Permchecker: A Toolchain for Debugging Memory Managers with Typestate

KARL CRONBURG and SAMUEL Z. GUYER*, Tufts University, USA

Dynamic memory managers are a crucial component of almost every modern software system. In addition to implementing efficient allocation and reclamation, memory managers provide the essential abstraction of memory as distinct objects, which underpins the properties of memory safety and type safety. Bugs in memory managers, while not common, are extremely hard to diagnose and fix. One reason is that their implementations often involve tricky pointer calculations, raw memory manipulation, and complex memory state invariants. While these properties are often documented, they are not specified in any precise, machine-checkable form. A second reason is that memory manager bugs can break the client application in bizarre ways that do not immediately implicate the memory manager at all. A third reason is that existing tools for debugging memory errors, such as Memcheck, cannot help because they rely on correct allocation and deallocation information to work.

In this paper we present Permchecker, a tool designed specifically to detect and diagnose bugs in memory managers. The key idea in Permchecker is to make the expected structure of the heap explicit by associating *typestates* with each piece of memory. Typestate captures elements of both type (e.g., page, block, or cell) and state (e.g., allocated, free, or forwarded). Memory manager developers annotate their implementation with information about the expected typestates of memory and how heap operations change those typestates. At runtime, our system tracks the typestates and ensures that each memory access is consistent with the expected typestates. This technique detects errors quickly, before they corrupt the application or the memory manager itself, and it often provides accurate information about the reason for the error.

The implementation of Permchecker uses a combination of compile-time annotation and instrumentation, and dynamic binary instrumentation (DBI). Because the overhead of DBI is fairly high, Permchecker is suitable for a testing and debugging setting and not for deployment. It works on a wide variety of existing systems, including explicit malloc/free memory managers and garbage collectors, such as those found in JikesRVM and OpenJDK. Since bugs in these systems are not numerous, we developed a testing methodology in which we automatically inject bugs into the code using bug patterns derived from real bugs. This technique allows us to test Permchecker on hundreds or thousands of buggy variants of the code. We find that Permchecker effectively detects and localizes errors in the vast majority of cases; without it, these bugs result in strange, incorrect behaviors usually long after the actual error occurs.

CCS Concepts: • **Software and its engineering** → **Software maintenance tools**; *Source code generation*; *Runtime environments*.

Additional Key Words and Phrases: typestate, debugging, language implementation, memory management, memory layout, compiler extension

ACM Reference Format:

Karl Cronburg and Samuel Z. Guyer. 2021. Permchecker: A Toolchain for Debugging Memory Managers with Typestate. *Proc. ACM Program. Lang.* 5, OOPSLA, Article 149 (October 2021), 28 pages. <https://doi.org/10.1145/3485526>

Authors' address: Karl Cronburg, karl@cs.tufts.edu; Samuel Z. Guyer, sguyer@cs.tufts.edu, Tufts University, 161 College Ave, Medford, Massachusetts, 02155-5593, USA, USA.

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

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/10-ART149

<https://doi.org/10.1145/3485526>

1 INTRODUCTION

In December of 2018, a bug was reported to Oracle's OpenJDK [Oracle 2006] development team. In this bug, the Java Virtual Machine (JVM) would deadlock on an object monitor, even though no threads held exclusive access to the monitor. This failure only occurred when the jemalloc [Evans 2006] was used in place of the default malloc/free implementation, leading the developers to suspect a memory safety error. However, the failure was difficult to reproduce and caused knock-on effects leading to incomplete and misleading stack traces. After ten months of on-again off-again investigation, the team had all but given up, with one comment added to the bug report as follows:

In the Fall of 2019, this developer finally figured it out using carefully hand-coded instrumentation of object monitors. An address mix-up was causing the monitor in question to be owned by two different threads. After almost a year from the initial report, the bug was found and fixed: an address comparison that should have been a greater-than was using a greater-than-or-equal-to operation instead.

This kind of bug is among the most difficult and time-consuming to diagnose because it involves the very subsystem that we ordinarily rely on to *enforce* memory safety: the memory manager, and an object model describing how objects and their metadata are laid out in memory. For this reason, many of the existing tools and techniques for tackling memory errors are ineffective or cannot be easily applied to code in the memory manager. At the same time, the programming of such code is very challenging and error prone, involving sensitive pointer manipulations, implicit address invariants, and complex memory layouts accessed by disparate subsystems. We normally expect runtime system code to support a diverse array of hardware and software environments, where concurrency is the norm and performance is at a premium. Although the number of memory manager implementations is relatively small, every single program that dynamically allocates memory relies on one to run correctly.

In this paper we present Permchecker, a toolchain designed to aid in the debugging of memory managers, including explicit allocators and garbage collectors. Permchecker's design is based on the observation that all memory managers perform the same basic task: they take a large array of bytes and partition it into chunks, giving structure to memory and managing the lifecycle of objects for the client application. The specific partitioning strategy of a memory manager gives each chunk of memory meaning, much the way types give meaning to data in a programming language. An error in a memory layout looks much like a type safety error: a chunk of memory designated for one purpose is accessed improperly or used for another purpose.

Memory chunk types, however, have important differences from traditional types. For one, chunk types often cannot be explicitly defined in the type system of the implementation language. Instead, the role of a chunk is implied by its relative spatial location in memory. For example, a generational garbage collector might classify a page of memory as part of the nursery because it resides in the nursery's reserved address range. Similarly, the argument to a manual free() implementation is an untyped pointer, and this function might need to inspect the pointer value

or metadata stored elsewhere in memory to determine, e.g., that the chunk belongs on a free list. Traditional types alone are a poor fit for these for these situations because the correct type information is not readily available from either type declarations or the operations performed.

Another difference from traditional types is that a chunk of memory changes type during the normal course of partitioning, allocation, deallocation, and coalescing. For example, when a page is divided up into cells, all subsequent code must access memory in the page as cells, even though they reside within a *page*. This behavior is different from dynamic typing, where a single *variable* can refer to values of different types, but the values themselves do not change type. The most closely related type feature is C/C++ union, which allows a single chunk of memory to be viewed as different types, but provides no mechanism to check and maintain type safety.

The key idea in Permchecker is to use *typestates* [Strom and Yemini 1986] to specify and track the structure of memory as it changes, and check that the memory manager accesses and updates this structure according to the expected specification (the “*permissions*”). The system consists of two main parts: (1) an API that developers can use to annotate memory management code with the expected types and state transitions in each operation (e.g., alloc, free, scan, collect, etc), and (2) a runtime system that associates typestates with real memory (a shadow map) and uses binary instrumentation to check that at every memory access the permissions associated with the code match the typestate of the memory. We show that this technique detects memory errors as soon as they happen, rather than when they manifest, which can be much later in a program execution.

The term typestate, as we apply it to a dynamic setting, is slightly different from its static counterpart. In a static analysis setting, a typestate requires three components. First, types describe the shape of data and the allowable operations over a variable. Second, states and accompanying state transitions narrow the allowable operations over the variable. Finally, a semilattice with a meet function defined over it approximates the true state of the variable at each static program point by defining how to combine state information at control-flow merge points.

With Permchecker **the first two components of typestate are defined, but the third is not needed**. First, types describe the in-memory layout of a data structure and the allowable operations over a memory location therein. Second, states and accompanying state transitions restrict the memory location to being accessed by specific program instructions. Finally, program annotations define code permissions which exactly track state information meaning no semilattice approximation is necessary.

A goal of this work is to build tools that work with existing memory managers across a range of platforms. Our approach is designed around this use case. We assume that developers cannot reimplement their systems in a new language or using a new methodology. Therefore, our API is generated from a separate specification and added to the memory manager code. To avoid high up-front costs, developers can introduce these annotations gradually, starting with a simple, but coarse model of memory, and increasing the precision as necessary with more fine-grained annotations.

Contributions. In this work, we apply *typestate* to a dynamic setting for the purpose of describing how a memory manager partitions and recycles memory. The goal of Permchecker is to provide a way for the developer of a runtime system to specify and check the lifecycles of memory typestates across all abstraction levels of the system. A relatively simple tracking API allows the developer to associate a typestate with a chunk of memory. A permissions API and annotations can then be used to declare access permissions. This formulation of dynamic typestate is a natural extension of the static typestate program analysis techniques developed in the 1980’s [Strom and Yemini 1986].

The contributions presented in this paper are as follows:

- A technique for specifying the hierarchical layout and decomposition of a memory manager’s memory.

- A tool for generating an API from a layout specification. The API tracks structures of memory via explicit tpestates transitions.
- Compile-time support, implemented in the Clang C/C++ compiler, for automatically translating common tpestate transitions from function annotations to API calls.
- A Pin-based dynamic binary instrumentation (DBI) tool that tracks the tpestate of every byte of memory and ensures that memory accesses from the anywhere in the program (both the memory manager and the application) respect the expected structure.
- An experimental methodology that includes both real bugs and injected bugs. Since bugs in memory managers are relatively uncommon, we use a technique in which we extract the essence of a real bug as a bug pattern, and then run hundreds of experiments in which similar defects are introduced in different places in the code.

2 DESIGN OVERVIEW

Permchecker takes a step towards verifiable memory management by providing a way to explicitly annotate the structure and state of a heap, and then check that the implementation of the memory manager properly maintains and respects that structure. Our design is inspired by Valgrind's Memcheck tool, but generalized to an arbitrarily hierarchical memory manager. At the application level, memory is either allocated or not, and Memcheck checks that application code only accesses allocated memory. Inside a memory manager, there are often numerous intermediate states of memory apart from allocated and free. Memory is obtained from the operating system in large chunks and carved into successively smaller pieces. Each partitioning is typically managed by an intermediate allocator which services requests from higher-level allocators and requests memory from lower-level allocators. Additionally, multiple independent allocators often service requests for specific categories of values, contributing further to the rich landscape of memory subsystems that compose a modern runtime system.

Allocation Policies & Alternative Hierarchies. As an example, a generational garbage collector might contain three independent allocators: a bump pointer for the nursery, a free-list allocator for the mature space, and a reference-counted allocator and collector for large objects. For instance, the free-list allocator has at least two allocation levels. First, a low-level block allocator carves off a large chunk of the virtual address space from the operating system. Second, the free-list's cell allocator extracts from the large chunk an object-sized cell and doles it out to the application. This process is further complicated by a myriad of allocation and collection policies at both the block and cell level. As a result a fully automated technique, like Memcheck, is infeasible in the absence of definitive memory type and allocator state information.

Tpestate Permissions. It is crucial for runtime system code to only ever access the kind of memory it is supposed to. For example, a free-list allocator should not access memory currently allocated by the nursery's bump pointer allocator. Given explicit expectations of these two subsystems, we can begin to verify that memory is managed safely by ruling out cross-contamination.

With Permchecker, each location in memory is assigned a type, or more precisely a tpestate since the type can change over the course of execution. After assignment, code which accesses the location must have permission to access the assigned tpestate. An error is thus reported as soon as a memory access occurs over a tpestate of memory the code was not expecting to access. The resulting error message is a detailed report of the impermissible memory access, and is described in the terminology of the particular system. For example, a garbage collector's nursery allocator might incorrectly attempt to allocate memory recently added to the mature space's free-list. The error reported by Permchecker would be something like the following: *"Observed type FREELIST_CELL expecting type NURSERY_CELL."*

Incrementality. A primary goal of this work is to provide practical debugging tools and techniques that can help diagnose bugs in existing systems. To this end, Permchecker is built as a lightweight C/C++ API and accompanying heavyweight instrumentation tool runnable on virtually any memory manager. We present our use of this tool to debug OpenJDK’s Hotspot VM, Oracle’s open-source version of their industrial-strength JVM. Not discussed experimentally in this paper is our use of Permchecker to debug memory bugs found in other manual and automatic memory managers. These include `dlmalloc` [Lea 1991], Doug Lea’s implementation of the standard `malloc/free` interface, and the Jikes Research Virtual Machine [IBM 2005], another implementation of a JVM. Without the trial-and-error process of studying this multitude of diverse systems, we would not have landed upon an *incremental* usage of Permchecker where the programmer starts out with a simple model of their heap, adding more and more tpestate information over time.

3 BACKGROUND & RELATED WORK

Memory safety is not a new concern. Over the years, a wide array of systems, tools, techniques, and languages have been developed to help diagnose and prevent these potentially catastrophic errors. One of the most successful and widely deployed techniques is automatic memory management (garbage collection) which, together with runtime bounds checking, has practically eliminated the most common coding errors that lead to memory corruption. But a nagging problem has remained: how to find and fix bugs in the memory manager itself. In this section, we discuss how this concern is different from ordinary memory safety and why that makes it such a hard problem to deal with.

Debugging an Abstraction. Memory safety ensures that a memory access respects the boundaries and lifetimes of objects and values in memory. At the level of raw memory (virtual addresses), however, there are no boundaries and lifetimes; memory is just one giant array of bytes. These abstractions are created and maintained by the memory manager, and unfortunately existing memory safety tools require this information in order to perform their checks. Valgrind [Nethercote and Seward 2007], for example, can perform highly detailed memory safety checks on an application, but it must be able to intercept (shim) calls to `malloc` and `free` in order to know when objects are allocated and freed, and their bounds. When the memory manager itself has an error, however, the memory checker becomes unreliable because it is using faulty information.

Missing Information. A memory manager bug is similar in spirit to a compiler bug: a rare “bug of last resort” that is often identified only after extensive application-level debugging fails. In a compiler bug, however, it is possible to inspect the assembly code output and verify that it correctly implements the input source code (or not, as the case may be). No such opportunity exists for errors in the heap layout of a memory manager. Most implementations do not include a formal description of what correct states of the heap can look like. Therefore, we cannot merely inspect the resulting contents of memory to determine if the memory manager has functioned correctly. For this reason, a preprocessing tool in the Permchecker toolchain is inspired by LoCal [Vollmer et al. 2019] and Floorplan [Cronburg and Guyer 2019], both of which are memory layout description languages intended to provide special purpose heap access abstractions.

GC Debugging. Recent research targeting, among others, the multicore OCaml garbage collector (GC) tackles the issue of unreliability in debugging tools for concurrent GC bugs. While memory corruption is one source of unreliability in producing accurate error reports, another source of unreliability is the non-deterministic nature of many GCs. The RR [O’Callahan et al. 2017] debugger, an extension to the GNU Debugger, alleviates nondeterminism while debugging. It does this by recording and replaying program executions with relatively low overhead. The developers of RR

report that a principle motivation in developing it the way they did was “to improve the [experience] of [using] user-space record-and-replay while being less invasive than [e.g.] a full kernel implementation”. These principles are shared by our approach to debugging garbage collectors. We studied the practicality of adding small amounts of GC instrumentation for the purpose of reliably reporting the violation of unobservable program invariants.

Static Solutions. Typestate finds numerous applications in both static analysis [Strom and Yemini 1986] and static programming language features. For instance typestate in Rust [Weiss et al. 2019] derives from the language’s support for linear and affine types at compile-time, by enforcing memory movement semantics statically in the context of a single variable. In contrast typestates supported by Permchecker derive from the toolchain’s tracking of shadow memory at run time, and enforcing memory access permissions dynamically in the context of an individual memory location. The key difference here is that a single variable cannot be aliased by constructing it from dynamic calculations, whereas a memory location *must* be aliasable via pointer operations in the implementation of a memory manager.

Smart Pointers. Much like linear and affine types for enforcing usage constraints at compile-time on variables, a smart pointer can check or enforce certain usage properties at run time *on a variable* as well. One canonical example is a reference-counted smart pointer for distinguishing live allocated objects from garbage. Such a reference counting scheme tracks the number of known pointers stored in memory in places where they can, in the future, be loaded into variables (registers, at run time) and referenced as such with a load or store instruction. This ability to access memory indicates the memory remains allocated. A memory manager however not only manages allocated memory: it also manages and requires access to any and all freed memory. Thus, without reachability as a proxy for whether or not a memory location should be accessible, reference-counted smart pointers alone cannot sufficiently check memory management code for safety.

Compiler Extensions. A compiler extension known as AddressSanitizer [Serebryany et al. 2012], like Valgrind’s Memcheck, is able to detect out-of-bounds accesses to memory locations with shadow memory. Unlike Memcheck, it relies on unaddressable “poisoned” padding to be added to heap, stack, and global chunks of memory. As a result this mechanism is able to detect a subset of inter-chunk corruptions, but not any intra-chunk corruptions. Detecting intra-chunk corruption requires a richer allocation distinction than just allocated/free, and fewer assumptions about where memory comes from. The key problem we face, in the domain of memory management, is that multiple memory allocation schemes share the same single virtual address space.

Bug Injection Methodologies. The process of generating a corpora of bugs is often motivated by a specific kind of mistake a programmer can make: off-by-one errors, operator selection defects [Rice et al. 2017], bounds-checking mistakes [Dolan-Gavitt et al. 2016], and more. These mistakes can be constructed by source code mutation [Roy et al. 2018]. What this bug creation technique does not take into account is the modeling of undesirable algorithmic operations of the system having bugs injected into it.

In a memory manager, we know of a relatively fixed and small set of algorithmic operations that are bad: use-after-free, overlapping allocations, among others. In contrast, the categories of linguistic (syntactic and semantic) mistakes a programmer can make are limited by which symbols the programmer types and the compiler accepts. The key difference is that injecting invalid algorithmic operations into a system realistically stresses a broad category of resulting downstream behaviors of the system, while injecting linguistic bugs does not. Injecting a linguistic bug requires both a far larger enumeration of sources of linguistic mistakes and clever tactics for eliminating repetitive bugs, such as ones which always crash the program.


```

1  <id>   := [A-Z][_a-zA-Z0-9]*
2  <ts>   := <id>    // A typestate is a unique identifier
3  <prim>  := bytes | words
4  <stmt> := <ts> -> <ts>
5      | <ts> -> seq  { <exp> , ... } // One or more ', '-sep exps
6      | <ts> -> union { <exp> | ... } // One or more '| '-sep exps
7      | # <exp> // Prefix notation for "zero-or-more" reps
8  <exp>  := <ts> | <stmt> | [0-9]+ <prim>
9  <spec> := <stmt>+

```

Fig. 1. Syntax for describing the layout of memory in terms of typestates. A layout <spec> is one or more layout statements describing a sequence of memory (seq), alternative views of the same memory (union), some amount of primitive memory (<prim>), or repetitions of a structure with (#).

4 TOOLCHAIN OVERVIEW BY EXAMPLE

In this section we present in modest detail the mechanisms by which a memory management programmer and the Permchecker toolchain orchestrate the checking of typestates in a memory manager. In doing so we illuminate the role of the compiler, VM annotations, and dynamic instrumentation in tracking and checking the typestates. The idea here is that each toolchain component is small or simple in nature and designed to serve a single purpose effectively: manage and check uniquely identifiable typestates over a process' address space. To start, the following are some important high-level definitions clarifying the scope of typestate in this work:

Typestate. This is a uniquely identifiable state of a memory location, coupled with the valid value types that can reside in that location. A mutable typestate on a heap memory location for the duration of a program's execution is similar to a dynamic type on a local variable for the duration of a function's execution. The key difference is that a type typically restricts a function call-site to referring to a specific set of variables, whereas a typestate restricts an assembly instruction to accessing a specific set of memory locations.

Set of typestates. This is a collection of typestates, each of which can be accessed by the same piece of polymorphic code. Such a collection associated with a specific piece of code permits access to memory locations containing the same types of values, without unnecessarily restricting the code to a highly particular typestate.

4.1 Allocation Hierarchies

In Figure 1 we define a minimal syntax for defining typestates of memory. The idea is that typestates are interrelated via the hierarchy by which they are allocated. For example in Hotspot, a page resource allocates for regions, the region manager allocates for objects, and object management code allocates object headers and various primitive application field types. To broadly describe this hierarchy with the syntax, we might write the following:

```

1  Pages -> # union { Region | ... }
2  Region -> seq { # Object, # words }
3  Object -> seq { Header, # Field }
4  Header -> union { 2 words | Array -> 3 words }
5  Field -> # words

```

(L1)

Code L1 defines 6 tpestates: Pages, Region, Object, Header, Array, and Field. The idea is that each and every memory location in use by the memory manager is in one of these tpestates at some point during runtime (or implicitly in an UNMAPPED tpestate). For example, an allocated heap region with a single 10-word object allocated into it involves three tpestates. The first 2 words of the region are in the Header tpestate, the next 8 words are Field, and the remaining memory is all Region. In this example, the first # repetition in the Region statement effectively takes on a value of 1 (a single object) and the second # repetition consumes the remaining words in a (fixed-size) region. Also note that the ellipsis on the first line is a placeholder for other allocation schemes that we do not discuss here.

The repetition (#) operator means *some number of* chunks of memory. Intuitively, this operator flood fills the surrounding space by repeating in memory zero or more times the layout expression it is applied to. Formally, the semantics of repetition are equivalent to those of the Kleene star regular expression operator. For example, the (#) operator can describe gaps in the address space of statically unknown size inbetween known structures. This enables the description of alignment padding on an object, dead objects fragmented among live allocated objects, and more.

4.2 Layout Description

In addition to defining a hierarchy, we need to be able to declare what the intended contents (values) of a memory layout are. In a memory manager, navigating a memory layout to obtain values in memory is performed by pointer manipulations and offset calculations. These calculations ultimately lead us to the underlying contents of memory. To declare the underlying contents of objects from Code L1, we might for example write the following:

```

1 Object -> union {
2     Free  -> # words
3     | Alloc -> seq { Header, # words } }
4 Header -> union {
5     Arr -> seq { MarkWord, KlassPtr, Len, Gap } // 3 words
6     | Cls -> seq { MarkWord, KlassPtr } }           // 2 words

```

(L2)

In this Code L2, we define the tpestates of a Hotspot Object and its Header. An object is either in the Free tpestate, or it is in the Alloc tpestate and therefore contains a header and some number of payload words of memory which we leave unrefined. The Header of such an object is either 3 or 2 words in size. This depends on whether or not the object is an array with an array length (Len) field. The Arr is this array type, whereas the Cls is an ordinary Java class object. A KlassPtr is a reference to C++ metadata about a Java class. The Gap field is unused alignment padding. Lastly, a MarkWord is a bit-field containing runtime metadata such as the object's hash, locking bits, number of object promotions or evacuations, and a mark bit. In some cases these fields overlap and change depending on the state of the object, especially in relation to per-thread object locking. Apart from the mark bit and promotion count ("age"), no space is reserved in the mark word for the GC. Assuming a 64-bit architecture, the header fields take on the following sizes:

```

1 MarkWord -> 1 words
2 KlassPtr  -> 1 words
3 Len       -> 4 bytes
4 Gap       -> 4 bytes

```

(L3)

Code L2 and L3 serve two purposes: (1) to define hierarchical relationships among named typestates, and (2) to associate sizes with them. With (1), the idea is that this code models where a piece of memory comes from in the context of a hierarchy of typestate mutators. One such typestate mutator is the code in Hotspot which stores a garbage object onto a free-list. This mutator code implicitly changes the typestate of the underlying memory from allocated to free. However, this is in some sense a lazy operation. Perhaps we could deem an object to be free as soon as a Java-level pointer update causes the object to no longer be accessible by the Java heap. While more precise, this definition is untrackable without invasive changes to the runtime system. The point is it is not always clear, based on a cursory analysis of the code, where a typestate mutation should occur. This state of affairs reflects our dynamic formulation of typestate for tracking memory *as implemented*, with minimal disruption to that implementation.

4.3 Typestate Identifiers: Code Generation

In this section we discuss how we annotated the Hotspot VM with typestate layout annotations generated from a version of Code L1. The idea is that we can generate much of the repetitive boilerplate necessary to track typestates at run time. This boilerplate includes a consistent centralized naming scheme for managing hierarchies of allocable typestates.

First, a preprocessing tool processes a layout description, generating a series of named typestates, among other code snippets to reduce boilerplate during annotation of a memory manager. From Code L2, we get out a series of typestates defined as follows:

```

1  #define PCK_Object ((TypeState)1)
2  #define PCK_Object_Free ((TypeState)2)
3  #define PCK_Object_Alloc ((TypeState)3)
4  #define PCK_Object_Alloc_Header ((TypeState)4)          (C1)
5  #define PCK_Object_Alloc_Header_Arr_MarkWord ((TypeState)5)
6  #define PCK_Object_Alloc_Header_Arr_KlassPtr ((TypeState)6)
...
30 #define PCK_MarkWord ((TypeState)30)
31 #define PCK_KlassPtr ((TypeState)31)
...
```

In this series of typestates in Code C1, TypeState is a class for unique identifiers. Each identifier is referenceable with a unique numeric value, as indicated by the implicit cast constructors referenced above. The number zero (0) is reserved for PCK_UNMAPPED memory. Any top-level defined <stmt> gets associated with it a nested hierarchy of addressable typestates. For example, a memory location can be part of a free (PCK_Object_Free) or allocated (PCK_Object_Alloc) object. Notably this means that no typestates of the form PCK_Alloc_* are generated, but indeed PCK_Object_* *are* generated. The key is that one can choose which typestate contexts are diagnostically helpful to track based on how memory is managed.

One important distinction is between the generated typestates PCK_Object_Alloc_Header_Arr_MarkWord and PCK_MarkWord. The latter, PCK_MarkWord, represents memory for a top-level C++ variable not located in the Java heap. The former is a fully-qualified mark word of an allocated array object header *located on the Java heap*. This typestate represents memory which was allocated by the proper Java object allocation code. The distinction here is necessarily self-enforced, based on annotation usage, but it is a diagnostically helpful distinction. The idea is that by default internal memory allocation context is all-but lost once an initialization is complete. With the two

```

1  namespace pck {
2  class Map;
3  class TypeState; // Includes an identifying member field of u64_t
4  Map create();
5  void destroy(Map m);
6  TypeState Map::get(Address a);
7  void      Map::put(Address a, TypeState ts); }

```

Fig. 2. The VM-level interface to initialize and track a memory to tpestate mapping. All functions here are in the pck (Permchecker) namespace.

tpestates however, we can now retain the context for validation or diagnosis purposes if and when a corruption occurs.

However, the tpestate distinction comes at a price. In order to specify when the distinction does *not* matter, we must annotate the code with all of the allowable tpestates. Instead of doing this manually, the code generator does this for us. The generator creates names for collections of tpestates based on all the qualified names by which any given “leaf” tpestate can be allocated. For example, for the mark word tpestates from Code L2 we get the following collection:

```

1  #define PCK_ALL_MarkWord PCK_MarkWord \
2    PCK_Object_Alloc_Header_Cls_MarkWord \
3    PCK_Object_Alloc_Header_Arr_MarkWord \
4    PCK_Header_Arr_MarkWord PCK_Header_Cls_MarkWord

```

(C2)

In this macro we have (1) a C++ variable mark word, (2) an allocated class’s or (3) array’s mark word on the Java heap, and (4) a C++ array variable’s or (5) Java class variable’s mark word. This collection is a fairly common idiom in memory managers. That is, a mark word accessor function typically does not care which variant it accesses. By annotating code with this PCK_ALL_MarkWord tpestate collection, we get the benefit of a coarse polymorphic tpestate when the precise distinction does not matter. At the same time, we still get the benefit of a detailed taint analysis when Permchecker reports an error in terms of the precise tpestate observed at run time.

4.4 Tpestate Tracking

The VM’s interface to the tpestate tracker is shown in Figure 2. Apart from its connection to dynamic instrumentation, this interface operates how one would expect it to. An individual shadow map can be created and destroyed, with each map maintaining an injective key-value mapping from addresses to tpestates. The special UNMAPPED tpestate is the default value for unaddressable memory, and UNMAPPED can be explicitly referenced as a tpestate. The interface further uses array-access overloading to provide the following map update and query syntax:

```

1  Address a; TypeState ts;
2  map[a] = ts;
3  pck::assert(map[a] == ts)

```

(C3)

In this Code C3, we have defined some address *a* and some tpestate *ts*. On the second line, an assignment operation calls `Map::put` to assign the r-value *ts* to the map location referenced by address *a*. On the third line, an equality comparison operation calls `Map::get` to query the tpestate at address *a* and check that it matches *ts*. Note that the equality and map-access operations return

proxy object types tracking what address was accessed, and the arguments to any comparisons. This proxy information allows `pck::assert` to report a helpful error message in terms of the address, observed tpestate, and expected tpestate of the comparison, rather than simply reporting that the assertion failed.

As we will discuss in the next section, the dynamic instrumentation effectively performs an assertion, like above, over the tpestate map for each and every memory access executed by the processor. To simplify our discussion, we focus on the (sufficient) use case of tracking a single tpestate map. In this case, the `create` and `destroy` functions get called near the beginning and end of program execution, respectively. In the case of Hotspot, some complication arises from the use of `fork/exec` system calls, requiring explicit map initialization after the main Java thread is created. A more careful integration with thread and process creation system calls could provide more automation of this process in the future. For now, this is unnecessary work to automate a very small amount of code that a memory management developer knows exactly where to put.

4.4.1 In Practice. We initially inserted 52 assertions related to object mark-words and klass pointers in Hotspot. Of these, 11 assertions were subsequently replaced with annotations on simple accessor functions, to be discussed in Section 4.5.1. A typical pattern in a memory manager is for a simple one-line accessor function to access a specific set of tpestates, and so a function-level annotation is ideal. In contrast, manual assertions are ideal when instrumenting (1) compiled Java code, (2) other VM components that cannot easily rely on the C++ compiler, and (3) when substantial code refactoring would be necessary to systematically instrument certain memory accesses.

4.4.2 Code Generation. For some tpestates, Permchecker’s preprocessor generates helper macros. One such macro manages all the tpestate updates necessary to allocate a series of numerous nested components. For example, for tracking the nested tpestates associated with an array header the preprocessor generates the following macro based on Code L2:

```
1 #define transition_Object_Alloc_Header_Arr (a) { \
2   map[words(a, 1)] = PCK_Object_Alloc_Header_Arr_MarkWord; \           (C4)
3   map[words(a + words(1), 1)] = PCK_Object_Alloc_Header_Arr_KlassPtr; \
4   map[bytes(a + words(2), 4)] = PCK_Object_Alloc_Header_Arr_Len; \
5   map[bytes(a + words(2) + bytes(4), 4)] \
6     = PCK_Object_Alloc_Header_Arr_Gap; }
```

In this Code C4 we see a series of array-update operations over the shadow map, as explained earlier. The `words` and `bytes` functions in this code return a C++ proxy class indicating the location and size of a piece of memory to update in the tpestate map. For the two-parameter versions, the first parameter is a starting address and the second parameter is a number of bytes to assign the tpestate to. For the one-parameter version, the parameter is just a memory size for use in offset calculations. For example, the offset calculation in the second map update above computes the address of a offset by one word of memory. The key is that this code fragment conceptually (if not literally, for performance) boils down to a series of appropriate calls to `Map::put`.

4.5 Permission Tracking

Next, we need to be able to assign permissions to a piece of code in terms of the tpestates the code is allowed to access. The primary mechanism the toolchain supports for this purpose is permission attributes on functions and classes in C++. The exact low-level mechanism to track this information is interception of specific function calls by a dynamic instrumentor. We conclude this section by discussing what bugs this dynamic instrumentor can and cannot detect as tpestate violations.

```

1 namespace pck {
2   enum Perms { RW, R, W, NONE };
3   Map::protect(InstanceState ts, Perms ps);
4   Map::unprotect(); }

```

Fig. 3. The VM-level interface to set memory access protections by tpestate, instead of a concrete address like with the POSIX `mprotect` system call. The `Map::protect` method here applies only to the calling thread, and temporarily overrides any previous protections for the given tpestate. These previous protections are restored by a subsequent matching call to `unprotect`.

		Total	Mark & Klass
1	<code>__attribute__((pck_R (InstanceState...)))</code>	Lines added	1830
2	<code>__attribute__((pck_W (InstanceState...)))</code>	Transitions	144
3	<code>__attribute__((pck_RW(InstanceState...)))</code>	Fncn Annots	64
4	<code>__attribute__((pck_NONE(InstanceState...)))</code>	Typestates	60
			16

Fig. 4. **Left:** Read, write, and read-write annotation forms which can be attached to either a function or a class to indicate that entity has permission to access the listed tpestates. The `pck_NONE` option allows one to explicitly *remove* access, notably on a member function where the class generally has permission but the member function should not. **Right:** Lines of code added/modified in the VM to annotate/instrument the most crucial parts of Hotspot (enabling Java object allocation). The “Mark & Klass” column counts lines of code relating to just `MarkWord` and `KlassPtr` variants. The bottom-most row is a count of unique tpestates referenced in the instrumentation.

The VM’s interface to the permission tracker is shown in Figure 3. This interface models memory access permissions by associating with each thread of execution a set of per-tpestate permissions. For example, each thread which calls an accessor function with read or write access to the mark word of an object header can temporarily be given permission to access any such mark word. This does not guarantee the correct mark word is accessed, just that the underlying memory was allocated as such. The key is that an individual function, class, or algorithmic operation typically either has permission to access the mark word, or it does not.

The `protect` and `unprotect` functions implement this permission model. These functions operate similarly to the `mprotect` POSIX system call, but with a few differences. A call to `protect` states that the currently executing thread now has permission to access any memory address in the given tpestate. A subsequent matching call to `unprotect` reverts the permissions to what they were before. These functions are thread-safe, because each thread has its own permissions. These functions are *not* safe with respect to non-reentrant control flow, such as runtime system exception handling. Future work could automatically support arbitrary control flow by associating permissions with instruction pointers instead of function entry and exit points.

In designing this part of Permchecker, a simpler model was considered where permissions could only be updated with `protect` and the previous permissions are forgotten. This simpler stateless model, however, fails to account for a nested call frame for which the parent and child frames execute functions both with permission to access the same tpestate. Calling `unprotect` during the child’s epilogue would therefore incorrectly wipe out the permissions the parent still needs to possess. Therefore the stateful implementation of `protect` in the dynamic instrumentor tracks a stack of bit-vectors corresponding to nested function annotations.

4.5.1 Function & Class Annotations. In Figure 4, on the left, we have a series of tpestate permission annotations for functions and classes. On a function, an annotation compiles to calls to

protect and unprotect in the prologue and epilogues of the function. On a class, the annotation gets distributed across functions in that class to the same effect. Based on this model, an annotated function calling some other function conservatively grants the callee the same permissions it has.

Also in Figure 4, on the right, is a breakdown of modifications made to the Hotspot VM to achieve permission checking of object header words. The modifications reported are the ones attributable to the memory management developer, with the **Total** column including modifications necessary to support tracking of the Java heap across the entire allocation hierarchy. These numbers represent a reasonable upper bound on the annotation burden of dealing with a small number of the most complex tpestates in a memory manager.

The lines of code added to Hotspot include modifications to the JIT compiler, assembler, argument parsing, imports, error reporting, foreign function interface, and write barrier. Apart from these tasks, 144 and 64 lines were respectively added to track tpestates and to apply permissions to code. Of those tpestate management lines of code, 49 deal with just a variety of 16 tpestates pertaining to the mark or klass word of an object. The remaining 44 of 60 tpestates provide diagnostically helpful allocation context for when the mark or klass word is accessed erroneously. Such tpestates include allocated object fields, free regions, reserved pages of memory, and others.

The amount of annotation effort required to validate layout safety for a tpestate correlates with the number of places in the code memory changes to/from the tpestate, plus the number of instructions allowed to access it. As a result simpler kinds of memory, like the page allocator in Hotspot, require relatively few annotations. The desired level of specificity then governs how many different kinds of memory are annotated, leading to a higher annotation effort.

Permchecker allows for progressively annotating each tpestate in isolation while still being effective for diagnosing violations of the annotated tpestate(s). Thus the value of Permchecker is directly proportional to the number of completely annotated tpestates. Annotating the appropriate places in the code for a single tpestate is the smallest unit of work for a developer to accomplish when getting started with Permchecker. In our experience such a task can take anywhere from a few minutes, to a few hours, to at most a few days of a single developer's time. In the next few subsections we go on to discuss how Permchecker reports helpful tpestate mismatches in the presence of either a specification or an implementation error.

4.5.2 True Negatives. Consider a VM function intended to access the mark word of an object, and we wish to give the corresponding assembly code permission to access mark words therein. Therefore we want the function to compile to an instruction sequence that looks like the following:

```

1 // Begin perms: stack (RW)
2 mov $0x21,%esi
3 mov $0x2,%edx
4 // Begin perms: stack (RW) & mark word (W)
5 call protect // Mangled name: _ZN3pck3Map7protectEyNS_5PermsE (A1)
6 mov -0x28(%rbp),%rax
7 movl $0x1,(%rax) // Write to a mark word
8 call unprotect // Mangled name: _ZN3pck3Map9unprotectEv
9 // Ending perms: stack (RW)
```

Code A1 executes as follows. We store immediate values for a mark word's tpestate (0x21) into register esi and for write access permission (0x2) into edx. On line 5 we call protect to register this permission. Next on line 6, we load an address to an object off the stack and into register rax. Finally, on line 7, we set four bytes of the object header to 0x1 to indicate it is a regular unlocked

object. After this, a call to `unprotect` reverts the thread's permissions to what they were at lines 2 and 3. The idea is that if line 7 writes to 4 bytes in a valid mark word typestate then the program does not have a memory error – a true negative check by Permchecker.

4.5.3 True Positives. Now consider a bug where the mark word annotation is in the wrong place, or the C++ compiler erroneously reorders a memory access to the object's mark word. Thus the memory access occurs outside the scope of the `protect` call as follows:

```

1 // Begin perms: stack (RW)
2 mov $0x21,%esi
3 mov $0x2,%edx
4 call protect // Begin perms: stack (RW) & mark words (W)      (A2)
5 call unprotect // Begin perms: stack (RW)
6 mov -0x28(%rbp),%rax
7 movl $0x1,(%rax) // Violation!
```

In this Code A2 Permchecker will report an access violation because line 7 here did *not* have permission to write to the mark-word even though line 7 is supposed to be able to do so. The key is that Permchecker does not make a claim of truth as to which permission was correct: the code or the memory. By doing so, Permchecker can effectively report localized memory errors without needing to assume *any* code is necessarily correct. Thus, not only is a bug in the implementation reported as a typestate mismatch, but a specification (or compiler) bug is reported as such too.

4.5.4 Benign Checks. In the two snippets of assembly code above, we in fact accessed two different kinds of memory: a mark word, and a pointer to that mark word found on the stack. For the true-negative, both memory accesses occur within the scope of our call to `protect`, but we only needed the mark word access to be checked. In checking both, we made a tradeoff in how Permchecker was designed. The idea is that it is often sufficient to check a memory access against a set of possibly unrelated typestates when doing so simplifies our permission specification.

4.5.5 Exceptions Limitation. When Permchecker instruments a function with calls to `protect` and `unprotect`, it does not handle the presence of throw-catch style C++ exceptions. Nor does it reason about any other dynamic feature which alters ordinary program control flow in a permission-annotated function. A long term goal for Permchecker is to obviate this limitation by directly relating typestate permissions with individual memory access assembly instructions in the code spaces of a process. This high level of individual detail, however, increases by default the upfront annotation cost required to annotate a memory manager. Therefore some care must ultimately go into the design of assembly instruction permission techniques for Permchecker to support.

4.5.6 Failure Modes. A false positive is generally not possible with Permchecker. This is because a typestate error indicates the presence of a bug in the memory manager, the layout specification, or both. Such a bug can even be present and reported as an error when the application runs correctly. Such a report indicates an at-least benign bug that *should* be fixed.

In contrast, a false negative *is* generally possible with Permchecker. For example, a false negative will occur in Code A1 when attempting to write the mark-word actually writes a value to stack memory. Permchecker does not report this as an error because the executing thread had permission to write to the stack when the purported mark word instruction clobbers the stack. To mitigate this kind of source of false negative, we have come up with a dynamic heuristic to understand the sensitivity of the specified code permissions.

Fig. 5. ASCII-art diagram depicting the difference between two memory layouts central to an off-by-one bug affecting memory safety.

The heuristic Permchecker provides is to generate a table of typestate usage statistics at the end of a program’s execution. This table includes which typestates each instruction¹ was observed to have accessed, and whether or not any permissible typestates were never accessed. The idea is that any given instruction in a mostly-bug-free memory manager will access a fixed set of typestates after sufficient testing. Then, any remaining permissible typestates are suspicious. The permission annotation could have been overly broad, blanketing a number of unrelated memory access instructions. Or, the instruction was insufficiently tested. In either case, there exist remedial actions the memory management developer can take which either incrementally improve the precision of the annotations, or improve the typestate coverage of the testing benchmarks.

4.6 Runtime Debugging

With typestate tracking and permissions in place, we now need to check the validity of each memory access. The goal is to ensure that for each address the only load and store instructions operating over it occur when the code has permission to access it. This section discusses how we achieve this goal by presenting an example where a header’s mark word is incorrectly accessed.

4.6.1 Mark Word Example. In Figure 5 we see two distinct memory layouts. On the right is the layout of an allocated Java object, as used in the production version of Hotspot. On the left is a modified version of this layout, with a word of memory containing a forwarding-pointer added to the beginning of the object’s header. In the middle of the diagram arrows pointing left-to-right indicate how the different typestates shift in the address space when the author of the left layout modified a part of the memory manager to use the layout on the right.

In Figure 6, we see the error observed and reported by Permchecker when two different pieces of code disagree on the expected layout of objects in memory. In this case, the VM allocates and initializes the contents of objects according to the production version of the VM where non-array objects consume 2 words of memory for the header. When the VM proceeds to access memory according to the development version with an extra header word, Permchecker reports that an offending memory access read a Klass pointer but expected (had permission to access) mark word typestates. The idea here is that Permchecker does not assume that *either* the observed layout or

¹Requires debugging symbols to get source locations.


```
[error][pck] /jdk13/src/.../oop.inline.hpp:123
Permchecker Violation in Thread #2:
  Expected tpestates: PCK_MarkWord, ...
  Observed tpestate: PCK_Object_Header_Alloc_KlassPtr
  Address = 0x00007ffb1c43e008
  Address is preceded by 1 words of PCK_Object_Header_Alloc_MarkWord
  Address starts 1 words of PCK_Object_Header_Alloc_KlassPtr
  This is followed by 6 words of PCK_Object_Alloc
  Dumping tpestates to /memdbg/permchecker.3952.log
  Dumping core file to /memdbg/core.3952
```

Fig. 6. The error reported by Permchecker when code assuming the left-layout of Figure 5 attempts to access the mark word of an object allocated and initialized by code assuming the layout on the right of that figure.

the expected layout are necessarily correct. In fact, both could be incorrect. Instead, we label code with intentions and Permchecker treats those annotations not as a ground-truth specification of expected behavior, but as a mechanism for pin-pointing inconsistencies.

4.6.2 Lightweight Debugging. Permchecker supports both a lightweight and a heavyweight mode. In the lightweight mode, only memory locations explicitly assigned a tpestate are checked for access permissions. In this mode, a program with no updates to the tpestate map will never produce a violation because all memory locations remain in the UNMAPPED tpestate. The idea is that the lightweight mode eliminates a lot of error detection noise when the memory manager is initially being annotated, favoring local sensitivity over system-wide sensitivity. It does this by only telling the developer when an unexpected piece of code accesses memory in a known tpestate. Memory accesses to locations with unknown tpestates are ignored.

4.6.3 Heavyweight Debugging. In the heavyweight mode, Permchecker reports an impermissible memory access for all tpestates, including UNMAPPED, as an error. In this mode, a program with no updates to the tpestate map will produce a violation for *every single* memory access because no instruction has permission to read or write unmapped memory. The idea is that the heavyweight mode increases the number of true positives once a related set of tpestates have been annotated in the memory manager, thus increasing system-wide sensitivity. While the lightweight mode allows the developer to build the lifecycle of an individual tpestate, the heavyweight mode discovers all buggy, unspecified, or improperly specified memory accesses.

4.6.4 Thread Permissions. In Permchecker's dynamic instrumentor, a per-thread permission tracker is implemented as two extensible bit-vectors representing read and write permissions each with one bit per tpestate. When a bit in the vector is set, the associated thread has that kind of (read or write) permission to the tpestate associated with the offset of the bit into the bit-vector. This mechanism directly relates to how the lightweight and heavyweight modes differ.

The lightweight mode can be thought of as an opt-in checking model, where all tpestates implicitly have their read and write bits set, except an explicitly named set of them. The one exception to this occurs when a thread's permissions are annotated as `pck_NONE(<ts>)`, which does cause the given `<ts>` to be checked. In contrast the heavyweight mode can be thought of as an opt-out checking model, where all tpestates implicitly have their read and write bits unset. As a result every single memory access is checked by default, and the developer must tell Permchecker to explicitly allow a memory access by annotating it.

Table 1. Each feature Permchecker touches, what system component it involves, and the feature's purpose with respect to typestate.

Feature	Component	Typestate Purpose
FFI	VM modification	Traffics typestates from non-native VM code with prerogative over typestate.
Annotations	Clang extension	Distributes a typestate permission over an entire VM component: functions and classes.
Layout Spec	Preprocessing tool	Allows for developer-defined relationships among typestates as an allocation hierarchy.
Codegen	Preprocessing tool	Supports the expression of common permission idioms and typestate transitions.
Macros	VM boilerplate	Manages error-prone calculations for offsets among typestates.
Proxy objects	C++ overloading	Integrates typestate operations with the VM's host language for ease-of-use.
Tracking API	pck namespace	Flexibility to manage typestates at non-function or class boundaries.
Assembly instr.	JIT compiler	Typestate checking of runtime generated code.
Checker	Dynamic Instrumentor	Applies a simple and rational checking model across every single memory access.
Shadow memory	Library	Maintains a snapshot of the typestate map.

5 DEBUGGING TOOLCHAIN ARCHITECTURE

Different components of a VM have access to widely varying mechanisms to support typestate tracking. A simple numeric typestate tracking system is a modest endeavor in and of itself, short of supporting fully featured contracts and logic. In achieving the former, Permchecker takes advantage of the language features listed in Table 1 in order to manage and track typestates.

5.1 Artifacts & Mechanisms

In Figure 7 we show how typestate information flows through the toolchain at different levels of abstraction. For instance, the dynamic binary instrumentor (DBI) intercepts function calls at run time to the typestate tracking interface. The DBI then uses Pin's API for instrumenting memory accesses, in order to check typestate permissions. The idea is that a shared universe of typestate identifiers applies across all abstraction levels. This is necessary because each level accesses and mutates the exact same pieces of memory according to specially crafted policies. These policies all require unsafe low-level access, unlike any other domain to our knowledge, in order to extract every ounce of performance from components including the JIT compiler, VM host-language, allocator, garbage collector, and heap mutator.

Many of the components listed in Figure 7 center around compilation of a C++ or C-like implementation language for the memory manager. The Permchecker toolchain most seamlessly allows for the instrumentation, compilation, and dynamic interception of typestate-related runtime events of a memory manager implemented in C++. We now discuss, in Sections 5.2 and 5.3, our experience using Permchecker components in two memory managers other than Hotspot.

Fig. 7. Architecture of the Permchecker toolchain, indicating how typestate tracking and permissions are implemented. Nodes are artifacts in the system which involve typestate, and arrows are mechanisms for enacting their purpose.

5.2 Jikes RVM

In the Jikes RVM, Permchecker aided in the detection and diagnosis of a latent memory layout bug in an uncommon configuration of the runtime system. This bug manifests when two header fields of an object are defined to be different sizes, resulting in a memory error. This error is caused by an incorrect size variable being used to compute the offset of one of the two object header fields. The bug however is benign in the default, common, system configuration because the incorrect size variable happened to be the correct value. Notably, the offset calculation and resulting errant memory access is implemented in Java.

Detecting this Java bug with the Permchecker components listed in Table 1 involves three conditions. First, a mature FFI to C++ or assembly is necessary to communicate typestate information to members of the pck namespace as compiled with Clang or GCC². We consider this burden minimal, as a FFI is an already integral part of memory manager implementations. This was the case

²Clang only required for function and class annotations.

with Jikes RVM, for which a readily extensible FFI to C++ allowed us to instrument the VM with tpestate transitions and code permissions.

The second condition is that the layout specification language’s codegen and accompanying tpestate annotations do not apply. Instead, custom VM-specific annotations for defining code permissions were implemented into the staged compilation process of Jikes RVM. Functionally these annotations are equivalent to the Clang extension for permissions. However, the presence of a separate bootstrapping Java compiler in the Jikes RVM build process complicates the detection of early memory corruption with annotations. For example, in order to achieve 100% memory coverage with tpestates a mechanism for tracking a memory location across bootstrapping stages is necessary. We leave this to future work.

The final condition arises from the fact that debugging information, such as filenames and line numbers, are not necessarily readily available to the permission-checking Pin tool in the form of DWARF debugging information. To address this, the Pin tool includes the ability to receive on the command line a mapping from instruction pointers to filenames and line numbers. While this ability is not utilized in our study of Hotspot, our experimentation with the Jikes RVM requires it to report helpful diagnostic information. As a part of Jikes RVM’s compilation process a debugging map is produced, thereby reducing the burden of enabling Permchecker’s diagnostic capabilities to mechanically translating the information into a reportable form.

Finally, with respect to Jikes RVM, the latent memory layout bug discussed earlier closely mirrors the Hotspot header layout bug causing the error reported in Figure 6. In the Jikes RVM bug, Permchecker reports a memory access to one object header field when the VM only had permission to access another (adjacent) field. Upon inspecting the recorded tpestates for surrounding memory locations it becomes apparent that the pointer access is in conflict with memory as-allocated, leading directly to the source of the bug.

5.3 dlmalloc

This section discusses the applicability of Permchecker’s dynamic tpestate checking capabilities to the implementation of a manual memory allocator.

Instrumentation Burden. In the dlmalloc we fully annotate and instrument tpestate transitions for memory related to a program’s heap as allocated through the standard malloc and free functions. The source code for unmodified dlmalloc version 2.8.6, ignoring mspace variants of the code, amounts to around 5,500 lines of code and documentation. After annotation, the source code contains an additional 600 lines of code including code permission annotations, tpestate transitions, and manual tpestate assertions.

Another 400 lines of boilerplate code consists of glue code for Permchecker, tpestate identifier definitions, and tpestate transition macros. Chronologically, we studied dlmalloc prior to developing any of the components listed in Table 1, except the last two rows comprising the dynamic checker and shadow memory. Therefore, the 400 lines of code added to dlmalloc correspond to code automated in our study of Hotspot. Of the 600 lines of code annotations and instrumentation listed, about 100 correspond to explicit calls to protect and unprotect in dlmalloc. This effort is reduced by half by the Clang-based function annotations, as applied to Hotspot.

Lessons Learned. **Annotating a legacy memory manager with tpestate permissions is a practice in pattern matching** as one spelunks through the code. For example, dlmalloc extensively uses C structs to “declare a view into memory allowing access to necessary fields at known

offsets from a given base.” [Lea 1991]³ Annotating access to such fields is a systematic refactoring-like process of annotating each field (de)reference operation.

It is ideal to test the annotation of a memory manager with tpestates gradually. While the appropriate tpestate for a function is typically readily apparent from textual context in the code, annotations are no less susceptible to typos, categorical mistakes, and missing annotations. We discovered this in our study of `dlmalloc` by testing the system after looking for instances of a code pattern (e.g. accessor functions), and annotating them appropriately.

The most common instrumentation failure mode we observed involved failing to adequately generalize an implementation pattern, resulting in missed annotations. Such failures become immediately apparent when tested by `Permchecker` in conjunction with inspection of code and memory coverage statistics. It is reassuring to recognize that these simple procedures can effectively rule out memory corruption — procedures including manual refactoring, iterative testing, and code and memory coverage inspection.

In first studying `dlmalloc` then switching to `Hotspot`, we pinpointed the language tools needed to instrument a memory manager (1) systematically, (2) gradually, and (3) progressively towards ruling out memory stomping. We then built and studied these tools, each listed in Table 1, in the context of `Hotspot`. Ongoing future work now involves developing more expressive tooling for post-hoc tpestate inspection, live (GDB-like) breakpoints, historical program tracing, and program slicing in terms of tpestates.

6 BUG INJECTIONS

In this section we present a methodology, and application thereof, for injecting bugs into a memory manager. The idea is that each injected bug is an instantiation of a template designed based on real bugs found in memory managers. Each template encapsulates the memory safety or correctness effects of the original bugs, with consistent reproducibility of many unique injectable bugs at modest scale. Having such unique bugs at scale then allows us to stress the system for a broad array of behaviors and failure modes.

6.1 Object Cloning Bug Injection

The core functionality of one bug template used to inject bugs into `Hotspot` is shown in Figure 8. The code from this figure is placed in the compaction phase of the Shenandoah garbage collector. More generally, the idea is that this particular template can be injected anywhere in the code where an object pointer (oop type) is available. The template works by delaying bug injection until some number of executions have happened. Once this happens the code injects a single non-recurring bug that will not simply crash the system outright.

To run this template, we must first decide on an initial value for `__CORRUPTION_COUNTDOWN`. This value tracks how many objects have been forwarded during concurrent compaction, and therefore changing it varies which exact objects are affected by the injected bug. By varying it over a large number of sufficiently long-running program runs, we now have numerous injected bugs each with differing effects on the downstream behavior of the system.

It should be noted that this bug template is not entirely divorced from tpestate annotations. In fact, the functions `is_forwarded`, `klass`, and `forwarder` are annotated with the appropriate access permissions, allowing ostensibly buggy code to access memory. However, it is evident these functions are used here in an appropriate, non-buggy, manner. They are all used to access header words of valid heap objects for which the tpestates match the usage of the values being read.

³Emphasis added, quotes removed.

```

1  oop p = ...;
2  if (DoInjectCorruption && InjectID == 1 && !__CORRUPTION) {
3      if (__CORRUPTION_SIZE > 0 && p->klass() == first_kls) { // 3rd case
4          HeapWord *src = (HeapWord*) p;
5          HeapWord *dst = first_forwardee;
6          Copy::aligned_conjoint_words(src, dst, size);
7          oop(dst)->init_mark_raw();
8          log_info(permchecker)("Corruption #1");
9          __CORRUPTION = 1;
10     } else if (__CORRUPTION_SIZE == 0 && // 2nd case
11               __CORRUPTION_COUNTDOWN == 0 && p->is_forwarded()) {
12         __CORRUPTION_SIZE = (size_t)p->size();
13         first = (HeapWord*) p;
14         first_kls = p->klass();
15         first_forwardee = (HeapWord*)p->forwardee();
16     } else if (__CORRUPTION_COUNTDOWN > 0) { // 1st case
17         __CORRUPTION_COUNTDOWN--; } }

```

Fig. 8. The core functionality of a template for injecting object cloning bugs into Hotspot. In the third conditional branch, we count down some number of executions of this code fragment. In the second branch, we then record an object located at pointer “p”. Finally in the first branch, we discover a second object of the same Klass type and (*incorrectly*) forward its contents to the destination of the first object. The variables `first`, `first_kls`, `first_forwardee`, `__CORRUPTION_COUNTDOWN`, `__CORRUPTION_SIZE`, `DoInjectCorruption`, `InjectID`, and `__CORRUPTION` are static class members.

It is not until the Copy operation that it becomes clear that a memory error is present. During this operation, we read the contents of a valid object from `src`. But as it turns out, we write these contents to a location which *already contains an object as well*. This fact about the tystate of the destination is already tracked by Permchecker when the destination first received a valid object, and therefore the operation is in error.

When the developer is told about this error, it might be tempting to think there is a specification error and that the developer should give the copy operation permission to overwrite allocated tystates. This option however makes no sense after minimal scrutiny, because transporting an object from one location to another should never be allowed to write to memory already in an allocated state. Instead, the developer must do one of two things. One, he can check that the destination is in the appropriate free tystate, transition it to the allocated tystates, and then perform a copy operation with permission to write to the allocated tystates. Or two, he can specify the copy operation has permission to write to the free tystate, and then transition the memory to the allocated tystates after copying.

The latter technique is preferable because it is less error-prone: the developer cannot forget to insert a tystate check, because the check for free memory is implied by a new annotation on the copy operation. The key is that with or without adding either the permission annotation or tystate transition, Permchecker reports an error. Based on this reasoning, it is generally better practice when instrumenting a memory manager to rely on automatic checks by the dynamic instrumentor than to use a technique that relies on the developer to check a tystate before changing it. Discovering this sort of best practice is one overarching goal of this work: to not only build correct systems, but to build error-resistant debugging and instrumentation techniques.

Table 2. Breakdown of bug templates injected into Hotspot, as detected with and without Permchecker. The second column group, Injections & Pck, indicate the total number of executions of the VM and how many of those were detected by Permchecker, respectively. The third column group indicates how the VM reacts in the absence of Permchecker.

Template	Injections	Pck	Ad-hoc	OS Error	Output	OOM	Benign
Object cloning	13, 812	13, 812	4, 820	0	8, 957	0	35
Use-after-free (R)	836	832	0	379	452	1	4
Use-after-free (W)	799	799	0	346	432	0	21
Overflow (padding)	8, 658	8, 658	313	0	0	0	8, 345
Overflow (no pad)	10, 026	10, 026	9364	662	0	0	0
Wasted Memory	1, 000	0	0	0	0	0	1, 000
Move Failure	976	649	562	31	2	5	376

6.2 Bug Template Results

In Table 2 we see how well Permchecker managed to detect the injection and execution of a series of bug templates, including the one just discussed in Section 6.1. A bug template when instantiated and tested with specific program inputs is an *error benchmark*. A key strategy we developed to create good error benchmarks was to confine the effects of a bug to a single execution of some operation by the memory manager. This strategy provides the following benefits:

- There is an increased chance of the bug evading traditional error detection mechanisms, mimicking the rarity of incidence exhibited by similar real bugs.
- No static analysis is necessary in order to synthesize or construct conditional statements which help rarify the bug's execution.
- The number of injections scales because the number of dynamic operations performed by the garbage collector is unbounded, unlike the number of static injection sites.

Based on this strategy, we created the bug templates listed below. For each template, the number of injections reported in Table 2 reflects the number of times we were able to get the template's preconditions to be satisfied. For example, the use-after-free templates require finding an object near the end of its containing region, and therefore required more compute power to accumulate the same number of injections as some of the other templates.

Object cloning - this bug template is the code fragment from Figure 8, which discovers two object pointers of the same Klass type, and overwrites the memory of one with the other. This template generally works in any VM function with access to at least one heap object pointer per execution.

Use-after-free (read) - this bug template discovers a single object pointer near the end of a region, waits for its contents to be evacuated, and then redirects the application's next field access to use the freed memory as a base pointer in its load instruction.

Use-after-free (write) - like above, but the field access is a write instruction causing garbage to be "corrupted" and the intended object field to not be updated.

Overflow (padding) - this bug template duplicates the Klass metadata for a single object instance when it gets allocated, modifying the duplicate Klass to indicate this one object has some multiple of an object's alignment less of memory than it requires based on the fields in the object.

Overflow (no padding) - like above, but the injection only happens when the last field of the object ends on an object alignment boundary.

Wasted Memory - this bug template operates similarly to the overflow templates, but the Klass is modified to indicate the object consumes *more* memory than it requires.

Move Failure - this bug template discovers an allocated object residing at the end of a memory region and skips compacting it, possibly causing other object(s) to point to a valid-looking object that the memory manager believes to be free memory that can be allocated into.

6.3 Choice of User-Level Application

One kind of memory bug we wish to study is one that does not crash the program at all. Instead, it affects the correctness of the user application such as program output. The incidence of output-based errors relies heavily on the domain and specifics of the user application. Characteristics affecting whether or not an output error will arise include object allocation patterns, drag time between an object's last use and when it becomes unreachable, and more. For the results in Table 2, we modified a version of GCBench [Ellis et al. 2014] to maximize fragility by minimizing drag.

Fragility, here, means how likely an application is to compute the wrong result if one of its objects gets corrupted. The original unmodified version of GCBench is a prime example of a non-fragile benchmark. GCBench by default iteratively constructs a number of binary trees of certain depths, in order to stress the *performance* characteristics of a garbage collector. The benchmark is not actually intended to compute anything of algorithmic value. Therefore, so long as the application terminates at the end of main, it “computed” the correct value.

In contrast our modified version of GCBench maximizes fragility. It does this by computing a hash over the contents of each node in a binary tree shortly before expecting the tree to become garbage. This hash iteratively accumulates over each node allocated. The hash is printed upon program completion. This strategy aggressively links correctness of application output to the presence of heap corruption. This strategy is further ripe for automation over any existing application in domains other than just performance analysis.

6.4 Analysis & Effectiveness

With the use-after-free (write) template, both pointer updates and primitive field updates were affected throughout the various error benchmarks. As a result a small number of the benchmarks terminated significantly early (with incorrect program output), presumably because a failed pointer initialization lopped-off a significant portion of a recursive data structure before it was processed by the application. The same behavior appears to have happened with the use-after-free (read) template, where some of the incorrect-output benchmarks terminated early because a load of a pointer to a large recursive sub-structure was redirected to a smaller one.

Of the 4 benign results with the use-after-free (read) template, all of them went undetected by Permchecker. This means that the freed memory was necessarily quickly reallocated by the memory manager, bringing it to a typestate validly accessible by the application mutator. Subsequently the application accessed a field at the wrong address, but happened to read the correct value so-as to be benign. We believe these 4 injections have to do with leaves of a binary tree being implemented as the NULL value. As a result a memory access to the left or right child of a node was corrupted, but would have received the NULL value it was supposed to read anyways.

This behavior exposes a limit to Permchecker's checking capability. Permchecker cannot detect a dangling pointer error when the memory pointed to by the dangling pointer is reallocated to a similarly allocated object of the same typestate and size. Valgrind's Memcheck tool is similarly unable to detect this subclass of dangling pointers, for a fundamentally identical reason. Neither tool checks how a pointer value is obtained, just that an observed access to the referenced memory is permissible in each tool's checking model.

6.4.1 Observable Behaviors. In Table 2 we break down the frequency with which each bug template causes a variety of observable behaviors in the absence of Permchecker. These behaviors are a

holistic collection of failure modes that a developer reasons about while debugging. The following list defines each failure mode:

- **Ad-hoc:** a VM assertion error, usually from a suspicious looking pointer or value.
- **OS Error:** a segmentation fault, bus error, or other OS error even if the VM intercepts it.
- **OOM:** an out-of-memory error reported by the VM.
- **Output:** the program successfully terminates with incorrect user-level textual output.
- **Benign:** the program successfully terminates with correct user-level textual output.
- **Deadlock:** corruption causing non-termination. This was not observed in testing.

We consider “benign” program behavior to be *bad* when a corruption is known to have occurred. In fact, all the benign bugs from Table 2 can rightly be considered failure modes – in each case the system does not exhibit signs of known corruption. The “Wasted Memory” benchmark, too, corrupts the layout of memory. This corruption occurs benignly in a way that the Permchecker annotations were unable to detect because no read or write instructions access the “extra” memory.

6.4.2 Utility of Error Reports. It is desirable for the contents of an error reported by Permchecker to be diagnostically “effective.” Subjectively, an effective error report is one which identifies information pertinent to the source of the error. With typestate checking, a simple first-order metric is whether or not either the permissible or the observed typestates, of a Permchecker error report, indicates the memory (or layout) which was actually corrupted.

In all of the Permchecker error reports pertaining to Table 2, each report included the pertinent (corrupted) typestate. For instance in the “Overflow (padding)” template, Permchecker always reports an error where the program had permission to access an object `Field`, but in fact accesses a `Padding` typestate. Similarly, with the “Overflow (no pad)” template, Permchecker always reports an error involving access to a `MarkWord` but observes a `Field` access, or vice-versa.

The takeaway is that typestate checking improves sensitivity over the existing ad-hoc error checks. Checking also exhibits increased diagnostic value by reporting an observed type, in-lieu of just an expected type implied by the stack trace of an ad-hoc check. In essence, Permchecker is capable of meaningfully diagnosing a generalized class of *initialization* bugs.

An initialization bug is a bug in which a resulting program error is traceable to a root cause based on the most recent component of the system which (1) initialized the memory to its observed state, or (2) should have initialized it to its expected state. Part (1) indicates Permchecker’s bug detection capabilities are similar to that of a taint analysis. Part (2) indicates Permchecker’s ability to detect some control-flow bugs, e.g. a missing deallocation call.

Notably missing from these two kinds of bugs are *value* and *logic* bugs. For example, a garbage collection (GC) algorithm failing to track all GC roots leads to a premature free. Permchecker then detects this situation as a memory access by the application to purportedly free memory. This safety manifestation of a decidedly logic-based GC bug illuminates the difference between diagnostic value and pure detection – Permchecker reports less effective information for value and logic bugs than it does for initialization bugs.

6.4.3 Latency of Error Reports. Prior to writing bug templates and testing their impact on the behavior of the VM, we had little basis for any kind of insight into how quickly typestate annotations would be able to detect a memory safety error compared to existing ad-hoc checks in the VM. In one formulation, we believed existing and often value-based sanity checks in the VM might excel at proactively preventing bad values from being operated over. In contrast, Permchecker’s typestate annotations might excel at detecting bad operations once they happen, after a bad value has already been created computationally rather than loaded from memory.

For the bug templates we chose, this formulation was entirely not the case. In all the injections from Table 2, **Permchecker detects an error accounted for in the third column prior to any corresponding error accounted for in the fourth (ad-hoc) column**. Methodologically this ordering is determined by running the VM with *both* forms of checks enabled, but where the program simply logs the error reports rather than terminating immediately. This methodology has the benefit of obviating any need to create strictly reproducible thread schedules for the VM. Furthermore it is valid to compare program behaviors this way because program checkers have ostensibly no side-effects or impact on values in the VM. The tradeoff is that this methodology cannot compare one-to-one results of Permchecker with any similarly purposed, but non-composable with Permchecker, dynamic instrumentation tools.

7 MOTIVATING OBJECT MONITOR BUG RESULTS

We now discuss the results of debugging and diagnosing the deadlock bug, as motivated at the beginning of Section 1. This discussion necessarily benefits from hindsight, as the bug was in fact diagnosed and fixed without the aid of Permchecker through manual instrumentation of Hotspot's code which manages object monitors. An object monitor is a locking mechanism which prevents two or more threads from executing Java methods of the same object instance at the same time.

7.1 Annotation Considerations

By instrumenting object monitor code with ad-hoc log statements, the original debugging developer ultimately observed an erroneous lock operation. This erroneous operation is unique because it is memory corruption. Memory corruption typically involves an access to memory containing data of the wrong type. In this object monitor bug, however, this was not the case. The root cause occurs in a function checking the boundedness of a memory location on the Java stack. This function incorrectly reports malloc'd memory adjacent to the stack as being part of the stack. Thus, a pointer to memory of the *correct* type (object monitor) is corrupted by code *allowed* to access it.

In applying Permchecker annotations to Hotspot code related to object monitors, we observe some benefits as well as drawbacks to tpestate annotations. First, a permission annotation on a function is applicable to *all* callers of the function. As a result, a (calling) context sensitive tpestate property is in tension with code reuse – tpestate annotations must invariably be placed higher and higher in the call stack, leading to more and more annotations at individual call-sites.

With respect to memory layout oriented tpestates, however, there exist actionable and systemizable mitigations to the code smell of widespread annotations. In the case of Hotspot's object monitor code, we mitigate this by specializing multiple instances of a `ObjectMonitor C++` class for each allocation context (or hierarchy). The idea here is that the specialization of the existing (object monitor) functionality maintains functional equivalence and therefore the modification cannot by itself introduce (non-Heisenbug) program errors.

7.2 Object Monitor Refactoring and Debugging Process

The allocation contexts for `ObjectMonitor` include certain C++ runtime system code managed by the C++ compiler, and compiled application code compiled by the Java compiler. An object monitor allocated for the former purpose is allocated into the malloc heap. In contrast, an object monitor allocated for the latter purpose is allocated onto the Java stack. By first separating out call-sites of accessors based on these two allocation contexts, Permchecker is then able to correctly validate layout safety in each context. The idea here is that this refactoring process is (1) minimally invasive with respect to functionality, (2) adds otherwise unknown information to the implementation / artifact, and (3) systematically applicable to all allocation hierarchies. The application of object monitor tpestates to Hotspot went as follows:

- We identify allocation sites of object monitors, and annotate them with typestate transitions. These sites include ones in C++ as well as generated assembly from the Java compiler. A grep for the `ObjectMonitor` class was a good first step here.
- We refactor the `ObjectMonitor` class to specialize two versions of it with annotations.
- We propagate refactoring up to call sites for each of the allocation contexts: stack & malloc.
- Finally, we validate the memory layout by running preliminary test applications which exercise the object monitor related runtime system code.

In the last step above, Permchecker reports as a violation any (1) access patterns we missed or applied annotations to incorrectly, and (2) true memory layout errors. A missed access pattern may appear to Permchecker as an invalid memory access. A true memory layout error for the object monitor deadlock bug appears as a call site to an object monitor accessor function intending to lock a monitor on the Java stack, but is observed by Permchecker to be on the malloc heap.

7.3 Boundedness and Diagnosability

The most applicable class of properties also validated by existing memory checkers is *boundedness*. Boundedness is a property which helps define whether or not a pointer references memory within the bounds of a specified chunk, or class of chunks, of memory. The technique of modeling object ownership [Weiss et al. 2019] places annotations governing boundedness into code on a per-variable (per memory chunk) basis. This high level of per-chunk specificity achieves local integrity at the cost of global safety checking (in a memory manager).

The Memcheck [Nethercote and Seward 2007] tool validates boundedness by partitioning addresses into unallocated and allocated⁴ classes (of chunks of memory). This per-class level of specificity achieves global integrity, such as an absence of certain use-after-free errors. However, low per-chunk specificity rules out the detection of unspecified boundedness properties.

8 SCOPE, LIMITATIONS, & FUTURE WORK

Threats to Validity. Permchecker targets the debugging and diagnosis of the most difficult to diagnose bugs in a memory manager. While we discussed limitations to Permchecker’s ability to diagnose value and logic bugs, those bugs remain detectable when they manifest as corruption. Bugs which do not do so, or cannot be reproduced in the first place, go undetected.

Permchecker’s dynamic instrumentation, in particular, can perturb a runtime system such that JIT compilation performance statistics, thread-scheduling, and the timing of garbage collections prevent a rare bug from manifesting. As a result, there exist memory corruption bugs which Permchecker will *never* be able to observe. This is especially problematic because the domain of memory manager implementations fall squarely into the class of programs which are highly sensitive to timing considerations. In this regard, typestate checking of memory managers will greatly benefit from the development of zero-overhead dedicated hardware to support live debugging.

Overhead. Permchecker’s dynamic instrumentor incurs a 10x to 50x slowdown in the tests of bug templates listed in Section 6.2. This observed slowdown primarily owes itself to the instrumentation of each and every memory access by the VM: stack, Java heap, and C++ metadata alike. At each memory access, the most common fast path of the shadow-memory implementation performs two or three additional memory accesses. This behavior leaves substantial room for improvement, either in smarter data structures or, optimally, hardware support similar in nature to page tables.

Thread Safety. Thread safety in the Java Native Interface (JNI / FFI), JIT compiler, and mutator slow paths of Hotspot are entirely handled by existing VM code. Thus the only place in the toolchain

⁴And initialization status.

where thread-safety became a non-trivial concern was in the implementation of shadow memory. As such we implemented shadow memory as a lock-free tree-like map data structure in which map updates and map allocations are atomic. In the presence of a data race involving an update and a read access to the tpestate of the same memory location, the shadow memory provides no guarantee as to which tpestate the read access observes: the old one or the new one.

First-Order Permissions. Permchecker is limited in scope to checking access permissions over a fixed set of uniquely identifiable tpestates. This scope limits our ability to check the validity of broader algorithmic contracts or predicates. For example, this predicate: “*the second write to some Field at address α should be preceded by memory of tpestate MarkWord containing the value 0b00*”. Such a check would require both a mechanism by which to communicate the desirable property to the dynamic instrumentor, as well as efficient algorithms and data structures to track necessary information and verify validity.

Other debugging tools, notably GDB, supports features like conditional watchpoints and step-wise debugging. These features allow a developer to know when the program accesses a particular memory address, and to inspect subsequent instructions. Similarly, there is value in knowing when the memory manager accesses a particular tpestate of memory. In this work we focused on building a toolchain to diagnose reproducible errors, but a natural future extension could involve more traditional debugging features capable of inspecting tpestate at run time.

ACKNOWLEDGMENTS

This work was funded by NSF Grant #1717373.

REFERENCES

- Karl Cronburg and Samuel Z. Guyer. 2019. Floorplan: Spatial Layout in Memory Management Systems. In *Proceedings of the 18th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences (Athens, Greece) (GPCE 2019)*. Association for Computing Machinery, New York, NY, USA, 81–93. <https://doi.org/10.1145/3357765.3359519>
- B. Dolan-Gavitt, P. Hulin, E. Kirda, T. Leek, A. Mambretti, W. Robertson, F. Ulrich, and R. Whelan. 2016. LAVA: Large-Scale Automated Vulnerability Addition. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Jose, CA, 110–121. <https://doi.org/10.1109/SP.2016.15>
- John Ellis, Pete Kovac, and Hans Boehm. 2014. https://hboehm.info/gc/gc_bench/ Accessed: 2021-04-13.
- Jason Evans. 2006. A Scalable Concurrent malloc(3) Implementation for FreeBSD.
- IBM. 2005. Jikes RVM. <http://www.jikesrvm.org/> Accessed: 2018-09-28.
- Doug Lea. 1991. A Memory Allocator. <http://g.oswego.edu/dl/html/malloc.html> Accessed: 2021-04-13.
- Nicholas Nethercote and Julian Seward. 2007. Valgrind: A Framework for Heavyweight Dynamic Binary Instrumentation. In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation (San Diego, California, USA) (PLDI '07)*. ACM, New York, NY, USA, 89–100. <https://doi.org/10.1145/1250734.1250746>
- Robert O’Callahan, Chris Jones, Nathan Froyd, Kyle Huey, Albert Noll, and Nimrod Partush. 2017. Engineering Record And Replay For Deployability: Extended Technical Report. CoRR abs/1705.05937 (2017), 20 pages. arXiv:1705.05937 <http://arxiv.org/abs/1705.05937>
- Oracle. 2006. OpenJDK Hotspot Division. <http://openjdk.java.net/groups/hotspot/> Accessed: 2021-04-13.
- Andrew Rice, Edward Aftandilian, Ciera Jaspán, Emily Johnston, Michael Pradel, and Yulissa Arroyo-Paredes. 2017. Detecting Argument Selection Defects. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 104 (Oct. 2017), 22 pages. <https://doi.org/10.1145/3133928>
- Subhajt Roy, Awanish Pandey, Brendan Dolan-Gavitt, and Yu Hu. 2018. Bug Synthesis: Challenging Bug-Finding Tools with Deep Faults. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Lake Buena Vista, FL, USA) (ESEC/FSE 2018)*. Association for Computing Machinery, New York, NY, USA, 224–234. <https://doi.org/10.1145/3236024.3236084>
- Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitriy Vyukov. 2012. AddressSanitizer: A Fast Address Sanity Checker. In *2012 USENIX Annual Technical Conference (USENIX ATC 12)*. USENIX Association, Boston, MA, 309–318. <https://www.usenix.org/conference/atc12/technical-sessions/presentation/serebryany>

- R. E. Strom and S. Yemini. 1986. Tpestate: A programming language concept for enhancing software reliability. *IEEE Transactions on Software Engineering* SE-12, 1 (1986), 157–171.
- Michael Vollmer, Chaitanya Koparkar, Mike Rainey, Laith Sakka, Milind Kulkarni, and Ryan R. Newton. 2019. LoCal: A Language for Programs Operating on Serialized Data. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) (*PLDI 2019*). Association for Computing Machinery, New York, NY, USA, 48–62. <https://doi.org/10.1145/3314221.3314631>
- Aaron Weiss, Daniel Patterson, Nicholas D. Matsakis, and A. Ahmed. 2019. Oxide: The Essence of Rust. *ArXiv abs/1903.00982* (2019).