.
- [13] Paul E. McKenney. *Is Parallel Programming Hard, And, If So, What Can You Do About It? (2018.12.08a Release)*. kernel.org, Corvallis, OR, USA, 2018.
- [14] Paul E. McKenney, Alan Jeffrey, Ali Sezgin, and Tony Tye. Out-of-thin-air execution is vacuous. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/p0422r0.html>, July 2016.
- [15] Paul E. McKenney and Jack Slingwine. Efficient kernel memory allocation on shared-memory multiprocessors. In *USENIX Conference Proceedings*, pages 295–306, Berkeley CA, February 1993. USENIX Association. Available: <http://www.rdrop.com/users/paulmck/scalability/paper/mpalloc.pdf> [Viewed January 30, 2005].
- [16] Paul E. McKenney and John D. Slingwine. Read-copy update: Using execution history to solve concurrency problems. In *Parallel and Distributed Computing and Systems*, pages 509–518, Las Vegas, NV, October 1998.
- [17] Maged M. Michael. Safe memory reclamation for dynamic lock-free objects using atomic reads and writes. In *Proceedings of the 21<sup>st</sup> Annual ACM Symposium on Principles of Distributed Computing*, pages 21–30, August 2002.
- [18] Maged M. Michael. Hazard pointers: Safe memory reclamation for lock-free objects. *IEEE Transactions on Parallel and Distributed Systems*, 15(6):491–504, June 2004.



## REFERENCES

---

- [19] Raúl Silvera, Michael Wong, Paul E. McKenney, and Bob Blainey. N2153: A simple and efficient memory model for weakly-ordered architectures. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2007/n2153.pdf>, January 2007.
- [20] Matthew D. Sinclair, Johnathan Alsop, and Sarita V. Adve. Chasing away RAts: Semantics and evaluation for relaxed atomics on heterogeneous systems. In *Proceedings of the 44th Annual International Symposium on Computer Architecture, ISCA '17*, pages 161–174, New York, NY, USA, 2017. ACM.