QVM: An Efficient Runtime for Detecting Defects in Deployed Systems

MATTHEW ARNOLD and MARTIN VECHEV, IBM Research
ERAN YAHAV, Technion and IBM Research

2

Coping with software defects that occur in the post-deployment stage is a challenging problem: bugs may occur only when the system uses a specific configuration and only under certain usage scenarios. Nevertheless, halting production systems until the bug is tracked and fixed is often impossible. Thus, developers have to try to reproduce the bug in laboratory conditions. Often, the reproduction of the bug takes most of the debugging effort.

In this paper we suggest an approach to address this problem by using a specialized runtime environment called *Quality Virtual Machine* (QVM). QVM efficiently detects defects by continuously monitoring the execution of the application in a production setting. QVM enables the efficient checking of violations of user-specified correctness properties, that is, typestate safety properties, Java assertions, and heap properties pertaining to ownership. QVM is markedly different from existing techniques for continuous monitoring by using a novel overhead manager which enforces a user-specified overhead budget for quality checks. Existing tools for error detection in the field usually disrupt the operation of the deployed system. QVM, on the other hand, provides a balanced trade-off between the cost of the monitoring process and the maintenance of sufficient accuracy for detecting defects. Specifically, the overhead cost of using QVM instead of a standard JVM, is low enough to be acceptable in production environments.

We implemented QVM on top of IBM's J9 Java Virtual Machine and used it to detect and fix various errors in real-world applications.

Categories and Subject Descriptors: D.2.5 [**Testing and Debugging**]: Diagnostics; D.2.4 [**Software/ Program Verification**]: Reliability; D.3.4 [**Programming Languages**]: Processors—*Runtime environments*

General Terms: Reliability

Additional Key Words and Phrases: Virtual machines, typestate, heap assertions, diagnosis, debugging

ACM Reference Format:
Arnold, M., Vechev, M., and Yahav, E. 2011. QVM: An efficient runtime for detecting defects in deployed systems. ACM Trans. Softw. Eng. Methodol. 21, 1, Article 2 (December 2011), 35 pages.
DOI = 10.1145/2063239.2063241 http://doi.acm.org/10.1145/2063239.2063241

1. INTRODUCTION

Despite increasing efforts and success in identifying and fixing software defects early in the development lifecycle, some defects inevitably make their way into production. The wide variety of deployment configurations and the diversity of usage scenarios is almost a certain guarantee that any large system will exhibit defects after it has been deployed.

M. Vechev is also affiliated with ETH Zurich.
A preliminary version of this article appeared in Proceedings of the 23rd Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications as Arnold et al. [2008].
Author's address: M. Arnold, M. Vechev, and E. Yahav; email: marnold@us.ibm.com; martin.vechev@inf.ethz.ch; yahave@cs.technion.ac.il.

DOI 10.1145/2063239.2063241 http://doi.acm.org/10.1145/2063239.2063241

Detecting and diagnosing defects in a production environment remains a significant challenge. Failures in such environments might occur with low frequency and be virtually impossible to reproduce. For example, a defect might occur due to a specific concurrent interleaving, a specific lengthy user interaction, or a slow resource leak that gradually degrades system performance leading to an eventual crash.

Existing tools for diagnosing defects "in the wild" are limited and usually incur an unacceptable overhead that significantly disrupts the operation of the deployed system. On the other hand, reproducing the failure in a test environment (if at all possible) may require considerable time and effort.

One way to detect rarely occurring defects is to continuously monitor a system for violations of specified correctness properties. For example, this can be achieved by using global property monitors and local assertions. However, the typical cost of these techniques prevents programmers from widely using them in production environments.

This work describes a runtime environment that is able to detect and help diagnose defects in deployed systems. Towards this end, we present the Quality Virtual Machine (QVM), a runtime environment that uses the technology and infrastructure available in a virtual machine to improve software quality. QVM provides an interface that allows software monitoring clients to be executed with a controlled overhead. Based on this interface, we present three such clients that continuously monitor application correctness by using a combination of simple global property monitors (typestate properties) and assertions. In addition, QVM automatically collects debug information that enables effective defect diagnosis.

We implemented QVM on top of IBM's J9 Java Virtual Machine. We used a number of large-scale real-world applications with QVM and found defects in many of them. We explain the design rationale behind QVM in Section 3.1.

1.1 Main Contributions

The contributions of this paper include the following.

— QVM is a runtime environment targeted towards defect detection and diagnosis in production systems.
— A novel overhead manager enforces an overhead budget on client analyses, while maintaining sufficient accuracy for detecting defects.
— We introduce property-guided sampling and in particular *object-centric* sampling to collect sampled profiles while preserving correctness of the analysis.
— A lightweight interface helps separate analysis clients from the details of the underlying VM, and transparently manages overhead of these clients.
— We use this infrastructure to implement three representative analysis clients: (i) tracking simple temporal safety properties and providing debug information; (ii) checking standard Java assertions; (iii) checking expressive heap queries pertaining to object ownership.
— We implemented QVM on top of IBM's production Java Virtual Machine (J9). We used QVM as our standard day to day virtual machine, running a wide range of applications without a noticeable slowdown. We show that QVM can be used to effectively detect defects in such applications, and help diagnose them. In addition, we evaluated the overhead on the standard SPECjvm98 and DaCapo benchmarks [Blackburn et al. 2006].

1.2 Overview

In this section we provide a brief informal overview of QVM components and our experimental evaluation.

Overhead Manager. QVM allows the user to specify an overhead that is considered acceptable for the current monitoring environment. The maximum acceptable overhead may be 5%–10% in a live deployed system, yet 100% overhead (factor of 2 slowdown) may be considered acceptable in a testing environment. Given an overhead budget, the QVM strives to collect as much useful information as possible from the executing program while staying within the specified budget.

QVM Interface (QVMI). A performance-aware profiling/monitoring interface that allows client analyses to remain decoupled from the VM, while maintaining efficiency. The design goal of this component is to enable development of powerful, yet efficient dynamic analyses. Technically, the overhead manager and the QVMI work together to provide clients with a transparent adaptive overhead management.

Analysis Clients. Using the QVM platform, we implement three analysis clients:

Typestate Properties. This analysis client enables the dynamic checking of typestate properties. Dynamic checking of typestate properties, as well as generalized multiple-object typestate, has been addressed before in Tracematches [Allan et al. 2005] and MOP [Chen and Roşu 2007]. We use the typestate client to demonstrate three contributions of our platform: (i) adaptive overhead management; (ii) collection of timing information for typestate transitions; (iii) collection of additional detailed debug information with low overhead.

Local Assertions. QVM allows efficient sampling of user assertions by intercepting standard Java assertions and managing their execution through the overhead manager.

Heap Probes and Operations. QVM enables the dynamic checking of various global heap properties such as object-sharing, ownership, thread-ownership and reachability. These properties are useful for both debugging and program understanding purposes.

Experimental Evaluation. To evaluate the usability of QVM in finding defects and diagnosing them, we focused on typestate properties that correspond to resource leaks. For that purpose, we set QVM as the default JVM used in our environment and used it to perform all of our daily tasks while recording its error reports. To further exercise QVM, we used a wide range of applications on a regular basis. Some of the applications considered are an instant-messenger (goim), newsfeed readers (feednread, rssowl), file management utilities (virgoftp, jcommander), large IBM internal applications, etc. For all of these applications, the overhead incurred by running them on top of QVM was unnoticeable to the user. Since the overhead of QVM is not noticeable by the user while using interactive applications, we use the SPECjvm98 and DaCapo benchmarks to evaluate the overhead manager's effectiveness.

In some of our experiments (e.g., Azureus, virgoftp, goim), we investigated each report manually, diagnosed the causes of the errors, and implemented fixes. For some applications, our defect reports were confirmed by the development team, and our fixes were incorporated into the codebase.

Outline. The rest of this article is organized as follows. Section 2 motivates our approach with a simple example. Section 3 provides a brief overview of the QVM architecture and the QVM Interface (QVMI). Section 4 describes the details of the overhead manager, which together with QVMI provides the foundations for client analyses. Section 5 describes our client analyses, and Section 6 provides details of

```
QVM ERROR:[Resource_not_disposed] object [0x98837030]
of class [org/eclipse/swt/graphics/Image]
allocated at site ID 2742 in method
[com/aelitis/azureus/.../ListView.handleResize(Z)V]
died in state [UNDISPOSED]
with last QVM method [org/.../Image.isDisposed()Z]
```

Fig. 1. A sample QVM error report for Azureus.

their implementation. Section 7 describes the results obtained running application on top of QVM. Section 8 discusses related work.

2. MOTIVATING EXAMPLE

Azureus[1] is an open-source implementation of the BitTorrent protocol. It supports several modes of user interaction, all implemented using the Standard Widget Toolkit (SWT)[2] Azureus is the #1 downloaded Java program from SourceForge, and has more than 360 million downloads to date. Azureus plays the role of both a client and a server for P2P file sharing, and is therefore a relatively long-running application.

Finding Bugs. We run Azureus with QVM, monitoring various correctness properties, including possible SWT resource leaks and IOStream leaks. Azureus runs on QVM with no apparent slowdown. Over the course of few hours, we check the QVM logs and observe that some errors were reported.

Figure 1 shows an example of an error reported by QVM while running Azureus. This is the actual error report as produced by QVM where some package names have been abbreviated. By itself, this error report provides useful information about the property being violated. In this case, the reported `Image` object has not been properly disposed before it became unreachable. Failure to properly dispose such SWT resources leads to leakage of OS-level resources and may gradually hinder performance and even lead to a system crash. The error report of Figure 1 provides the basic information necessary to track down the error: the method in which the object was allocated, the object's last state, and the last method invoked on the object.

Generally, the only user specification required for QVM to report this kind of errors is a typestate specification like the one shown later in Figure 7. QVM has built-in specifications for detecting resource leaks of SWT resources, as well as other resources managed by standard Java libraries.

Diagnosing the Cause. The QVM error report above notifies the user that there is an error, but understanding the cause of the error and introducing a fix is still nontrivial. The programmer needs to track the flow of the object through the program to identify why dispose was not called. To assist the programmer in this task, QVM provides additional, more detailed, debug information in the form of a *typestate history*. A typestate history for an object shows all the methods that have been invoked with that object as a receiver, over the course of the object's lifetime, from allocation to collection. For every method invocation, the invocation history collects the contexts in which it was invoked. (We provide a more elaborate description of the typestate history in Section 5.1.)

To maintain a low runtime overhead, a typestate history is only collected for *some* of the tracked objects. Whenever an allocation site is identified as allocating a number of objects that violate a property, QVM starts recording typestate histories for a sample of objects allocated at that site. This object-centric sampling is one of the features that makes it possible to collect detailed debug information with low overhead.

[1]Azureus: Java BitTorrent Client. http://azureus.sourceforge.net/.
[2]Eclipse. Standard widget toolkit (swt). http://www.eclipse.org/swt/.

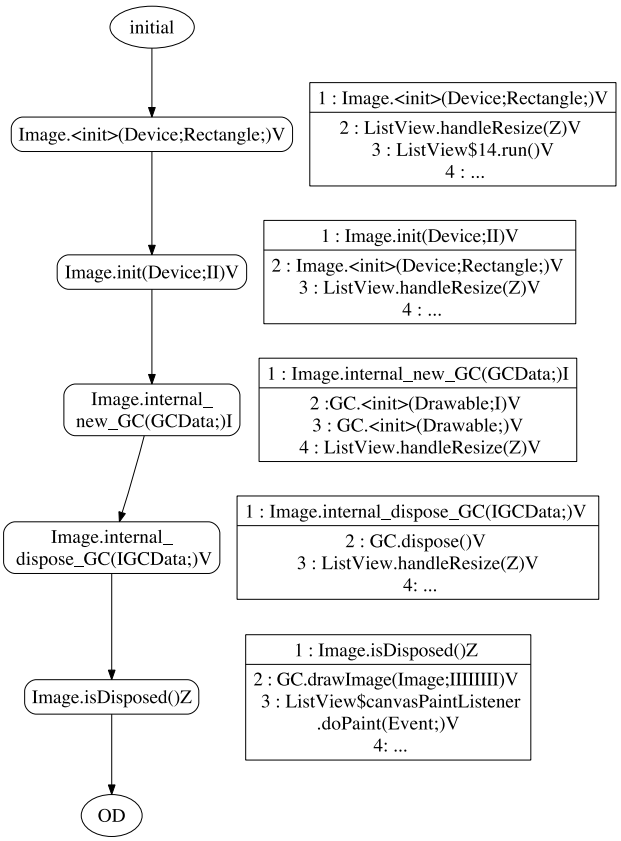

Fig. 2. Sample typestate history for a single instance of `Image` that was reported as non-disposed in Figure 1. The figure only shows a single sample stack trace for every method invoked on the object.

Figure 2 shows an example of a typestate history for an object allocated at the same site as the object reported in Figure 1. The typestate history abstracts the history of methods invoked on the object. Technically, the typestate history is a directed graph with labeled nodes and labeled edges. A node in the graph represents the state of the object after a specific method has been invoked on it. There is a single node in the graph for each method invoked on the object (summarizing all invocations of that method). A node in the graph is labeled by the name of the invoked method, and by a set of (bounded) contexts, representing the context in which the method was invoked. An edge between nodes m_1 and m_2 in the graph represents the fact that the method corresponding to m_2 has been invoked immediately after the method corresponding to m_1 has been invoked. Note that this directed edge only denotes the order in time between the two methods. It does not say that m_2 is called from m_1.

Next, we show how we used the debug information provided by QVM to find the cause of an error. In the example of Figure 2, there are 5 methods that have been invoked on the tracked object. First, the object is initialized by invoking `<init>` and `init`. Then, a graphical context (`GC`) is created around the image (`internal_new_GC`) and disposed (`internal_dispose_GC`). Finally, `isDisposed` is invoked on the image. The method `Image.dispose()` that is required for properly disposing the image is never invoked.

In this simple example, there is only one context in which each method has been invoked. The context is shown inside a rectangle next to its corresponding graph

```
class ListView extends ... {
 private Image imgView = null; // ...
 protected void handleResize(boolean bForce) { // ...
  if (imgView == null || bForce) {
   imgView = new Image(listCanvas.getDisplay(), clientArea);
   lastBounds = new Rectangle(0, 0, 0, 0);
   bNeedsRefresh = true;
 } else {
  // ...
 }
  //   ...
 }
}
```

Fig. 3. Azureus code fragment leaking SWT Image objects.

node. Considering the contexts in which the methods in this example were invoked, we can see that most of the operations on the tracked object are performed through the handleResize method in which it was allocated. The only exception is the call to isDisposed() which originates in a paint event of the list view.

We therefore focus our attention on the handleResize method in azureus.ui.swt.views.list.ListView. The typestate history serves as a guide to the execution in which the property was violated. Following the sequence of calls in the debug information we further focus attention to the code excerpt shown in Figure 3.

The problem in this method represents a common source of leaks: a new image is stored into the field imgView without properly disposing the previous image that was stored in the field. In this example, handleResize mixes the case of imgView == null (no previous image is known for taking previous bounds) with the case of forced resize (bForce == true). As a result, there are cases in which a new Image is created without properly disposing the previous Image stored in imgView.

The number of Image objects leaked as a result of this bug directly depends on user interaction. Since this leak is associated with a resize event, it may not occur with high frequency. However, the cumulative effect of a large number of small leaks may be fatal. In Section 7.1, we discuss additional problems found on Azureus by QVM, and show that some of these occur very frequently and result in significant resource leaks.

Developing a Fix. Now that we have diagnosed the bug as being caused by not disposing the old Image object stored in imgView, the question is how do we introduce a fix. What we would like to do is to invoke dispose on the object stored in imgView before we stored the newly allocated image into the field. Unfortunately, we do not know the source of the Image stored in imgView, and in particular, whether this image is *shared* with other GUI components. In SWT, it is common for resources such as images, fonts, and colors to be shared between multiple GUI components. The convention is that whoever allocates the resource is responsible for its safe disposal. When we reach the point of allocating a new Image and storing it into imgView, we do not know whether the previous value of imgView was allocated in this method. Furthermore, we do not know whether other GUI components are still using the image.

At this point, we leverage QVM's heap assertions and check that the object pointed to by imgView is not shared (i.e., does not have any references other than imgView pointing to it). We introduce disposal code preceded by an assertion that makes sure that we are not disposing a shared resource. (The disposal of a shared resource might end up crashing the application at a later point when the user takes an action that uses the resource.) The modified handleResize method is shown in Figure 4.

```
protected void handleResize(boolean bForce) { // ...
 if (imgView == null || bForce) {
  assert(!QVM.isShared(imgView));
  if(imgView != null && !imgView.isDisposed())
   imgView.dispose();
  imgView = new Image(listCanvas.getDisplay(), clientArea);
  lastBounds = new Rectangle(0, 0, 0, 0);
  bNeedsRefresh = true;
 } else {
  // ...
 }
 //  ...
}
```

Fig. 4. A fix to the `Image` leak in `handleResize` of Figure 3.

The QVM heap assertion is not a required part of the fix. The intention is to use it during the testing period of the fix and remove it before deployment. This will allow the fixed code to execute on any JVM, and not only on QVM.

We now run the fixed version of this method with QVM for a few weeks, and observe that the previously reported leak does not occur. Our assertion also makes sure that the disposal of the `Image` does not affect any other GUI component.

We reported this leak and its fix, as well as other problems mentioned in Section 7, to the Azureus development team. The problems were confirmed as real bugs, and our suggested fixes were incorporated into the project's codebase.

We provide an elaborate description of some of the bugs we found in our benchmarks in an online supplement to this article, available at: `http://www.research.ibm.com/qvm/papers/qvm-notes.htm`.

3. QVM PLATFORM

In this section we describe the QVM platform. First, we provide some background and design rationale, then we briefly describe the overall QVM architecture and its main components. Finally, in Section 3.3, we describe the QVM interface (QVMI).

3.1 Design Rationale: Modifying a VM

Today's production-grade virtual machines employ sophisticated techniques and optimizations to achieve maximal application performance. In contrast, there is little support for application correctness in a production environment besides checking low-level properties such as absence of null dereferences and array index bounds. While rich in functionality, current debug and monitoring interfaces (e.g., JVMTI) are also not applicable as they incur a slowdown that is unacceptable in production mode.

The goal of this work is to extend a production-grade virtual machine to provide software-quality services while maintaining competitive performance. We would like a solution to provide: (i) high performance and low overhead; (ii) maximal separation of analysis clients from the details of the underlying VM.

There is an apparent tension between requirement (i) and (ii). We resolve this tension by providing a generic interface (QVMI) that manages functionality common to all analysis clients, but in addition, we allow clients to cut through abstraction layers and use other VM services when appropriate.

QVMI is a development that could serve as the next generation JVMTI interface, currently provided by modern day virtual machines. We believe that QVMI can become the de-facto standard for communication between the JVM and the native client code.

However, our technique still requires VM modifications, and modifying a production-grade virtual machine is a nontrivial task. A virtual machine is a large, complex system. Moreover, implementing the quality services inside a specific VM

makes them nonportable and ties users of the system to the specific VM version. In contrast, using pure bytecode instrumentation at the language level or a standard profiling interface such as JVMTI[3] is portable across virtual machines.

Despite these disadvantages, there are a number of advantages in having at least part of an analysis reside within a production VM, as we describe below.

VM-only information. Having access to the runtime allows the client analyses to utilize information that is not readily available at the language level. For example, analyses can use free bits in object headers, directly examine the heap, quickly access structures like thread-local storage, re-use existing VM code (such as garbage collection heap traversal logic) to perform a slightly different functionality. Analyses can utilize low-level profile data and infrastructure, such as hardware performance monitors (HPM) and fine-granularity timing (for example, see overhead monitor in Section 4).

Performance. Having access to the dynamic optimizer (JIT) ensures that the critical code paths are well optimized. The JIT can also use advanced optimization techniques for fast and slow paths (thin guards [Arnold and Ryder 2002], code patching [Suganuma et al. 2001], full duplication [Arnold and Ryder 2001], etc.). The system can also make use of profile data already collected by the VM to optimize and tune a dynamic analysis.

Dynamic updating. By using advanced techniques such as code patching and on-stack replacement (OSR) [Fink and Qian 2003], VMs can support efficient dynamic updating of instrumentation during an application run.

Deployment. The deployment process becomes trivial because the required features become as ubiquitous as the VM. There is no need to "install" an analysis (recompile the program source to add instrumentation, etc.) which is particularly difficult for large production applications that might make heavy use of custom class loaders. Our analysis can be run by simply enabling a command line flag on the VM.

In the next section, we provide an overview of the QVM architecture and show how we hide the complexity of the underlying VM from most analysis clients by using the generic QVMI interface.

3.2 QVM Architecture

Figure 5 shows the overall architecture of QVM. At a high level, the QVM extends the VM execution engine with three main components:

(1) *QVM Interface (QVMI).* A performance-aware profiling/monitoring interface that allows client analyses to remain decoupled from the VM, while maintaining efficiency. The design goal of this component is to enable quick and easy development of powerful, yet efficient dynamic analyses. QVMI is described in Section 3.3.
(2) *Overhead Manager (OHM).* The overhead control system enables users to bound the overhead incurred by QVM clients. The system does fine-grained monitoring of the time spent in the clients and adapts the sampling to stay near or below overhead bounds. OHM is described in Section 4.
(3) *QVM Clients.* A flexible set of clients that leverage the QVMI. In this paper we describe three example clients that enable checking of a variety of correctness properties with controlled overhead. Clients are discussed in Section 5.

[3]Sun Microsystems. Jvmtm tool interface, version 1.0.
http://java.sun.com/j2se/1.5.0/docs/guide/jvmti/jvmti.html.

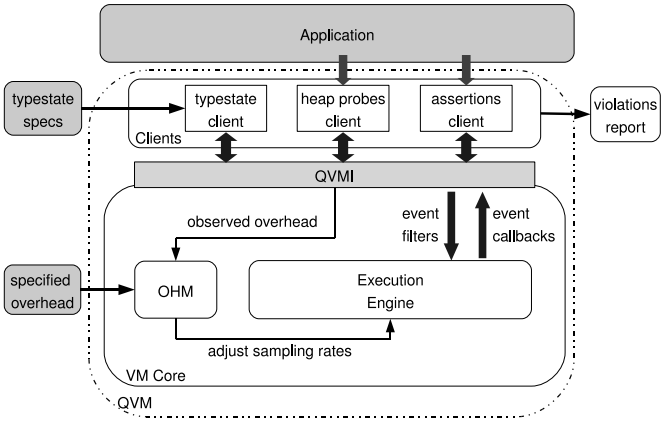

Fig. 5. Overall architecture of QVM.

In this architecture, the overhead manager and the QVMI work together to provide clients with a transparent adaptive overhead management. The clients use QVMI without the need to be aware of overhead management mechanisms (but with the ability to partially control it when desired).

The OHM uses the information collected by QVMI to adjust the sample rate such that the overhead matches the desired overhead specified by the user.

3.3 QVMI: The QVM Interface

Various profiling interfaces such as JVMTI make it easy to write monitoring clients. The client specifies the events of interest, and these events are provided by the interface. Clients are kept separate from the internal VM implementation that collects the events. Similarly, although our profiling clients are packaged as part of the VM, keeping a clear abstraction interface between the core VM details and the profiling clients is important for maintenance and ease of adding clients in the future.

The primary limitation with existing and general profiling interfaces is performance. For example the granularity at which events are requested is too coarse. With existing interfaces such as JVMTI, if a client wants to receive method callbacks for some subset of the method invocations, it must register to receive callbacks for *all* method invocations, and filter out the unnecessary callbacks on the client side of the interface. This introduces significant overhead that is completely unnecessary if the analysis needs only a subset of the methods.

Filtering on the VM side. To address this problem, the QVM interface is designed such that an efficient implementation is possible. The key difference from existing profiling interfaces is that it is structured with the goal of allowing as much filtering as possible to occur on the VM side of the interface. For example, if an analysis client needs method callbacks, it must specify what methods callbacks are necessary. This allows the remainder of the program to run at full speed. Similarly, the client may request method callbacks only for a subset of the objects in the program. The VM can use its suite of dynamic optimization techniques to achieve an efficient implementation of the sampled profile.

Table I shows a partial list of the operations supported by QVMI. Clients that register with QVMI have to support a similar set of operations (as described below). In addition to the operations listed in Table I, QVMI has similar callbacks for field read

Table I. A Partial List of the Operations Supported by QVMI

Method	Description
void registerClient(Client c)	registers a client to receive callbacks
TrackLevel isTrackedAlloc(AllocSite as)	should the specified allocation site be tracked
CallTrackLevel isTrackedCallSite(CallSite cs)	should the specified call site be tracked
boolean shouldExecute(Site s)	should this site fire an event (based on sampling info)
void allocEvent(AllocSite as)	tracked allocation event
void invocationEvent(CallSite as)	tracked invocation event
void objectDeath(Object o)	object death event

and writes, exceptions being thrown, and other events supported by standard interfaces such as JVMTI.

In the table, we separate operations of different stages of the execution by double horizontal lines. The manner in which these operations are used is illustrated below.

On VM initialization. Upon startup of the virtual machine, the clients have to register themselves with QVMI to receive callbacks by calling `registerClient`.

On method compilation. During the compilation of a method, the VM queries the QVM agents to determine whether the code being compiled needs any form of instrumentation. This insures that maximal filtering occurs; instrumentation is not inserted on any program statements if it is not required by at least one client.

This querying is done by invoking QVMI operations such as `isTrackedAlloc` and `isTrackedCallSite`, which query all of the registered QVM clients to obtain a `TrackLevel`, which determines what level of instrumentation is needed. For example, for our typestate client, the compiler prompts QVMI to check whether allocation or method call sites in the code should be tracked. Further details on how the typestate client is implemented via QVMI is discussed in Section 6.2.

During execution. Depending on the tracking level, the VM fires events for tracked sites by invoking operations such as `allocEvent` and `invocationEvent`. When an object is collected by the garbage collector, QVMI is notified by calling `objectDeath`.

3.4 Property-Guided Sampling

One of the major features provided by QVMI is the ability to perform *property-guided sampling*. Sampling is a key mechanism QVM uses to reduce analysis overhead, but for many analyses using naive random sampling would render the analysis useless because the analysis relies on certain relationships between events.

For example, if a dynamic analysis detects files that are opened but not closed, and tracking of method invocations were sampled randomly, QVM would report false positives any time file `open` was sampled, but its corresponding file `close` was not. To address this problem, QVM performs property-guided sampling, ensuring that the sampled profile maintains sufficient properties to make the dynamic analysis meaningful. For typestate properties, it is sufficient to maintain the relationship between events (method invocations) that occur for the same receiver object.

Object-centric sampling. QVM supports a novel feature called *object-centric sampling*. This technique allows an analysis to sample at the object instance level; an object can be marked as *tracked* and the analysis can receive all profile events for this object, while receiving no events for untracked objects. The choice on whether to track an object or not is made at allocation time (but can be more generally toggled on and off

during an object's lifetime). This allows overhead reduction via sampling, without destroying the profile properties needed for the dynamic analysis to produce meaningful results.

We refer to the points in the execution at which sampling decisions are made (i.e., whether an object is tracked, whether an assertion is executed) as *origins*.

Allocation sites are origins in our implementation of object-centric sampling. The decision of whether an object is tracked is made at allocation time; if sampled, a bit is set in the object header to mark the object as tracked. A short inlined code sequence checks this tracked bit on calls to QVM methods to determine whether a callback is needed.

3.5 Extensions

Our current interface is not intended to be complete, but is sufficient to cover a broad range of clients, including those described in this paper. The clients we implemented are built as part of the VM, but the interface could also be exposed to enable external clients. A full spec that could be published as a performance-aware alternate to the JVMTI is left for future work.

4. OVERHEAD MANAGER

Traditional dynamic analyses typically operate under the model that the user defines an analysis, then evaluates it to determine whether the overhead is acceptable. The instrumentation that is used to implement the analysis is fixed, and the overhead incurred is a function of the program that is executed.

The QVM *Overhead Manager*, or OHM, reverses this mentality by allowing the user to specify an overhead that is considered acceptable for the current monitoring environment. The maximum acceptable overhead may be 5%–10% in a live deployed system, yet 100% overhead (factor of 2 slowdown) may be considered acceptable in a testing environment.

Thus, the acceptable overhead is one of the inputs to QVM. Given an overhead budget, the QVM strives to collect as much useful information as possible from the executing program while staying within the specified budget. If the maximum overhead specified is too low, QVM may not report any useful information. This is obviously not the desired outcome, but in many cases it is more desirable than losing control of the overhead and having a performance crisis as a result.

There are three components to the overhead manager, each of which are discussed in the sections that follow.

(1) Monitoring: measures the overhead imposed by the QVM clients;
(2) Sampling strategy: a strategy for sampling each origin (e.g., allocation site or an assertion site) to ensure the system stays within the overhead budget;
(3) Controller: adjusts the sampling strategies for each origin based on the measured overhead.

4.1 Monitoring

The overhead monitor uses fine granularity timers on entry and exit to all QVMI calls to record the time spent in QVM clients and in the QVMI itself. The time is maintained separately for each origin (see Section 3.4) so that the sample rate of each origin can be adjusted independently.

Timer accuracy. The most important step in managing overhead is having the ability to measure overhead accurately. The overhead controller cannot be expected to make reasonable decisions if it is being given incorrect timing data as input.

Measuring overhead for coarse grained events (such as garbage collection time) is relatively easy; a number of system timing routines can be used to obtain reasonable results. However, timing short, frequently executed regions is more difficult and requires having a timer that is both accurate and efficient.

Using an inefficient timer mechanism has two serious problems: 1) it can cause significant overhead if called frequently (which can be the case with some QVM clients), and 2) the error can be significant when timing short regions and these timing errors will accumulate.

To address these problems, our OHM implementation uses inline assembly to read the cycle counter using the Intel's RDTSC (Read Time Stamp Counter) instruction. This mechanism results in very fast and accurate time stamping on entry and exit of the QVMI. Our initial implementation used the system call `gettimeofday()` and it created significant inaccuracies, as described in Section 7.

Measuring total application time. The timers measure time spent performing QVM tasks. To compute overhead relative to the non-QVM application, the OHM must also measure the total execution time. Using wall clock time, rather than process time, would be grossly incorrect for two reasons. First, interactive applications would create significant error because idle time would be counted as application time. Second, wall clock would be wrong for multi-threaded applications running on multi-processor machines. QVM time is measured and accumulated from all running threads, thus the total time must be the sum of the time spent executing on all processors.

For these reasons, we compute total time by using the `getrusage()` Linux system call to obtain the total time used by the JVM process. This solves the problems associated with using wall clock time discussed above and works well in practice for most applications. However, it is still not a fully robust solution when QVM activity is not evenly distributed across the application threads.

For example, consider an application with 2 threads running for 1 second each in parallel on a 2-processor machine; `getrusage()` will report 2 seconds of total execution time. Assume that QVM was given a 10% overhead budget, which translates to 0.2 seconds allocated to QVM. If all of the QVM callback activity takes place in one of the two application threads, one thread will run for 1.2 seconds while the other runs for 1 second. Although the total CPU time is increased by 10% a user of the program would observe the program terminating after 1.2 seconds, a 20% increase.

The most robust solution to this problem is to perform overhead tracking at the thread level. If overhead budgets are tracked and enforced per thread, total overhead as perceived by the user will always be within budget as well. A similar approach of using per thread metrics has been employed by real time systems to track time spent performing system services [Auerbach et al. 2008]. We leave an implementation of this approach within QVM as part of future work.

Base overhead. Even when accurately measuring the time spent in the QVM clients, there are still two potential sources of errors: 1) checking overhead, and 2) indirect effects.

The main sources of checking overhead is the inlined filtering. For example:

— virtual method calls (or inlined method bodies) for methods relevant to QVM clients filter samples by checking a bit in the object header.
— origin sites (i.e., allocation sites) check their sampling strategy (described in Section 4.2) to decide whether the allocated object is tracked.

These checks are short inlined code sequences and contribute very little to overall overhead (see Section 7); however, for very aggressive instrumentation, such as

instrumenting all calls in the program, the base overhead can potentially become significant.

Although not easy to measure online while the application is executing, base overhead can be estimated by observing the frequencies of the checks, and using a model of performance to estimate the overhead. Using a model is less desirable than direct measurement, but can still be used as a way of avoiding large performance surprises for aggressive clients.

Our implementation does not yet perform this modeling to avoid large base overhead, and it is left as part of future work.

The second source of base overhead is indirect effects on performance, such as cache pollution, or optimization in the JIT that are hindered by the presence of instrumentation. These sources of overhead are very difficult to measure without having two separate versions of the code and using techniques such as performance auditor [Lau et al. 2006] to identify the performance differences.

Although in general the base overhead might be high, in our experiments we observed a low base overhead. In Section 7.2, we show that the base overhead in our experiments was at most 2.5%. (See Figure 12).

4.2 Sampling Strategy

The QVMI maintains separate overhead statistics for each origin (see Section 3.4), allowing the OHM to increase or decrease the sample rate independently for each origin. Having origin-specific sample rates enables significant advantages for the client analysis. Maintaining a single sample rate would be sufficient for managing total overhead, but would be likely to miss origins in infrequently executed code. With origin-specific sampling, the controller can reduce overhead by scaling back hot origin sites, but continues to exhaustively track objects from cold sites, thus allowing the client analysis to see a broader view of the program execution. As shown in Section 7, this sampling strategy results in increased error coverage for a given overhead budget. This approach is similar to Hauswirth and Chilimbi [2004],which uses inverse sampling to avoid missing memory leaks in cold code.

Our implementation achieves sampling by maintaining a `sampleCounter` and a `sampleCounterReset` for each origin. At runtime, the checking code at each origin site decrements and checks `sampleCounter`; if it is less than zero, the origin is selected to be tracked and the counter is reinitialized by the value in `sampleCounterReset`.

The `sampleCounterReset` for each origin is adjusted by the Overhead Controller to change the sample frequency for that origin, thus reducing or increasing its overhead.

Emergency shutdown. Object-centric sampling is most effective for managing overhead when there are a large number of objects contributing to total overhead. If the majority of execution is dominated by method calls on a single, long-lived object, tracking this object will result in large overhead. To avoid severe performance degradation when a hot, long-lived object is tracked, the QVM supports the notion of an *emergency shutdown*. On each QVMI callback for allocations and invocations, the system checks a flag to determine whether an emergency shutdown is needed. If so, it disables the monitoring bit in the object header such that the object will no longer be sampled. The client analysis may now need to discard this object, as the method callbacks are not complete. However, this mechanism allows the system to ensure that overhead can be controlled.

4.3 Overhead Controller

The job of the Overhead Controller is to periodically check the QVM overhead, and adjust the sampling frequencies accordingly. The OHM is a basic feedback loop: if the

```
while (1) {

    sleep(OHM_SLEEP_INTERVAL);

    // 1) Adapt global threshold
    if (OBSERVED_OVERHEAD > OVERHEAD_BUDGET) {
        // Exceeding budget, sample less
        if ((OBSERVED_OVERHEAD - OVERHEAD_BUDGET) > CRITICAL_BUDGET_OVERSHOOT)
            // Too far over budget, shut everything down
            triggerEmergencyShutdown()
        else
            // Common case
            ORIGIN_OVERHEAD_TARGET = reduceThreshold(ORIGIN_OVERHEAD_TARGET);
    }
    else if (OBSERVED_OVERHEAD < OVERHEAD_BUDGET) {
        // Under budget, sample more
        ORIGIN_OVERHEAD_TARGET = increaseThreshold(ORIGIN_OVERHEAD_TARGET);
    }

    // 2) check each origin's overhead against ORIGIN_OVERHEAD_TARGET
    for each origin
        if (getObservedOverhead(origin) > ORIGIN_OVERHEAD_TARGET)
            decreaseSampleRate(origin);
        else
            increaseSampleRate(origin);
}
```

Fig. 6. Pseudocode for the overhead manager.

overhead is above the budget, sample frequencies are reduced; if the overhead is below budget, the frequencies are increased.

To avoid oscillation and large spikes in overhead, the controller monitors not only total overhead, but *recent* overhead. Recent overhead is computed via exponential decay; a second copy of application time and QVM time are maintained, and multiplied by a decay factor each time the controller wakes up. This gives more weight to recent timings, effectively measuring the overhead over a previous window of execution. The primary focus of the controller is keeping the overhead below the overhead budget. Maximizing the client executing time within that budget is also a goal, but it is secondary. Thus the controller reduces sample frequencies if either the total overhead or recent overhead exceed their budgets.

If the overhead deviates too high above the budget, the controller enacts the emergency shutdown to stop profiling in the current set of objects, and starts tracking new objects once the overhead is within budget.

Origin-specific adjustment. The QVMI maintains separate overhead statistics for each origin (see Section 4.2), allowing each origins's sample rate to be increased or decreased independently. Therefore, when adjusting sample rates the overhead manager must make an independent decision for each origin, whether the sample rate should be increased or decreased, to achieve the desired *total* overhead.

The overhead manager achieves this goal by maintaining a second overhead threshold, called `ORIGIN_OVERHEAD_TARGET`, which is the target threshold that the overhead of each individual origin is compared against. An origin's sample rate is not reduced unless it is overhead is above this target threshold, allowing infrequent sites to be monitored exhaustively. More active origins will have their sample rates reduced until their overhead is below this threshold, thus ensuring that a) the overhead budget is distributed (roughly equally) over a all origins, b) the total overhead does not exceed the global budget. The overhead manager logic (simplified for presentation) is shown in pseudocode form in Figure 6.

The parameter CRITICAL_BUDGET_OVERSHOOT controls when emergency-shutdown is triggered. The value for this parameter can be set by the user, and the default value for this parameter is 50% of the target overhead. This triggers an emergency shutdown when the observed overhead exceeds that target overhead by 50%. (The effect of emergency shutdown is shown clearly in Figure 13 where the observed overhead reaches 15% and the target overhead is set to 10%.)

5. QVM CLIENTS

In this section, we describe three clients built on top of the QVM platform. We have implemented a number of clients in order to cover a range of user properties: ranging from local assertions to continuous monitoring using temporal safety properties.

5.1 Typestate

In this section, we show how QVM is used to dynamically check typestate properties.

Typestate [Strom and Yemini 1986] is a framework for specifying a class of temporal safety properties. Typestates can encode correct usage rules for many common libraries and application programming interfaces (APIs). For example, typestate can express the property that a Java program should dispose a native resource before its Java object becomes unreachable and is collected by the garbage collector.

Dynamic checking of typestate properties, as well as generalized multiple-object typestate (also known as "first-order properties" [Ramalingam et al. 2002; Yahav and Ramalingam 2004]), have been addressed before in Tracematches [Allan et al. 2005] and MOP [Chen and Roşu 2007]. We use the typestate client to demonstrate three contributions of our platform: (i) adaptive overhead management; (ii) timed typestate transitions; (iii) collection of additional detailed debug information with low overhead.

Using the QVM platform to implement dynamic typestate checking also provides us with an advantage in getting object-death callbacks directly from the garbage collector and not relying on a finalizer method to be called. This guarantees that object-death events are fired in a timely manner (which is not guaranteed to happen when using finalizers) and allows us to measure resource-drag (as follows) more precisely.

Definition 5.1. A typestate property $\mathcal{F}$ is represented by a deterministic finite state automaton (DFA) $\mathcal{F} = \langle \Sigma, \mathcal{Q}, \delta, init, \mathcal{Q} \setminus \{err\} \rangle$ where Σ is the alphabet of observable operations, $\mathcal{Q}$ is the set of states, δ is the transition function mapping a state and an operation to a successor state, $init \in \mathcal{Q}$ is a distinguished *initial state*, $err \in \mathcal{Q}$ is a distinguished *error state* for which for every $\sigma \in \Sigma$, $\delta(err, \sigma) = err$, and all states in $\mathcal{Q} \setminus \{err\}$ are accepting states.

QVM uses a simple input language to let the user specify a finite-state automaton that represents the typestate property, and the types to which it applies. We refer to a type that appears in at least one typestate property as a *tracked type*. Once the tracked type is specified, our implementation instruments every object of this tracked type with additional information that maps the object to its typestate. During execution, QVM updates the typestate of each tracked object, and when an object reaches its error state, QVM records an error report (as the one shown in Figure 1) in a designated log file.

Example 5.2. Figure 7 shows a typestate property (represented as a finite state automaton) that identifies when an SWT resource has not been disposed prior to its garbage collection, thus possibly leaking native resources such as GDI handles. The tracked types are not shown in the figure, as this property applies to a large number of types (e.g., org/eclipse/swt/widgets/Widget). Since all states other than the designated error state are accepting, we simplify notation by not using a special notation for accepting states. We label edges of the finite-state automaton with regular expressions

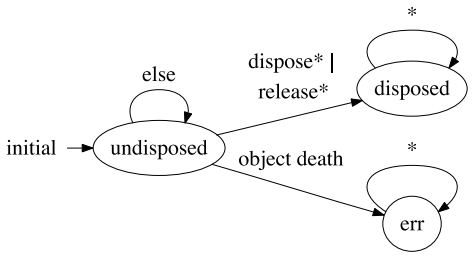

Fig. 7. A typestate property tracking proper disposal of SWT resources. Names of tracked types are not shown.

that define when the transition is taken. For example, the transition from undisposed to disposed occurs when invoking a method whose name begins with `dispose` or `release`. We use else to denote a transition that is fired when no other transition from the state can be matched (note that the automaton is deterministic).

In Section 7.1 we report experimental results for such properties.

For every typestate property, QVM tracks the number of times it has been violated. When the number of violations passes a specified threshold, QVM starts recording additional debugging information in the form of a *typestate history*.

As mentioned in Section 2, a typestate history of an object o is an abstraction of the sequence of method invocations performed during execution with o as a receiver. We use the name typestate history because we summarize the sequence of method invocations as an annotated DFA, similar to a typestate property.

We now formally define the notion of a typestate history. First, we define our notion of a method invocation. A method invocation s is a pair $\langle sig, t\rangle$ recording the method signature sig invoked and the time t in which it has been invoked. We use $m(s) = sig$ to denote the method signature of an invocation s, and $time(s) = t$ denote the time in which it took place. Given a sequence of method invocations $S = s_0, \ldots, s_k$, the *bounded context* of depth n for the i-th invocation s_i in S is the reverse subsequence starting at s_i and going backwards for n steps or until reaching s_0 (whichever comes first). That is, $context(S, i, n)$ is $s_i, s_{i-1}, \ldots, s_{i-n}$ when $n \leq i$, or $s_i, \ldots, s_0$ when $n > i$. The context of a method signature sig is the set of contexts of all of its invocations $context(sig, S, n) = \{context(S, i, n) \mid m(s_i) = sig\}$. We denote the set of all possible contexts in S by CTX_S.

The following provides a (declarative) definition of a typestate history constructed from a sequence of method invocations.

Definition 5.3. Let $S = s_0, \ldots, s_k$ be a sequence of method invocations over a tracked object o. A typestate history for the object o and the sequence S is represented as an annotated DFA $\mathcal{H}_S = \langle \Sigma, \mathcal{Q}, \delta, I, \mathcal{Q}, count, last, ctx\rangle$ where:

— The alphabet $\Sigma = \bigcup_{0\leq i\leq k} m(s_i) \cup \{\text{ODE}\}$ is the set of signatures of methods invoked in the sequence and a special event ODE corresponding to object-death.
— The set $\mathcal{Q} = \bigcup_{0\leq i\leq k} m(s_i) \cup \{I\} \cup \{\text{OD}\}$ is the set of states, where there is a single state for each method signature appearing in S, and two designated states: an initial state (I) and an object-death (OD) state.
— The transition function δ is defined by the sequence S: $\delta(I, m(s_0)) = m(s_0)$, for any $0 \leq i < k$, $\delta(m(s_i), m(s_{i+1})) = m(s_{i+1})$, and $\delta(m(s_k), \text{ODE}) = \text{OD}$.
— $I \in \mathcal{Q}$ is a distinguished *initial state*.
— All states in $\mathcal{Q}$ are accepting states.

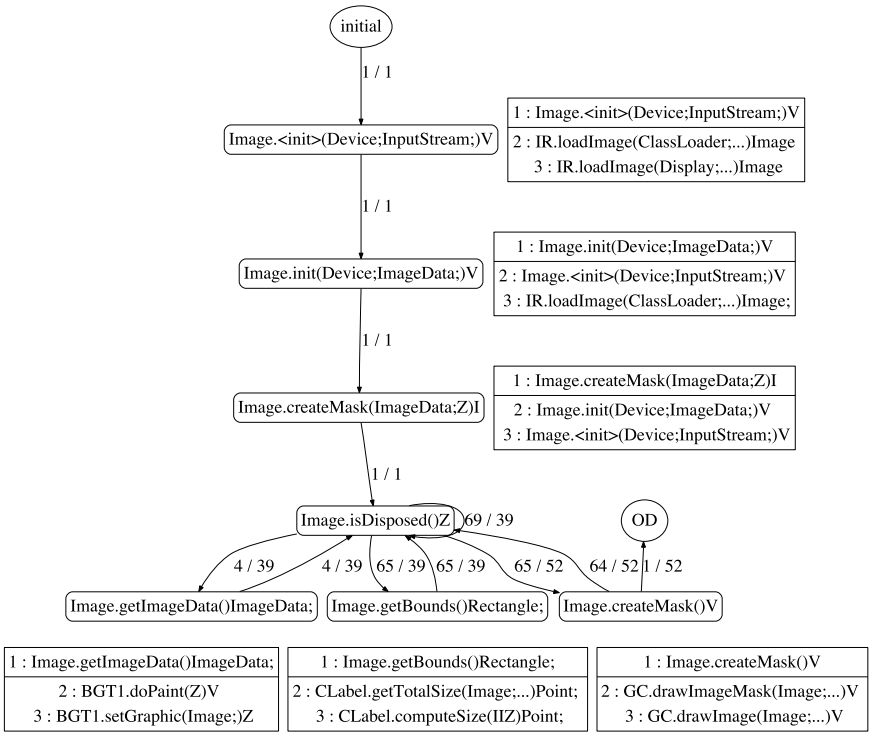

Fig. 8. An example typestate history for a leaking `Image` in Azureus. For brevity, we only show sample contexts and omit the context for `isDisposed`.

— $count : \mathcal{Q} \times \Sigma \rightarrow \mathcal{N}$ records the number of times the invocation corresponding to an edge has been taken in S.
— $last : \mathcal{Q} \times \Sigma \rightarrow \mathcal{N}$ records the last time the invocation corresponding to an edge has been taken in S.
— $ctx : \mathcal{Q} \rightarrow 2^{CTX_S}$ maps a state (corresponding to a method) to the set of contexts in which the method is invoked in S.

Intuitively, a state in the typestate history represents the state of the object after a specific method has been invoked on it. A state in the history is labeled with a set of (bounded) contexts, representing the contexts in which the method has been invoked. A transition between states m_1 and m_2 in the history represents the fact that the method corresponding to m_2 has been invoked immediately after the method corresponding to m_1 has been invoked.

A typestate history therefore provides information about the way a single object that violates the property was used in the program. This helps the programmer to diagnose the cause of the reported violation.

Example 5.4. Figure 8 shows an example typestate history produced by QVM. This provides an account of the behavior of a single object that violates the property. In the figure, we have abbreviated the type name `BufferedGraphicTableItem1` to `BGT1`, and the type name `ImageRepository` to `IR`. In figures of typestate histories we do not show method signatures on the edges because the label of an edge is always identical to the label of its target state.

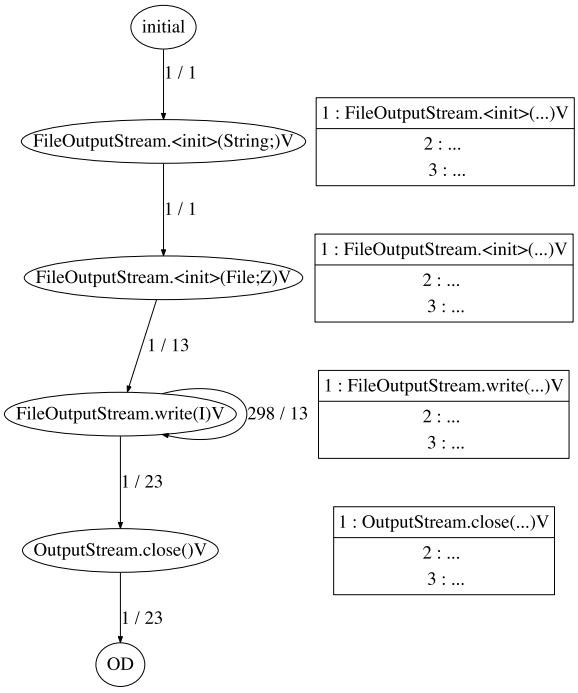

Fig. 9. An example typestate history showing resource lag and drag.

Unlike the simple typestate history of Figure 2, the typestate history of Figure 8 contains cycles and multiple invocations of methods. The label on a transition edge represents the *number of times* this transition occurred in the execution and the *last time* when it occurred. For example, the transition from the state in which `createMask` is the last method invoked on the object to the state in which `isDisposed` is the last method invoked on the object occurs `64` times in the execution summarized by the history of Figure 8. The last time in which the transition occurred is 52, where time is measured as the number of allocations performed by the program. In the figures, we show the time counter divided by 1024.

Resource Drag and Lag. Since QVM tracks the last time each transition took place, it can be used to identify when a resource is not released in a timely manner (known as resource drag). In such cases it is sometimes possible to improve performance by releasing the resource earlier. Similarly, since QVM also tracks calls to constructors and object-death events, it can be used to identify when an object is allocated too early (memory lag) or kept reachable for a longer time than necessary (memory drag). The information collected by QVM can also be used to find objects that are not allocated in a timely manner, that is, a long period of time passes between their allocation and their first use. In such cases, it is sometimes possible to improve performance by allocating the objects lazily.

Technically, we only track the last time in which a transition is taken, but since a constructor is only called once, the outgoing transition from the last constructor state only occurs once, and will therefore provide us with the time of the first use.

Example 5.5. Figure 9 shows a typestate history for a `FileOutputStream` in which the allocation was performed at time 1, the first use (transition from the constructor state)

Table II. QVM Heap Probes. Probe Returns *true* if the Condition Holds, *false* Otherwise

Probe Name	**Condition Checked**
isHeap(Object o)	object o is pointed to by a heap object
isShared(Object o)	object o is pointed to by two or more heap objects
isObjectOwned(Object o_1, Object o_2) isObjectOwned(Object o)	o_1 dominates o_2 the object pointed to by `this` dominates o
isThreadOwned(Thread t_a, Object o) isThreadOwned(Object o)	t_a dominates o current thread dominates o, *false* otherwise
isUniqueOwner(Object *root*)	*root* dominates all objects that are transitively reachable from *root*
isReachable(Object src, Object dst)	object *dst* is reachable from object *src*

took place in time 13, the last use of the stream (last write) also took place at time 13, closing the stream was performed at time 23, and its actual death (collection) took place at time 23 as well. This example exhibits some amount of resource lag (long time passed between allocation and first use) and some amount of resource drag (stream was not closed immediately after the writes). It does not exhibit memory drag, as the object is collected soon after it is closed.

Idle Objects. A special case of an inefficient use of allocated resources is that of objects for which no method invocation occurred other than their construction. We refer to these as *idle objects*.

Example 5.6. Running QVM on Eclipse, we found a large number of idle objects allocated by the subclipse plugin. In a typical synchronization with SVN, over 145,000 idle objects of type `DataInputStream` were created. Tracking the source of these idle objects, we found that they are (eagerly) allocated in the constructor of `StatusFromBytesStream`. We manually inspected the code and confirmed that this allocation can be modified to be done lazily when methods of `StatusFromBytesStream` require the stream rather than when a `StatusFromBytesStream` object is created.

Extensions. Next, we describe several possible extensions to our system. Our current implementation supports single-object typestate properties. A useful extension is handling multiple object typestates. Multiobject typestate has been actively studied in Tracematches [Allan et al. 2005] and MOP [Chen and Roşu 2007].

Another possible extension is using static analysis to eliminate some of the dynamic checks performed by QVM. This will lead to reduced runtime overhead. For example, in some cases, we can use static analysis [Bodden et al. 2007; Fink et al. 2006] to verify that a typestate property is never violated, or that some transitions of a typestate property never occur in the program.

5.2 Heap Probes

QVM enables the dynamic checking of various global heap properties such as object-sharing, heap-ownership, thread-ownership and reachability. These properties are useful for both debugging and program understanding purposes.

QVM provides a library that exports a set of methods, one for each heap property. We refer to these library methods as *heap probes*. The programmer can invoke heap probes from her program to inspect the shape of the heap at a program point. The library uses various components of the underlying runtime to obtain an answer. Our list of currently supported probes is shown in Table II.

In this paper, we describe heap probes at a high level, and focus on how they are used. Some additional details are provided in Arnold et al. [2008] and Vechev et al.

```
canvas.addDisposeListener(new DisposeListener() {
 @Override
 public void widgetDisposed(DisposeEvent arg0) {
  if (img != null && !img.isDisposed()) {
   assert(QVM.isObjectOwned(img));
   img.dispose();
  }
 }
});
```

Fig. 10. Using QVM to check that an SWT resource is not shared before attempting to dispose it.

[2010], but the implementation involves many subtle details that are beyond the scope of this article and the scope of Arnold et al. [2008] and Vechev et al. [2010]. We expect such details to appear in future work.

Similarly to nonheap probes, our heap probes can be sampled by the overhead manager to allow adjustment of overhead, and can therefore evaluate to one of three possible values: true, false, and unknown. The return value of a heap probe can be used in a standard Java assertion. When a heap probe is used inside an assertion we refer to it as a *heap assertion*.

In the case where the assertion is skipped (due to sampling) the return value of the assertion is *unknown*. In our implementation we map this value to *true*. That is, the net effect in this case is as if the assertion has not been executed. (This default behavior can be controlled by the user.)

Example 5.7. Disposal of SWT resources is based on two key principles: (i) a resource is disposed by calling a method on the object that allocated the resource. (ii) disposing resources of the parent object leads to disposing resources of its children.

These principles work well in cases where large numbers of allocated resources form an immutable containment tree. In such cases, disposing the resources of the parent leads to disposing the resources of its children. However, the treatment of shared resources such as Color, Fonts, and Images, is more complicated and error prone.

For shared resources, it may be rather challenging to find the correct program point at which it is safe dispose of the resource. A programmer may have a conjecture about the last program point in the application execution where a resource is used, but this conjecture may turn out to be incorrect. The heap may contain other references to, and future usage of, the resource.

Figure 10 shows how a QVM assertion can be used by the programmer to ensure that the resource is not shared by any other object except the current object (the object *this*). The code fragment shown here corresponds to a common idiom for disposing a resource by a dispose listener. This particular code fragment is taken from a fix we introduced for the Azureus benchmark as described in Section 7.1.

5.2.1 Ownership and Alias Control. Ownership simplifies reasoning about object-oriented programs by controlling the permitted aliasing. Ownership has been used in many settings. It has been used to ensure representation independence [Banerjee and Naumann 2005], to guarantee thread safety [Boyapati et al. 2002], and to enable modular reasoning [Rinetzky et al. 2007].

A wide variety of static approaches have been proposed for enforcing ownership (see Clarke [2003] for a nice survey). These approaches typically impose strict restriction on ownership transfer, requiring, e.g., that uniqueness [Aldrich et al. 2002; Baker 1995; Hogg 1991] holds on transfer, or impose a high annotation burden.

Our approach to ownership assertions complements static approaches for enforcing ownership. In particular, our approach may enable a type system to balk at some cases

```
public class SimpleWebServer ... {
 public void run() {
  while (_running) {
   Socket wsocket = _serverSocket.accept();
   RequestThread rt = new RequestThread(wsocket, _rootDir);
   wsocket = null;
   rt.start();
  }
 }
}

public class RequestThread {
 private Socket _socket;
 ...
 public RequestThread(Socket socket, File rootDir) {
  _socket = socket;
  _rootDir = rootDir;
 }
 public void run() {
   assert(QVM.isThreadOwned(_socket));
   ...
 }
}
```

Fig. 11. QVM thread ownership assertion checking that the socket used by a request thread is owned by the thread.

when ownership cannot be established, leaving an ownership check to be performed at runtime.

In addition, the QVM support for ownership properties can provide an alternative efficient implementation to the runtime support required by some ownership type systems [Müller and Rudich 2007].

Stack Confinement. It is often desirable to check that an object does not escape from a procedure to be stored in the heap. This is particularly important in a concurrent setting where exposing a heap-reference to an object might lead to an undesirable concurrent modification. Using the probe $isHeap(o)$, QVM allows the user to check that the only references to an object are stack references, and that the object is not pointed to from the heap.

Thread Confinement. QVM allows the user to check that an object is owned by a given thread. The probe $isThreadOwned(t, o)$ checks whether object o can be pointed by any object that is not transitively reachable from the thread t.

Example 5.8. Figure 11 shows a code fragment taken from SimpleWebServe[4] with an additional QVM assertion. In the SimpleWebServe, a new thread is created for every request received by the web-server. The new `RequestThread` is passed a `Socket` through which it communicates with the client. The QVM assertion guarantees that the `Socket` passed to a `RequestThread` is *owned* by the thread.

6. IMPLEMENTATION

In this section we provide the implementation details of object-centric sampling, as well as QVM clients of Section 5.

6.1 Object-Centric Sampling

There are two key components to the efficient implementation of object-centric sampling.

[4]Jibble Web Server. http://www.jibble.org/jibblewebserver.php.

The first one is the ability to obtain a single free bit in the object header. Compared to the approach of reserving a word in the object, this approach has an advantage of better space efficiency and increased locality.

Once identified as a tracked object, QVM clients need the ability to associate analysis data with an object. We implemented this in QVM by creating an OBJECTINFO for every tracked object. This ObjectInfo is then passed to the client on all object-related callbacks so the client can lookup or store data associated with the object (such as DFA state, etc).

The mapping from object to ObjectInfo is performed via a hashtable lookup. On allocation of an object, the corresponding ObjectInfo is created and inserted into the hashtable; on object death, they are removed. QVMI callbacks that require access to the ObjectInfo obtain it by doing a hash lookup.

As mentioned, an alternate implementation would be to reserve a word in the object header to point to the object's ObjectInfo. While this provides faster lookup, it is not necessarily the superior design because it reduces locality by increasing object size, and this overhead is regardless of the sample rate. A hashtable lookup is significantly slower, but the hashtable lookup is performed only for sampled objects; the inlined fast path only checks the tracked bit in the object header.

So although the hashtable implementation is slower for tracked objects, it allows a lower base overhead that is converged upon when when the sampling rate is reduced.

Because the goal of QVM is to target low-overhead scenarios, the hashtable design was chosen.

6.2 Typestate Client

Upon VM startup, the typestate module loads all of the user supplied properties, parses and stores that information in its own internal data structures. The typestate module then registers itself with the runtime via the `QVMI.registerClient` call.

On method compilation, the QVMI interface is called by the JIT via the `isTrackedAlloc` and `isTrackedCallSite` functions to determine whether instrumentation is needed for allocations and calls. These functions return a value of type `TrackLevel`. This type can take on one of three totally ordered values: NEVER (the minimal value), SOMETIMES and ALWAYS (the maximal value). All of the registered QVM clients are queried and the return result is computed by taking the maximal value from all of the client responses to ensure that sufficient instrumentation is inserted.

QVM then adjusts the instrumentation based on the tracking level. If the tracking level is ALWAYS or SOMETIMES, QVM instruments the code with a callback to report the event that occurred. In the case of SOMETIMES, QVM inserts inlined logic to decide (during execution) whether the callback gets invoked. If the tracking level is NEVER, no code instrumentation is performed by QVM for the site. If the tracking level is ALWAYS the callback is executed exhaustively and sampling is disabled.

For allocations sites marked with track level SOMETIMES, the inlined sampling logic consults the sampling strategy for that origin (see Section 4.2). If selected for sampling, the typestate allocation handler is called via the QVMI `allocEvent` call. The handler creates its internal QVM tracking structure for the allocated object, and marks the object as tracked by setting a bit in the object header. Note that there could be multiple tracking structures per object (e.g., the object is part of multiple typestate properties).

For method invocations tagged with SOMETIMES, the inlined code sequence checks whether the receiver is a tracked object by checking the tracked bit in the header. This check is executed even for inlined methods to ensure that callbacks are not optimized away by the JIT. If the object's tracked bit is set, QVMI's `invocationEvent` is invoked, which then calls the typestate invocation handler. The handler is passed the receiver

object, that object's ObjectInfo, and the method that was invoked. This handler updates the tracking structure for each DFA the object participates in.

In our implementation for typestate, we have used the object-centric tracking and sampling capabilities provided by QVMI (Section 3.4) and have inlined check of whether the object is tracked. This keeps overhead low by ensuring that QVMI is invoked only for tracked (sampled) objects. There are many other such property-specific optimizations that can be made. For example, if we know that the tracked object is in an error state that will not be exited, QVM does not need to invoke any other callbacks on this object.

On Object Death. We have instrumented the garbage collector to provide precise death events. Whenever an object is detected to be unreachable during the sweep phase of the collector, the collector calls the QVMI's `objectDeath` function. That function leads to calling the typestate module's handler for death events, where all object tracking information is freed (if the object is tracked), ensuring no memory leakage. If the object is found to be in a nonaccepting state, an error is reported.

6.2.1 Collecting Typestate Histories. In typestate histories, we use a notion of "time" to record when events occurred. We measure the time as the number of allocations performed by the program. To provide a scalable and efficient implementation of a global clock, each thread maintains a local allocation counter, and these are aggregated to a single global (approximate) time every 10 millisec. The precision of the aggregate global clock can be adjusted by the user by changing the frequency of aggregation operations (at the cost of a performance hit when using higher frequency).

6.2.2 Discussion. Although the typestate module is written as part of the VM, it is completely isolated from the VM via the QVMI interface; this interface can be used to easily write clients to check properties other than typestate. By having access to an unused bit in the object header bits, QVM is able to efficiently perform object-centric sampling without needing to store additional words in the object. Moreover, the ability to precisely intercept object death events frees us from having to rely on technique such as finalizers and weak references.

7. EXPERIMENTAL EVALUATION

Our experimental evaluation focuses on typestate properties. In Section 7.1, we conduct a preliminary study of the effectiveness of QVM in tracking typestate properties with low overhead. This part of the study focuses on the usability of interactive applications running with QVM. In Section 7.2, we evaluate the overhead manager by tracking a large number of typestate properties over standard non-interactive benchmarks.

7.1 Detecting Resource Leaks

In our experiments we focused on typestate properties that correspond to resource leaks. We monitor leaks of SWT resources and of IO streams. In these experiments the goal was to see if we can detect typestate violations that occur over an extended period of time. It is likely that massive leaks would have been detected and fixed in the testing phase, and therefore what we expect to find in these experiments is mostly a small number of leaks that accumulate over time. For that purpose, we used a range of applications on a regular basis to perform our daily tasks.

This part of the evaluation is intended as a preliminary proof of concept. We are aware that it would be hard to reproduce these exact experiments as they depend on

Table III.

Sources of typestate violations in our application For every application, we indicate the number of sources that are executed in a high-frequency (corresponding to critical leaks).

Application	Description	SWT Resources	IOStreams	High Frequency	Fixed
azureus	bittorrent client	11	0	4	5
etrader	trading platform	17	0	2	0
feednread	newsfeed reader	1	7	0	0
goim	IM client	3	0	1	3
ibm app 1	large scale product	0	0	0	0
ibm app 2	medium scale tool	3	2	0	0
jcommander	file manager	9	0	0	0
juploader	flickr upload tool	0	1	0	0
nomadpim	personal information manager	2	0	0	0
rssowl	RSS news reader	8	3	0	0
tvbrowser	TV program guide	0	5	0	0
tvla	program analysis framework	0	4	0	0
virgoftp	FTP client	6	0	0	6
Total		**60**	**22**	**7**	**14**

specific user interactions. However, it is precisely this inherent inability to find and reproduce defects in a lab setting which calls for systems which are designed for low overhead application monitoring. QVM is one such system prototype.

Some of the applications considered are an instant-messenger (goim), newsfeed readers (feednread, rssowl), file management utilities (virgoftp, jcommander), large IBM internal applications, etc. For all of these applications our strategy was to simply run them over QVM and record the reported errors. In some of our experiments we investigated each report manually, diagnosed the causes of the errors, and implemented fixes. This was an important exercise for evaluating and refining the debug information we collect (e.g., the typestate history).

7.1.1 Applications and Results. Table III summarizes the number of sources of typestate violations found in our applications. Rather than counting the number of objects that violate the property, we count the allocation sites in which such objects were allocated. This is a more objective measure of the number of bugs in the program than the number of objects exhibiting the violation, which usually depends on the duration of program execution. To measure the significance of a violation, we record whether it occurs frequently in the program execution.

Counting the number of violation sources has to be done carefully as the sources are not necessarily independent. For example, a whole sub-tree of components may leak due to a single missing dispose operation on the root of the tree.

In some of our experiments we investigated the errors and introduced appropriate fixes. Column fixed in the table reports the number of fixes we have introduced and tested.

Azureus. Azureus is the #1 downloaded Java program from SourceForge, and has more than 160 million downloads to date. Using QVM we detected 11 sources of resource leaks in this application. We fixed 5 of these and reported them to the Azureus development team. The reports were confirmed by the development team, and the fixes were incorporated into the codebase.

At least 4 of the reports correspond to leaks that were occurring rather frequently. One particularly high-frequency case was a method `Utils.getFontHeightFromPX(...)` that was allocating a `Font` object to compute font height and was not properly disposing the `Font` object upon its return. This method is frequently called and resulted with thousands of leaking fonts even for short executions. This method was very likely created by copying another method in the class that has similar functionality but returns the `Font` object. Among our other fixes, we fixed the frequently leaking method `getFontHeightFromPX(...)` and our fix was incorporated into the Azureus codebase.

Eclipse Trader. This application uses a frequently-updating UI to present streams of stock information, and as a result may be particulary sensitive to resource leaks. Using QVM, we detected 17 sources of resource leaks. For eclipseTrader, there are several allocation sites that are used in different contexts. For example, the method `Settings.getColor(Color)` returns a new `Color` object, and is used in a large number of contexts that fail to properly dispose the color. We count this method as a single violation source that occurs with high frequency (there are tens of thousands leaking objects that are allocated in this method in a typical execution of eclipsetrader).

Feed'N Read. In this application, SWT resources are mostly properly managed. There are some resources that are not disposed before the program exits, but these are resources that are supposed to be live throughout program execution. Although QVM reports these as violations, we do not count them here because this seems to be acceptable treatment of such resources (resources will be returned to the OS when the application terminates). feednread has minor problems in closing IO streams when managing archived feeds.

GOIM. We used GOIM[5] running on QVM to communicate between team members for a few days. We detected 3 sources of leaks and introduced fixes to all of them. We tested our fixed version of GOIM and confirmed that all previously reported leaks have been resolved. The fixes we introduced in GOIM were rather involved as we had to add new disposal code in places where no such code existed.

IBM Applications. We used QVM to run a development version of a large scale IBM product on a daily basis for a period of a few weeks. For this application, no problems were reported by QVM. This is not surprising as the development team is putting a lot of emphasis on preventing the kind of leaks we are tracking.

We used QVM to run a development version of another smaller IBM tool that makes heavy use of SWT. For this application, we found 5 source of violations. The leaks are associated with user actions like opening a new file.

JUploader. This application uses a small number of SWT resources. For this application we found a single source of leaks causing a frequent leak of `EventOutputStream` objects.

TVLA. Running TVLA [Lev-Ami and Sagiv 2000] we found two input streams that are not closed by the parser processing input files, and two streams that are not closed when producing analysis output. These are low frequency leaks that only create one leaking object per execution of the analysis engine.

[5]GOIM: Gamers own instant messenger. available at http://goim.us/wiki/show/GOIM.

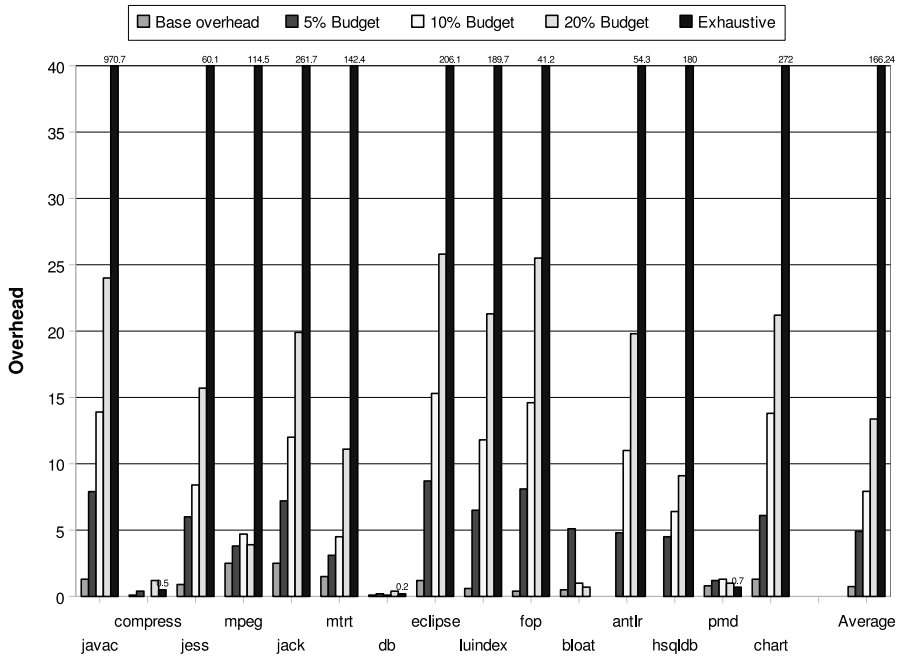

Fig. 12. Overhead with budget.

VirgoFtp. VirgoFTP is a multiplatform, graphical FTP client using SWT. For this application QVM reported 6 sources of leaks. We introduced fixes to all of these leaks, and tested that the fixed version resolves them.

One source of a low-frequency leak in VirgoFTP is a pattern that repeats across many SWT applications. Changing the color/font preferences in an application often causes the leak of the previous colors/fonts used. This kind of leaks occur with such a low frequency, that programmers are probably choosing to ignore the disposal of resources in this case. Fixing this problem in VirgoFTP required significant refactoring of the code.

7.2 Overhead Evaluation

Methodology. For overhead measurements we use the SPECjvm98 and DaCapo benchmark suites.[6] For DaCapo, one run of the large inputs was measured. For SPECjvm98, which is shorter running, the benchmarks were configured to iterate for roughly one minute to create a reasonable usage scenario, and total time was measured. 20 runs of each benchmark were used to reduce noise.

We created a set of representative typestate properties that incur a significant overhead. We instrumented classes such as Java Collections, Enumerations, Vectors, and Streams.

Results. Figure 12 reports the overhead of the typestate monitoring client when applied to our benchmarks suite with a range of overhead budgets (5%, 10%, and 20%). The rightmost bar for each benchmark shows the overhead when the typestate client

[6]Dacapo version was dacapo-2006-10.jar. Jython and xalan were excluded from the study because they do not run properly on the developmental version of the VM used for this work (independent of the QVM modifications).

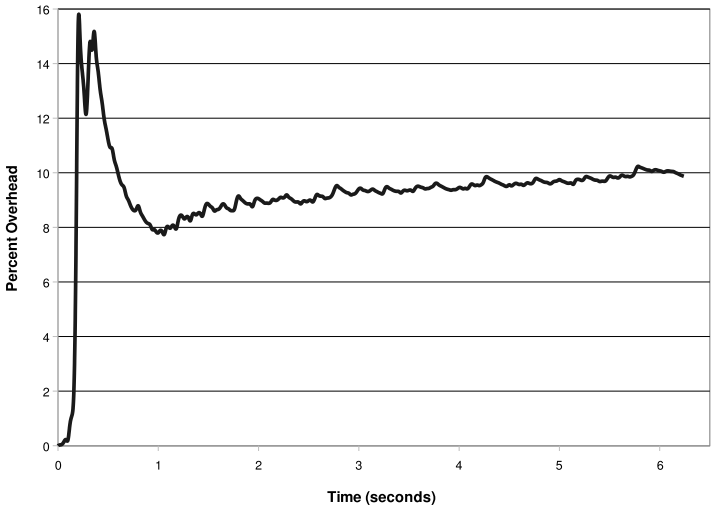

Fig. 13. Overhead over time with target overhead set at 10%.

is applied exhaustively, i.e., without sampling. The leftmost bar shows the base overhead, which represents the base checking overhead that is incurred when no sampling takes place (see Section 4.1).

The overhead incurred when checking these typestate properties exhaustively is high (up to 10x slowdown, with 7 of the benchmarks over 2x slowdown). Heavyweight properties that introduce frequent callbacks were selected intentionally to allow us to evaluate the effectiveness of the sampling infrastructure.

The base overhead (leftmost bar) is low, at most 2.5%. Having the base overhead be low is critical, as this is the overhead that is the lowest overhead that can be achieved when sampling is disabled.

The middle three bars show overhead incurred when QVM was run with a specific overhead budget. Although there is some fluctuation in the overhead achieved, it is generally quite close to the requested budget. Achieving accuracy at this level is quite challenging because the whole process takes place online and within a single execution of the benchmark. These results demonstrate not only the overhead monitor's ability to measure the overhead introduced, but the overhead controller's ability to keep the overhead close to the desired budget.

Figure 13 shows an example of the overhead manager adapting the overhead of the typestate client online for the javac benchmark and a 10% overhead budget. The x-axis shows time in seconds, and the y-axis shows percent overhead, as measured online by the QVM overhead monitor. The spike around 0.5 seconds occurs because there is some lag before the overhead monitor can react and reduce the sample rates. However, once the controller throttles the tagged objects at the hot allocation sites the overhead converges on the desired budget of 10%.

The goal of QVM is not just to have low overhead, but to collect as much useful information as possible within the overhead budget. The sampling strategy employed by the overhead manager (see Section 4.2) strives to distribute the samples across the allocation sites in the program, to help find bugs that may occur in cold code. Figure 14 compares the coverage of allocation sites achieved with 5% budget when using origin-specific sampling, as well as global sampling, where all sites are sampled equally. Origin-specific sampling enables nearly 100% coverage for all benchmarks,

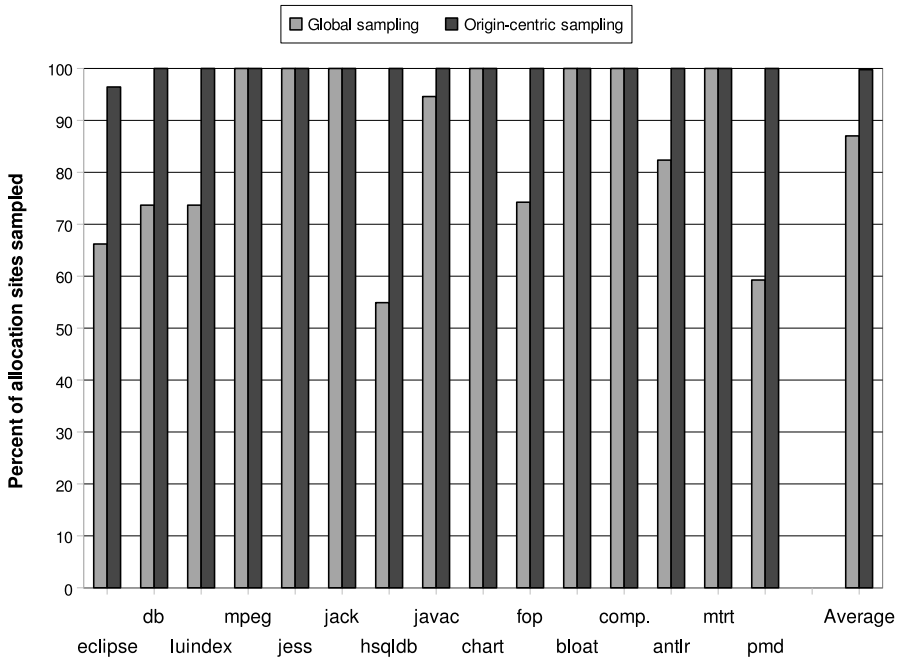

Fig. 14. Allocation Site Coverage: Percentage of allocation sites (of tracked types) that allocate at least one tracked object.

while global sampling misses a significant percentage of the allocation sites for at least half of the benchmarks.

QVM uses sampling to reduce overhead so there is no expectation that all objects will be tracked, however in many cases the sampling mechanism allows the dynamic number of tracked objects to be significantly higher than one might anticipate. In turn, this may lead to a higher-percentage of allocation sites being covered by our sampling. Table IV reports the percent of objects allocated (of the tracked types) that are sampled to be tracked by the typestate monitor.

Consider the program javac. Previously in Figure 12 we saw that our example set of typestate properties introduces overhead of around 970% when checked exhaustively. However, Table IV shows that with an overhead budget of 100% slowdown (more than a factor of 9 less than the exhaustive slowdown) 49% of the objects allocated (of tracked types) were still selected for tracking. This can be explained when a relatively small number of objects contribute significantly to the overhead; once sampling at these sites is throttled, the number of remaining allocations that can be tracked within the overhead budget may be large.

Some benchmarks (db, compress, bloat) report 100% for all overhead budgets because their exhaustive overhead for the typestate properties we selected is below 1% (see Figure 12).

7.3 Discussion

Wrapper Streams. For a large number of applications QVM reports violations of stream types that do not hold real resources but violate the contract of the `InputStream` and `OutputStream` API specification. An example that is widely reported by QVM is the `LEDataInputStream` from the package `swt.internal.image`. This stream is a wrapper around an `InputStream` and is often not closed because closing the wrapper closes the

Table IV.

Object Coverage: Percent of allocated objects (of tracked types) that are selected by QVM for typestate monitoring.

	Overhead Budget						
Benchmark	1%	2%	5%	10%	20%	50%	100%
db	100	100	100	100	100	100	100
mpegaudio	98	100	100	100	100	100	100
jess	63	76	85	87	95	100	100
jack	22	37	45	52	71	100	100
javac	0.4	1	4	9	31	41	49
compress	100	100	100	100	100	100	100
mtrt	39	46	66	83	90	93	94
antlr	13	19	34	68	67	92	98
eclipse	4	7	12	28	44	66	67
luindex	5	51	79	97	99	99	100
hsqldb	7	13	16	30	43	31	75
chart	40	64	85	88	93	94	97
fop	47	70	42	66	100	100	100
bloat	100	100	100	100	100	100	100
pmd	81	99	99	99	99	100	100

underlying InputStream. In many cases, the underlying stream outlives the wrapper stream and is therefore closed directly without ever invoking `close()` on the wrapper stream.

In addition, streams such as `ByteArrayInputStream` and `ByteArrayOutputStream` are simply wrappers around a byte array. Invoking `close` on such streams has no effect (although it is required by the streams API in principle), and programmers therefore avoid this redundant method call. We do not consider these to be real violations and do not include them in our QVM reports.

Library Objects vs. Application Objects. Our initial specification for SWT resources was not the one shown in Figure 7. Our initial specification required that `dispose()` be invoked on every SWT Widget, as this is the public method that an application code can invoke to dispose a resource. However, in SWT, widgets are arranged into an ownership structure in which a widget may have a parent that is responsible for its disposal. When the parent is disposed, it disposes all of its children, but instead of invoking the (public) method `dispose` to do so, it directly calls the (protected) internal method `release`. We therefore had to refine our specification to be aware of the internal library implementation and the fact that an SWT widget could be also released by an invocation of `release` that originates in library code.

Additional refinement of the specification is required to avoid objects that are allocated in the library for internal library use, and their lifetime is not managed (and should not be managed) by the application. For example, `Font` objects allocated by the static method `Font.gtk_new()` are managed by the library.

8. RELATED WORK

Practical Dynamic Analyses. Several tools such as Valgrind [Nethercote and Seward 2007] support dynamic analysis during development. These tools may allow a programmer to specify various clients (e.g. cachegrind). They are typically not used for production purposes as they have significant space and time overheads.

Aspects and Monitoring. Dynamic tools such as Tracematches [Allan et al. 2005], and MOP [Chen and Roşu 2007] are able to detect violation of typestate properties, and in particular detect resource leaks. For example, in Chen and Roşu [2007], JavaMOP was used to successfully detect a number of resource leaks in Eclipse. These tools extend aspect-oriented programming with the ability to specify declarative patterns against the history of the program, rather than against single events as in traditional aspects. Optimizing the performance of code generated from these declarative specifications is a challenging task and is currently an active area of research. Avgustinov et al. [2007a], the authors concentrate on dynamic optimizations that consider only the specified declarative pattern and not the program on which it is applied. Such optimizations include avoidance of memory leaks and better representation of the typestate automata. Alternatively, in Bodden et al. [2007], the authors take the program into account and perform static optimizations, for instance, removing unnecessary instrumentation points from the program. Unfortunately, despite these optimizations, there are cases where the overhead is still unacceptable for some properties. In Bodden et al. [2007], the authors propose two techniques: spatial and temporal partitioning. In the first optimization, assuming multiple users of the application, the instrumentation points are partitioned into sets optimizing the per-user overhead. However, it is still possible to partition the points in a way that some set has a hot point. The second optimization spawns a monitoring thread which can switch the instrumentation on and off at various times. The intervals defining when the point should be on or off are predetermined off-line and given to the thread as parameter. It seems that our approach of automatically adjusting the overhead online for a particular set of control sites will be beneficial to the second optimization.

Dwyer and Purandare [2007, 2008] show how dynamic typestate checking can leverage a preceding static typestate checking phase to reduce runtime costs. Their approach can be combined with QVM to further lower overhead and increase sampling coverage.

Sampling for Scalable Monitoring. Previous work has focused on low overhead techniques for sampling instrumentation [Arnold and Ryder 2001] and collecting such profiles in bursts [Chilimbi and Hirzel 2002]. However these techniques turn sampling on and off based on time or code execution frequency, and do not support a technique such as our object-centric sampling.

In Jump et al. [2004], profiling is limited to objects that are tagged at allocation time, an approach that is similar to our object-centric sampling, but is applied in the context of profiling for pre-tenuring.

In the cooperative bug isolation (CBI) project [Liblit 2007], the overhead of monitoring program execution is mitigated by using sparse random sampling and collecting information from a large number of users exercising the code. Collaborative techniques could be combined into QVM to collect application errors from a wider group of users. We believe that the ubiquity of QVM provides a natural channel for wider adoption of CBI-based techniques.

Typestate Verification and Static Leak Detection. A number of sound static tools target detection or prevention of memory and resource leaks [DeLine and Fahndrich 2001; DeLine and Fähndrich 2002; Fink et al. 2006, 2008; Foster et al. 2002; Heine and Lam 2003; Shaham et al. 2003]. Some tools specifically target detection of SWT resource leaks [Livshits 2005], and others target automatic generation of resource management code [Dillig et al. 2008]. In principle, most of these approaches are capable of detecting cases where an object is leaked or double disposed. In practice, however, these approaches do not scale to industrial-sized applications, and produce a large percentage

of false alarms. In addition, some of these approaches either require additional (potentially cumbersome) annotations or restrict the class of programs that may be written, e.g. by restricting aliasing [DeLine and Fahndrich 2001; Foster et al. 2002].

Heap Properties. Mitchell [2006] provides concise and informative summaries of real world heap graphs arising in production applications. The summaries are done offline and follow a set of useful heuristical patterns for summarizing graphs. In contrast, our goal is to check various user specified heap properties online. Subsequent work by Mitchell and Sevitsky [2007] study offline heap snapshots with the goal of finding inefficiencies in memory usage enforced by a particular program design.

Shacham et al. [2009] introduced CHAMELEON, a tool for profiling Java collections for finding inefficient use of collection in terms of space ("collection bloat"), and time. Their approach combines information from the garbage collector with collection usage statistics.

Chilimbi and Ganapathy [2006] provide a two-stage framework suitable for testing, where in the first stage a set of likely heap invariants based on node degree are computed at a small number of program points. Then the instrumented program is executed and checked against these invariants and a bug is reported if a deviation is observed.

Various works have relied on the garbage collector to find memory leaks. Jump and McKinley [2007] use the collector to help in suggesting potential leaks. Bond and McKinley [2006] study efficient leak detection for Java. Similarly to us, they make use of available bits in the object header and the adaptive profiling techniques from Hauswirth and Chilimbi [2004] applied on object use sites to reduce the space and time overheads. We see these advances as potential QVM clients, which could manage the overall overhead for them.

In recent work Aftandilian and Guyer [2008, 2009], use a sequential garbage collector to check several assertions specified by the programmer. Variants of two of the assertions proposed here, namely *isShared* and *isObjectOwned*, have been implemented in their system, albeit with different semantics. Our assertions are evaluated at the program point where they are issued, while in their work, assertions are evaluated when GC operates. Evaluating assertions only during GC allows to reduce the overhead of evaluation by batching all assertions up to that point in time. In contrast, we control the overhead of heap assertions using sampling. Further, our assertions are evaluated in parallel, enabling full use of the underlying multicore processor.

Current approaches typically study in detail a single useful property (such as leak detection) and modify the VM to support that property. In contrast, QVM can support clients checking many properties and automatically manage the performance overhead of such clients.

JVMTI. The current set of heap probes can also be implemented using JVMTI. However, even ignoring the already mentioned disadvantages of JVMTI, it is often difficult to implement what is required with a fixed set of functions without incurring significant overhead. For example, for our probe isShared() we avoid synchronization during traversal when parallel threads are used. In the case of JVMTI however, the notion of a parallel traversal thread is abstracted away (not seen at the callback level) and hence it would be required to synchronize internally to invoke the callback. Other probes such as isThreadOwned() requires the computation of the transitive closure of a given thread. Such a function is not provided in JVMTI (although the transitive closure of the object is provided), but can be simulated through other methods. Another problem are further optimizations such as concurrent processing of heap probes or filters on write barriers (to check if an escape bit in the object header is set, allowing the

avoidance of traversal). It is clear that if we are to support a language of core heap probes in the future, a flexible and high performance implementation is necessary, preferably realized through a specialized extendable interface at the virtual machine level.

9. FUTURE WORK

The overhead manager is a key component of the QVM. Our current implementation uses simple strategies that work well in practice, but do not guarantee any sort of optimality or enforce provable bounds. In the future, we plan to investigate how techniques from control theory can be used to provide a robust theoretical foundation for the overhead manager.

While our preliminary experience with heap assertions is promising, a thorough evaluation of these assertions is required on two aspects: (i) the appeal of heap assertions to programmers; (ii) the performance impact of heap assertions written in practice. We plan to address these questions in future work.

ACKNOWLEDGMENTS

We would like to thank Joshua Auerbach for helpful discussions on tracking overhead using per-thread metrics. We would also like to thank Noam Rinetzky for commenting on an earlier version of this manuscript.

REFERENCES

AFTANDILIAN, E. AND GUYER, S. Z. 2008. Gc assertions: Using the garbage collector to check heap properties. In *Proceedings of the ACM SIGPLAN workshop on Memory Systems Performance and Correctness (MSPC)*. ACM, 36–40.

AFTANDILIAN, E. E. AND GUYER, S. Z. 2009. GC assertions: Using the garbage collector to check heap properties. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'09)*. ACM, New York, NY, 235–244.

ALDRICH, J., KOSTADINOV, V., AND CHAMBERS, C. 2002. Alias annotations for program understanding. In *Proceedings of the 17th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'02)*. ACM, 311–330.

ALLAN, C., AVGUSTINOV, P., CHRISTENSEN, A. S., HENDREN, L., KUZINS, S., LHOTÁK, O., DE MOOR, O., SERENI, D., SITTAMPALAM, G., AND TIBBLE, J. 2005. Adding trace matching with free variables to Aspectj. In *Proceedings of the 20th Annual ACM SIGPLAN Conference on Object Oriented Programming, Systems, Languages, and Applications (OOPSLA'05)*. ACM, 345–364.

ARNOLD, M. AND RYDER, B. G. 2001. A framework for reducing the cost of instrumented code. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'01)*. ACM, New York, NY, 168–179.

ARNOLD, M. AND RYDER, B. G. 2002. Thin guards: A simple and effective technique for reducing the penalty of dynamic class loading. In *Proceedings of the 16th European Conference on Object-Oriented Programming*. B. Magnusson Ed., Lecture Notes in Computer Science Series, vol. 2374, 498–524.

ARNOLD, M., VECHEV, M., AND YAHAV, E. 2008. Qvm: An efficient runtime for detecting defects in deployed systems. In *Proceedings of the 23rd ACM SIGPLAN conference on Object-oriented programming systems languages and applications (OOPSLA'08)*. ACM, 143–162.

AUERBACH, J., BACON, D., CHENG, P., GROVE, D., BIRON, B., GRACIE, C., MCCLOSKEY, B., MICIC, A., AND SCIAMPACONE, R. 2008. Tax-and-spend: Democratic scheduling for real-time garbage collection. In *Proceedings of the International Conference on Embedded Software*. ACM, New York, NY.

AVGUSTINOV, P., TIBBLE, J., AND DE MOOR, O. 2007. Making trace monitors feasible. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object Oriented Programming Systems and Applications (OOPSLA'07)*. ACM, 589–608.

BAKER, H. G. 1995. 'Use-once' variables and linear objects - storage management, reflection and multithreading. *SIGPLAN Notices 30*, 1, 45–52.

BANERJEE, A. AND NAUMANN, D. A. 2005. Ownership confinement ensures representation independence for object-oriented programs. *J. ACM 52*, 6, 894–960.

BLACKBURN, S. M., GARNER, R., HOFFMAN, C., KHAN, A. M., MCKINLEY, K. S., BENTZUR, R., DIWAN, A., FEINBERG, D., FRAMPTON, D., GUYER, S. Z., HIRZEL, M., HOSKING, A., JUMP, M., LEE, H., MOSS, J.

E. B., PHANSALKAR, A., STEFANOVIĆ, D., VANDRUNEN, T., VON DINCKLAGE, D., AND WIEDERMANN, B. 2006. The DaCapo benchmarks: Java benchmarking development and analysis. In *Proceedings of the 21st Annual ACM SIGPLAN conference on Object-Oriented Programing, Systems, Languages, and Applications (OOPSLA'06)*. ACM, 169–190.

BODDEN, E., HENDREN, L. J., AND LHOTÁK, O. 2007. A staged static program analysis to improve the performance of runtime monitoring. In *Proceedings of the European Conference on Object-Oriented Programming*. 525–549.

BODDEN, E., HENDREN, L. J., LAM, P., LHOTÁK, O., AND NAEEM, N. A. 2007. Collaborative runtime verification with tracematches. In *Proceedings of the 7th International Workshop on Runtime Verification (RV)*. Lecture Notes in Computer Science, vol. 4839, 9–21.

BOND, M. D. 2008. Diagnosing and tolerating bugs in deployed systems. Ph.D. thesis, The University of Texas at Austin.

BOND, M. D. AND MCKINLEY, K. S. 2006. Bell: Bit-encoding online memory leak detection. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, New York, NY, 61–72.

BOND, M. D., NETHERCOTE, N., KENT, S. W., GUYER, S. Z., AND MCKINLEY, K. S. 2007. Tracking bad apples: Reporting the origin of null and undefined value errors. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-Oriented Programming Systems and Applications (OOPSLA'07)*. ACM, New York, NY, 405–422.

BOYAPATI, C., LEE, R., AND RINARD, M. 2002. Ownership types for safe programming: Preventing data races and deadlocks. In *Proceedings of the 17th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'02)*. ACM, New York, NY, 211–230.

CHEN, F. AND ROŞU, G. 2007. MOP: An efficient and generic runtime verification framework. In *Proceedings of the Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA'07)*.

CHILIMBI, T. M. AND GANAPATHY, V. 2006. Heapmd: Identifying heap-based bugs using anomaly detection. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 219–228.

CHILIMBI, T. M. AND HIRZEL, M. 2002. Dynamic hot data stream prefetching for general-purpose programs. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'02)*. ACM, New York, NY, 199–209.

CLARKE, D. G. 2003. Object ownership and containment. Ph.D. thesis, New South Wales, Australia.

DELINE, R. AND FAHNDRICH, M. 2001. Enforcing high-level protocols in low-level software. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'01)*. ACM Press, New York, NY, 59–69.

DELINE, R. AND FÄHNDRICH, M. 2002. Adoption and focus: Practical linear types for imperative programming. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'02)*. 13–24.

DILLIG, I., DILLIG, T., YAHAV, E., AND CHANDRA, S. 2008. The closer: Automating resource management in Java. In *Proceedings of the 7th International Symposium on Memory Management (ISMM'08)*. ACM, New York, NY, 1–10.

DWYER, M. B. AND PURANDARE, R. 2007. Residual dynamic typestate analysis exploiting static analysis: Results to reformulate and reduce the cost of dynamic analysis. In *Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering (ASE'07)*. ACM, New York, NY, 124–133.

DWYER, M. B. AND PURANDARE, R. 2008. Residual checking of safety properties. In *Proceedings of the 15th International SPIN Workshop on Model Checking Software*. Lecture Notes in Computer Science, vol. 5156, Springer, 1–2.

FINK, S. J. AND QIAN, F. 2003. Design, implementation and evaluation of adaptive recompilation with on-stack replacement. In *Proceedings of the International Symposium on Code Generation and Optimization (CGO'03)*. 241–252.

FINK, S., YAHAV, E., DOR, N., RAMALINGAM, G., AND GEAY, E. 2006. Effective typestate verification in the presence of aliasing. In *Proceedings of the International Symposium on Software Testing and Analysis (ISSTA'06)*. ACM Press, New York, NY, 133–144.

FINK, S. J., YAHAV, E., DOR, N., RAMALINGAM, G., AND GEAY, E. 2008. Effective typestate verification in the presence of aliasing. *ACM Trans. Softw. Eng. Methodol. 17*, 2.

FOSTER, J. S., TERAUCHI, T., AND AIKEN, A. 2002. Flow-sensitive type qualifiers. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'02)*. 1–12.

HAUSWIRTH, M. AND CHILIMBI, T. M. 2004. Low-overhead memory leak detection using adaptive statistical profiling. In *Proceedings of the 11th International Conference on Architectural Support for Programming Languages and Operating Systems*. ACM, 156–164.

HEINE, D. L. AND LAM, M. S. 2003. A practical flow-sensitive and context-sensitive c and c++ memory leak detector. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'03)*.

HOGG, J. 1991. Islands: Aliasing protection in object-oriented languages. In *Proceedings of the Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA'91)*. ACM, New York, NY, 271–285.

JUMP, M. AND MCKINLEY, K. S. 2007. Cork: Dynamic memory leak detection for garbage-collected languages. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'07)*. ACM, New York, NY, 31–38.

JUMP, M., BLACKBURN, S. M., AND MCKINLEY, K. S. 2004. Dynamic object sampling for pretenuring. In *Proceedings of the 4th International Symposium on Memory Management (ISMM'04)*. ACM, New York, NY, 152–162.

LAU, J., ARNOLD, M., HIND, M., AND CALDER, B. 2006. Online performance auditing: Using hot optimizations without getting burned. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'06)*. ACM, 239–251.

LEV-AMI, T. AND SAGIV, M. 2000. TVLA: A framework for Kleene based static analysis. In *Proceedings of the International Static Analysis Symposium*. Lecture Notes in Computer Science, vol. 1824, Springer-Verlag, 280–301.

LIBLIT, B. 2007. *Cooperative Bug Isolation*. Lecture Notes in Computer Science, vol. 4440, Springer.

LIVSHITS, V. B. 2005. Turning Eclipse against itself: Finding bugs in Eclipse code using lightweight static analysis. In *Eclipsecon'05 Research Exchange*.

MITCHELL, N. 2006. The runtime structure of object ownership. In *Proceedings of the European Conference on Object-Oriented Programming*. D. Thomas Ed., Lecture Notes in Computer Science, vol. 4067, Springer, 74–98.

MITCHELL, N. AND SEVITSKY, G. 2007. The causes of bloat, the limits of health. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object Oriented Programming Systems and Applications (OOPSLA'07)*. ACM, New York, NY, 245–260.

MÜLLER, P. AND RUDICH, A. 2007. Ownership transfer in universe types. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object Oriented Programming Systems and Applications (OOPSLA'07)*. ACM, New York, NY, 461–478.

NETHERCOTE, N. AND SEWARD, J. 2007. Valgrind: A framework for heavyweight dynamic binary instrumentation. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'06)*. ACM, New York, NY, 89–100.

RAMALINGAM, G., WARSHAVSKY, A., FIELD, J., GOYAL, D., AND SAGIV, M. 2002. Deriving specialized program analyses for certifying component-client conformance. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'02)*. ACM, New York, NY, 83–94.

RINETZKY, N., POETZSCH-HEFFTER, A., RAMALINGAM, G., SAGIV, M., AND YAHAV, E. 2007. Modular shape analysis for dynamically encapsulated programs. In *Proceedings of the European Symposium on Programming*. 220–236.

SHACHAM, O., VECHEV, M., AND YAHAV, E. 2009. Chameleon: Adaptive selection of collections. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'09)*.

SHAHAM, R., YAHAV, E., KOLODNER, E., AND SAGIV, M. 2003. Establishing local temporal heap safety properties with applications to compile-time memory management. In *Proceedings of the Static Analysis Symposium*.

SHOHAM, S., YAHAV, E., FINK, S., AND PISTOIA, M. 2007. Static specification mining using automata-based abstractions. In *Proceedings of the International Symposium on Software Testing and Analysis (ISSTA'07)*. ACM, 174–184.

SHOHAM, S., YAHAV, E., FINK, S. J., AND PISTOIA, M. 2008. Static specification mining using automata-based abstractions. *IEEE Trans. Softw. Engin. 34*, 5, 651–666.

STROM, R. E. AND YEMINI, S. 1986. Typestate: A programming language concept for enhancing software reliability. *IEEE Trans. Softw. Engin. 12*, 1, 157–171.

SUGANUMA, T., YASUE, T., KAWAHITO, M., KOMATSU, H., AND NAKATANI, T. 2001. A dynamic optimization framework for a Java just-in-time compiler. In *Proceedings of the 16th ACM SIGPLAN Conference*

on Object Oriented Programming, Systems, Languages, and Applications (OOPSLA'01). ACM, New York, NY, 180–195.

VECHEV, M., YAHAV, E., AND YORSH, G. 2010. Phalanx: Parallel checking of expressive heap assertions. In *Proceedings of the International Symposium on Memory Management (ISMM'10)*.

YAHAV, E. AND RAMALINGAM, G. 2004. Verifying safety properties using separation and heterogeneous abstractions. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'04)*. ACM Press, 25–34.

Received April 2009; revised September 2009, February 2010; accepted March 2010