



Engineering Record And Replay For Deployability

Robert O'Callahan and Chris Jones, *unaffiliated*; Nathan Froyd, *Mozilla Corporation*;
Kyle Huey, *unaffiliated*; Albert Noll, *Swisscom AG*; Nimrod Partush, *Technion*

<https://www.usenix.org/conference/atc17/technical-sessions/presentation/ocallahan>

This paper is included in the Proceedings of the
2017 USENIX Annual Technical Conference (USENIX ATC '17).

July 12–14, 2017 • Santa Clara, CA, USA

ISBN 978-1-931971-38-6

Open access to the Proceedings of the
2017 USENIX Annual Technical Conference
is sponsored by USENIX.

Engineering Record And Replay For Deployability

Robert O’Callahan*

Chris Jones*

Nathan Froyd
Mozilla Corporation

Kyle Huey*

Albert Noll*
Swisscom AG

Nimrod Partush*
Technion

Abstract

The ability to record and replay program executions with low overhead enables many applications, such as reverse-execution debugging, debugging of hard-to-reproduce test failures, and “black box” forensic analysis of failures in deployed systems. Existing record-and-replay approaches limit deployability by recording an entire virtual machine (heavyweight), modifying the OS kernel (adding deployment and maintenance costs), requiring pervasive code instrumentation (imposing significant performance and complexity overhead), or modifying compilers and runtime systems (limiting generality). We investigated whether it is possible to build a practical record-and-replay system avoiding all these issues. The answer turns out to be yes — if the CPU and operating system meet certain non-obvious constraints. Fortunately modern Intel CPUs, Linux kernels and user-space frameworks do meet these constraints, although this has only become true recently. With some novel optimizations, our system RR records and replays real-world low-parallelism workloads with low overhead, with an entirely user-space implementation, using stock hardware, compilers, runtimes and operating systems. RR forms the basis of an open-source reverse-execution debugger seeing significant use in practice. We present the design and implementation of RR, describe its performance on a variety of workloads, and identify constraints on hardware and operating system design required to support our approach.

1 Introduction

The ability to record a program execution with low overhead and play it back precisely has many applications [14, 15, 19] and has received significant attention in the research community. It has even been implemented in products such as VMware Workstation [28], Simics

[20], UndoDB [1] and TotalView [22]. Unfortunately, deployment of these techniques has been limited, for various reasons. Some approaches [17, 20, 28] require recording and replaying an entire virtual machine, which is heavyweight. Other approaches [6, 14, 26, 29] require running a modified OS kernel, hindering deployment and adding security and stability risk to the system. Requiring compiler and language runtime changes [29] also hinders deployment, especially when applications include their own JIT compilers. Some approaches [24, 30, 35] require custom hardware not yet available. Many approaches [1, 7, 22, 34] require pervasive instrumentation of code, which adds complexity and overhead, especially for self-modifying code (commonly used in polymorphic inline caching [25] and other implementation techniques in modern just-in-time compilers). A performant dynamic code instrumentation engine is also expensive to build and maintain.

We set out to build a system that maximizes deployability by avoiding all these issues: to record and replay unmodified user-space applications with stock Linux kernels, compilers, language runtimes, and x86/x86-64 CPUs, with a fully user-space implementation running without special privileges, and without using pervasive code instrumentation. We assume RR should run unmodified applications, and they will have bugs (including data races) that we wish to faithfully record and replay, but these applications will not maliciously try to subvert recording or replay. We combine techniques already known, but not previously demonstrated working together in a practical system: primarily, using `ptrace` to record and replay system call results and signals, avoiding non-deterministic data races by running only one thread at a time, and using CPU hardware performance counters to measure application progress so asynchronous signal and context-switch events are delivered at the right moment [33]. Section 2 describes our approach in more detail.

With that in place, we discovered the main perfor-

*Majority of work carried out while supported by Mozilla Research.

mance bottleneck for low-parallelism workloads was context switching induced by using `ptrace` to monitor system calls. We implemented a novel *in-process system-call interception* technique to eliminate those context switches, dramatically reducing recording and replay overhead on important real-world workloads. This optimization relies on modern Linux kernel features: `seccomp-bpf` to selectively suppress `ptrace` traps for certain system calls, and `perf` context-switch events to detect recorded threads blocking in the kernel. Section 3 describes this work, and Section 4 gives some performance results, showing that on important application workloads RR recording and replay slowdown is less than a factor of two.

We rely on hardware and OS features designed for other goals, so it is surprising that RR works. In fact, it skirts the edge of feasibility; in particular it cannot be implemented on ARM CPUs. Section 5 summarizes RR’s hardware and software requirements, which we hope will influence system designers.

RR is in daily use by many developers as the foundation of an efficient reverse-execution debugger that works on complex applications such as Samba, Firefox, Chromium, QEMU, LibreOffice and Wine. It is free software, available at <https://github.com/mozilla/rr>. This paper makes the following research contributions:

- We show that record and replay of user-space processes on modern, stock hardware and software without pervasive code instrumentation is possible and practical.
- We introduce an *in-process system-call interception* technique and show it dramatically reduces overhead.
- We show that for low-parallelism workloads, RR recording and replay overhead is reasonably low, lower than other approaches with comparable deployability.
- We identify hardware and operating system design constraints required to support our approach.

An extended technical report containing additional technical details and a retrospective on “lessons learned” during the development and use of RR is available [32].

2 Design

2.1 Summary

Most low-overhead record-and-replay systems depend on the observation that CPUs are mostly deterministic. We identify a boundary around state and computation,

record all sources of nondeterminism within the boundary and all inputs crossing into the boundary, and reexecute the computation within the boundary by replaying the nondeterminism and inputs. If all inputs and nondeterminism have truly been captured, the state and computation within the boundary during replay will match that during recording.

To enable record and replay of arbitrary Linux applications, without requiring kernel modifications or a virtual machine, RR records and replays the user-space execution of a group of processes. To simplify invariants, and to make replay as faithful as possible, replay preserves almost every detail of user-space execution. In particular, user-space memory and register values are preserved exactly, with a few exceptions noted later in the paper. This implies CPU-level control flow is identical between recording and replay, as is memory layout.

While replay preserves user-space state and execution, only a minimal amount of kernel state is reproduced during replay. For example, file descriptors are not opened, signal handlers are not installed, and filesystem operations are not performed. Instead the recorded user-space-visible effects of those operations, and future related operations, are replayed. We do create one replay thread per recorded thread (not strictly necessary), and we create one replay address space (i.e. process) per recorded address space, along with matching memory mappings.

With this design, our recording boundary is the interface between user-space and the kernel. The inputs and sources of nondeterminism are mainly the results of system calls, and the timing of asynchronous events.

2.2 Avoiding Data Races

With threads running on multiple cores, racing read-write or write-write accesses to the same memory location by different threads would be a source of nondeterminism. Therefore we take the common approach [17, 28, 1, 15] running only one thread at a time. RR preemptively schedules these threads, so context switch timing is nondeterminism that must be recorded. Data race bugs can still be observed if a context switch occurs at the right point in the execution (though bugs due to weak memory models cannot be observed).

This approach is much simpler and more deployable than alternatives [7, 18, 34, 39, 29], avoids assuming programs are race-free [14, 29], and is efficient for low-parallelism workloads. There is a large slowdown for workloads with a consistently high degree of parallelism; however, even for applications which are potentially highly parallel, users often apply RR to test workloads with relatively small datasets and hence limited parallelism.

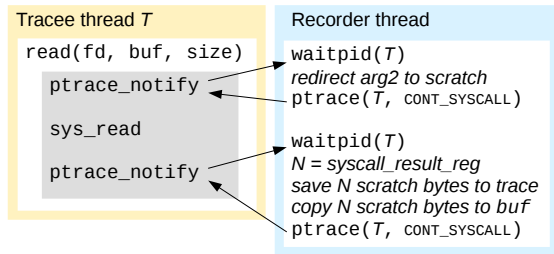


Figure 1: Recording simple system call

2.3 System Calls

System calls return data to user-space by modifying registers and memory, and these changes must be recorded. The `ptrace` system call allows a process to supervise the execution of other “tracee” processes and threads, and to be synchronously notified when a tracee thread enters or exits a system call. When a tracee thread enters the kernel for a system call, it is suspended and RR is notified. When RR chooses to run that thread again, the system call will complete, notifying RR again, giving it a chance to record the system call results. RR contains a model of most Linux system calls describing the user-space memory they can modify, given the system call input parameters and result.

As noted above, RR normally avoids races by scheduling only one thread at a time. However, if a system call blocks in the kernel, RR must try to schedule another application thread to run while the blocking system call completes. It’s possible (albeit unlikely) that the running thread could access the system call’s output buffer and race with the kernel’s writes to that buffer. To avoid this, we redirect system call output buffers to per-thread temporary “scratch memory” which is otherwise unused by the application. When we get a `ptrace` event for a blocked system call completing, RR copies scratch buffer contents to the real user-space destination(s) while no other threads are running, eliminating the race.

Figure 1 illustrates recording a simple `read` system call. The gray box represents kernel code.

During replay, when the next event to be replayed is an intercepted system call, we set a temporary breakpoint at the address of the system call instruction (recorded in the trace). We use `ptrace` to run the tracee thread until it hits the breakpoint, remove the breakpoint, advance the program counter past the system call instruction, and apply the recorded register and memory changes. This approach is simple and minimizes the number of context switches between RR and the tracee thread. (Occasionally it is unsafe and we fall back to a more complicated mechanism.)

Some system calls manipulate threads or address

spaces and require special handling during replay. For example a recorded `mmap` is replayed with `MAP_FIXED` to ensure the mapping is created at the correct address.

2.4 Asynchronous Events

We need to support two kinds of asynchronous events: preemptive context switches and signals. We treat the former as a special case of the latter, forcing a context switch by sending a signal to a running tracee thread. We need to ensure that during replay, a signal is delivered when the program is in exactly the same state as it was when the signal was delivered during recording.

As in previous work [17, 33, 10] we measure application progress using CPU hardware performance counters. Ideally we would count retired instructions leading up to an asynchronous event during recording, and during replay program the CPU to fire an interrupt after that many instructions have been retired — but this approach needs modifications to work in practice.

2.4.1 Nondeterministic Performance Counters

We require that every execution of a given sequence of user-space instructions changes the counter value by an amount that depends only on the instruction sequence, not system state invisible to user space (e.g. the contents of caches, the state of page tables, or speculative CPU state). This property (commonly described as “determinism” [40]) does not hold for most CPU performance counters in practice [17, 40]. For example it does not hold for any “instructions retired” counter on any known x86 CPU model (e.g. because an instruction triggering a page fault is restarted and counted twice).

Fortunately, modern Intel CPUs have exactly one deterministic performance counter: “retired conditional branches” (“RCB”), so we use that. We cannot just count the number of RCBs during recording and deliver the signal after we have executed that number of RCBs during replay, because the RCB count does not uniquely determine the execution point to deliver the signal at. Therefore we pair the RCB count with the complete state of general-purpose registers (including the program counter) to identify an execution point.

In general that *still* does not uniquely identify an execution point (e.g. consider the infinite loop `label: inc [global_var]; jmp label;`). However, in practice we have found it works reliably; code that returns to the same instruction with no intervening conditional branch must be very rare, and it only matters to RR if an asynchronous event occurs at such an instruction — in which case replay would probably diverge and fail.

2.4.2 Late Interrupt Firing

The other major problem is that, although CPUs can be programmed to fire an interrupt after a specified number of performance events have been observed, the interrupt does not fire immediately. In practice we often observe it firing after dozens more instructions have retired. To compensate for this, during replay, we program the interrupt to trigger some number of events earlier than the actual RCB count we are expecting. Then we set a temporary breakpoint at the program counter value for the state we're trying to reach, and repeatedly run to the breakpoint until the RCB count and the general-purpose register values match their recorded values.

2.5 Shared Memory

By scheduling only one thread at a time, RR avoids issues with races on shared memory as long as that memory is written only by tracee threads. It is possible for recorded processes to share memory with other processes, and even kernel device drivers, where that non-recorded code can perform writes that race with accesses by tracee threads. Fortunately, this is rare for applications running in common Linux desktop environments, occurring in only four common cases: applications sharing memory with the PulseAudio daemon, applications sharing memory with the X server, applications sharing memory with kernel graphics drivers and GPUs, and `vdso` syscalls. We avoid the first three problems by automatically disabling use of shared memory with PulseAudio and X (falling back to a socket transport in both cases), and disabling direct access to the GPU from applications.

`vdso` syscalls are a Linux optimization that implements some common read-only system calls (e.g. `gettimeofday`) entirely in user space, partly by reading memory shared with the kernel and updated asynchronously by the kernel. We disable `vdso` syscalls by patching their user-space implementations to perform the equivalent real system call instead.

Applications could still share memory with non-recorded processes in problematic ways, though this is rare in practice and can often be solved just by enlarging the scope of the group of processes recorded by RR.

2.6 Nondeterministic Instructions

Almost all CPU instructions are deterministic, but some are not. One common nondeterministic x86 instruction is `RDTSC`, which reads a time-stamp counter. This particular instruction is easy to handle, since the CPU can be configured to trap on an `RDTSC` and Linux exposes this via a `prctl` API, so we can trap, emulate and record each `RDTSC`.

Other relatively recent x86 instructions are harder to handle. `RDRAND` generates random numbers and hopefully is not deterministic. We have only encountered it being used in one place in GNU `libstdc++`, so RR patches that explicitly. `XBEGIN` and associated instructions support hardware transactional memory. These are nondeterministic from the point of view of user space, since a hardware transaction can succeed or fail depending on CPU cache state. Fortunately so far we have only found these being used by the system `pthread` library, and we dynamically apply custom patches to that library to disable use of hardware transactions.

The `CPUID` instruction is mostly deterministic, but one of its features returns the index of the running core, which affects behavior deep in `glibc` and can change as the kernel migrates a process between cores. We use the Linux `sched_setaffinity` API to force all tracee threads to run on a particular fixed core, and also force them to run on that core during replay.

We could easily avoid most of these issues in well-behaved programs if we could just trap-and-emulate the `CPUID` instruction, since then we could mask off the feature bits indicating support for `RDRAND`, hardware transactions, etc. Modern Intel CPUs support this (“`CPUID` faulting”); we are in the process of adding an API for this to Linux.

2.7 Reducing Trace Sizes

For many applications the bulk of their input is memory-mapped files, mainly executable code. Copying all executables and libraries to the recorded trace on every execution would impose significant time and space overhead. RR creates hard links to memory-mapped executable files instead of copying them; as long as a system update or recompile replaces executables with new files, instead of writing to the existing files, the links retain the old file data. This works well in practice.

Even better, modern filesystems such as XFS and Btrfs offer copy-on-write logical copies of files (and even block ranges within files), ideal for our purposes. When a mapped file is on the same filesystem as the recorded trace, and the filesystem supports cloning, RR clones mapped files into the trace. These clone operations are essentially free in time and space, until/unless the original file is modified or deleted.

RR compresses all trace data, other than cloned files and blocks, with the `zlib` “deflate” method.

With these optimizations, in practice trace storage is a non-issue. Section 4.4 presents some results.

2.8 Other Details

Apart from those major issues, many other details are required to build a complete record-and-replay system, too many to mention here. Some system calls (e.g. `execve`) are especially complex to handle. Recording and replaying signal delivery are complex, partly because signal delivery has poorly-documented side effects on user-space memory. Advanced Linux kernel features such as `unshare` (kernel namespaces) and `seccomp` require thoughtful handling. Many of these details are interesting, but they do not impact the overall approach.

3 In-process System-call Interception

The approach described in the previous section works, but overhead is disappointingly high (see Figure 5 below). The core problem is that for every tracee system call, as shown in Figure 1 the tracee performs four context switches: two blocking `ptrace` notifications, each requiring a context switch from the tracee to RR and back. For common system calls such as `gettimeofday` or `read` from cached files, the cost of even a single context switch dwarfs the cost of the system call itself. To significantly reduce overhead, we must avoid context-switches to RR when processing these common system calls.

Therefore, we inject into the recorded process a library that intercepts common system calls, performs the system call without triggering a `ptrace` trap, and records the results to a dedicated buffer shared with RR. RR periodically flushes the buffer to its trace. The concept is simple but there are problems to overcome.

3.1 Intercepting System Calls

A common technique for intercepting system calls in-process is to use dynamic linking to interpose wrapper functions over the C library functions that make system calls. In practice, we have found that method to be insufficient, due to applications making direct system calls, and fragile, due to variations in C libraries, and applications that require their own preloading [37, 3]).

Instead, when the tracee makes a system call, RR is notified via a `ptrace` trap and it tries to rewrite the system-call instruction to call into our interception library. This is tricky because on x86 a system call instruction is two bytes long, but we need to replace it with a five-byte `call` instruction. In practice, frequently executed system call instructions are followed by a few known, fixed instruction sequences; for example, many system call instructions are followed by a `cmp $0xffffffff, %eax` instruction testing the `syscall` result. We added five hand-written stubs to our interception

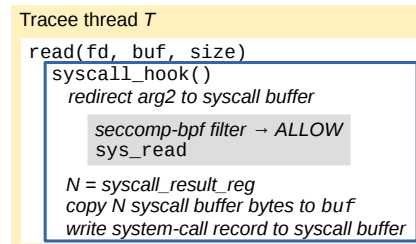


Figure 2: Recording with system-call interception

library that execute post-system-call instructions before returning to the patched code. On receipt of a `ptrace` system-call notification, RR replaces the system call instruction and its following instruction with a call to the corresponding stub.

We (try to) redirect all system call instructions to the interception library, but for simplicity it only contains wrappers for the most common system calls, and for others it falls back to doing a regular `ptrace`-trapping system call.

3.2 Selectively Trapping System Calls

`ptrace` system-call monitoring triggers traps for all system calls, but our interception library needs to avoid traps for selected system calls. Fortunately, modern Linux kernels support selectively generating `ptrace` traps: `seccomp-bpf`. `seccomp-bpf` was designed primarily for sandboxing. A process can apply a `seccomp-bpf` filter function, expressed in bytecode, to another process; then, for every system call performed by the target process, the kernel runs the filter, passing in incoming user-space register values, including the program counter. The filter's result directs the kernel to either allow the system call, fail with a given `errno`, kill the target process, or trigger a `ptrace` trap. Overhead of filter execution is negligible since filters run directly in the kernel and are compiled to native code on most architectures.

Figure 2 illustrates recording a simple `read` system call with in-process system-call interception. The solid-border box represents code in the interception library and the grey box represents kernel code.

RR injects a special page of memory into every tracee process at a fixed address (immediately after `execve`). That page contains a system call instruction — the “untraced instruction”. RR applies a `seccomp-bpf` filter to each recorded process that triggers a `ptrace` trap for every system call — except when the program counter is at the untraced instruction, in which case the call is allowed. Whenever the interception library needs to make an untraced system call, it uses that instruction.

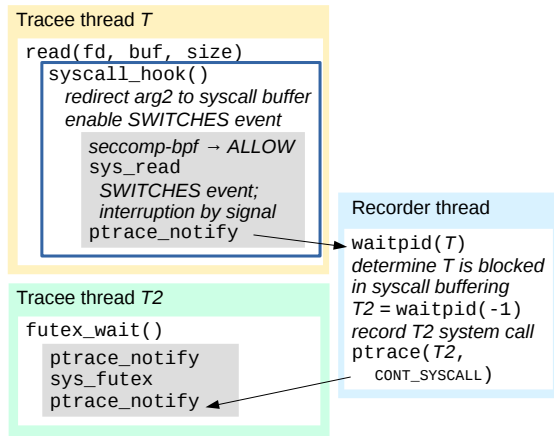


Figure 3: Recording a blocking system call

3.3 Detecting Blocked System Calls

Some common system calls sometimes block (e.g. `read` on an empty pipe). Because RR runs tracee threads one at a time, if a thread enters a blocking system call without notifying RR, it will hang and could cause the entire recording to deadlock (e.g. if another tracee thread is about to `write` to the pipe). We need the kernel to notify RR and suspend the tracee thread whenever an untraced system call blocks, to ensure we can schedule a different tracee thread.

We do this using the Linux `perf` event system to monitor `PERF_COUNT_SW_CONTEXT_SWITCHES`. The kernel raises one of these events every time it deschedules a thread from a CPU core. The interception library monitors these events for each thread and requests that the kernel send a signal to the blocked thread every time the event occurs. These signals trigger `ptrace` notifications to RR while preventing the thread from executing further. To avoid spurious signals (e.g. when the thread is descheduled due to normal timeslice expiration), the event is normally disabled and explicitly enabled during an untraced system call that might block. Still, spurious `SWITCHES` can occur at any point between enabling and disabling the event; we handle these edge cases with careful inspection of the tracee state.

Figure 3 illustrates recording a blocking `read` system call with system-call interception. The kernel deschedules the thread, triggering a `perf` event which sends a signal to the thread, *rescheduling* it, interrupting the system call, and sending a `ptrace` notification to the recorder. The recorder does bookkeeping to note that an intercepted system call was interrupted in thread *T*, then checks whether any tracee threads in blocking system calls have progressed to a system-call exit and generated a `ptrace` notification. In this example *T2* has

completed a (not intercepted) blocking `futex` system call, so we resume executing *T2*.

Resuming an intercepted system call that was interrupted by a signal (e.g. *T*'s `read` call in Figure 3) is more complicated. Explaining that requires understanding Linux system call restart semantics, which are too complicated to explain in the space available here.

3.4 Handling Replay

Conceptually, during recording we need to copy system call output buffers to a trace buffer, and during replay we need to copy results from the trace buffer to system call output buffers. This is a problem because the interception library is part of the recording and replay and therefore should execute the same code in both cases. (Previous work with user-level system call interception [11, 36, 21, 29] avoided these problems by having less strict goals for replay fidelity.)

For this reason (and to avoid races of the sort discussed in Section 2.3), the interception library redirects system call outputs to write directly to the trace buffer. After the system call completes, the interception library copies the output data from the trace buffer to the original output buffer(s). During replay the untraced system call instruction is replaced with a no-op, so the system call does not occur; the results are already present in the trace buffer so the post-system-call copy from the trace buffer to the output buffer(s) does what we need.

During recording, each untraced system call sets a result register and the interception library writes it to the trace buffer. Replay must read the result register from the trace buffer instead. We use a conditional move instruction so that control flow is perfectly consistent between recording and replay. The condition is loaded from an `is_replay` global variable, so the register holding the condition is different over a very short span of instructions (and explicitly cleared afterwards).

Handling “in-out” system call memory parameters is tricky. During recording we copy the input buffer to the trace buffer, pass the system call a pointer to the trace buffer, then copy the trace buffer contents back to the input buffer. Performing that first copy during replay would overwrite the trace buffer values holding the system call results, so during replay we turn that copy into a no-op using a conditional move to set the source address copy to the destination address.

We could allow replay of the interception library to diverge further from its recorded behavior, but that would have to be done very carefully. We’d have to ensure the RCB count was identical along both paths, and that register values were consistent whenever we exit the interception library or trap to RR within the interception library. It’s simplest to minimize the divergence.

3.5 Optimizing Reads With Block Cloning

When an input file is on the same filesystem as the recorded trace and the filesystem supports copy-on-write cloning of file blocks, for large block-aligned reads the system call interception code clones the data to a per-thread “cloned-data” trace file, bypassing the normal system-call recording logic. This greatly reduces space and time overhead for file-read-intensive workloads; see the next section.

This optimization works by cloning the input blocks and then reading the input data from the original input file. This opens up a possible race: between the clone and the read, another process could overwrite the input file data, in which case the data read during replay would differ from the data read during recording, causing replay to fail. However, when a file read races with a write under Linux, the reader can receive an arbitrary mix of old and new data, so such behavior would almost certainly be a severe bug, and in practice such bugs do not seem to be common. The race could be avoided by reading from the cloned-data file instead of the original input file, but that performs very poorly because it defeats Linux’s read-ahead optimizations (since the data in the cloned-data file is never available until just before it’s needed).

4 Results

4.1 Workloads

Benchmarks were chosen to illuminate RR’s strengths and weaknesses, while also containing representatives of real-world usage. They were tuned to fit in system memory (to minimize the impact of I/O on test results), to run for about 30 seconds each (except for *cp* where a 30s run time would require it to not fit in memory).

cp duplicates a git checkout of *glibc* (revision 2d02fd07) using `cp -a` (15200 files constituting 732MB of data, according to `du -h`). *cp* is single-threaded, making intensive use of synchronous reads and a variety of other filesystem-related system calls.

make builds DynamoRio [8] (version 6.1.0) with `make -j8` (`-j8` omitted when restricting to a single core). This tests potentially-parallel execution of many short-lived processes.

octane runs the Google Octane benchmark under the Mozilla Spidermonkey Javascript engine (Mercurial revision 9bd900888753). This illustrates performance on CPU-intensive code in a complex language runtime.

htmltest runs the Mozilla Firefox HTML forms tests (Mercurial revision 9bd900888753). The harness is excluded from recording (using `mach mochitest -f plain --debugger RR dom/html/test/forms`). This is an example

from real-world usage. About 30% of user-space CPU time is in the harness.

sambatest runs a Samba (git revision 9ee4678b) UDP echo test via `make test TESTS=samba4.echo.udp`. This is an example from real-world usage.

All tests run on a Dell XPS15 laptop with a quad-core Intel Skylake CPU (8 SMT threads), 16GB RAM and a 512GB SSD using Btrfs in Fedora Core 23 Linux.

4.2 Overhead

Table 1 shows the wall-clock run time of various configurations, normalized to the run time of the baseline configuration. *octane* is designed to run for a fixed length of time and report a score, so we report the ratio of the baseline score to the configuration-under-test score — except for replay tests, where the reported score will necessarily be the same as the score during recording. For *octane* replay tests we report the ratio of the baseline score to the recorded score, multiplied by the ratio of replay run time to recording run time. Each test was run six times, discarding the first result and reporting the geometric mean of the other five results. Thus the results represent warm-cache performance.

“Single core” reports the overhead of just restricting all threads to a single core using Linux `taskset`.

“Record no-intercept” and “Replay no-intercept” report overhead with in-process system-call interception disabled (which also disables block cloning). “Record no-cloning” reports overhead with just block cloning disabled.

“DynamoRio-null” reports the overhead of running the tests under the DynamoRio [8] (version 6.1.0) “null tool”, to estimate a lower bound for the overhead of using dynamic code instrumentation as an implementation technique. (DynamoRio is reported to be among the fastest dynamic code instrumentation engines.)

4.3 Observations

Overhead on *make* is significantly higher than for the other workloads. Forcing *make* onto a single core imposes major slowdown. Also, *make* forks and execs 2430 processes, mostly short-lived. (The next most prolific workload is *sambatest* with 89.) In-process system-call interception only starts working in a process once the interception library has been loaded, but at least 80 system calls are performed before that completes, so its effectiveness is limited for short-lived processes.

Figure 4 shows the overall recording and replay overhead for workloads other than *make*. Error bars in figures show 95% confidence intervals; these results are highly stable across runs.

Workload	Baseline duration	Record	Replay	Single core	Record no-intercept	Replay no-intercept	Record no-cloning	DynamoRio-null
cp	1.04s	1.49×	0.72×	0.98×	24.53×	15.39×	3.68×	1.24×
make	20.99s	7.85×	11.93×	3.36×	11.32×	14.36×	7.84×	10.97×
octane	32.73s	1.79×	1.56×	1.36×	2.65×	2.43×	1.80×	crash
htmltest	23.74s	1.49×	1.01×	1.07×	4.66×	3.43×	1.50×	14.03×
sambatest	31.75s	1.57×	1.23×	0.95×	2.23×	1.74×	1.57×	1.43×

Table 1: Run-time overhead

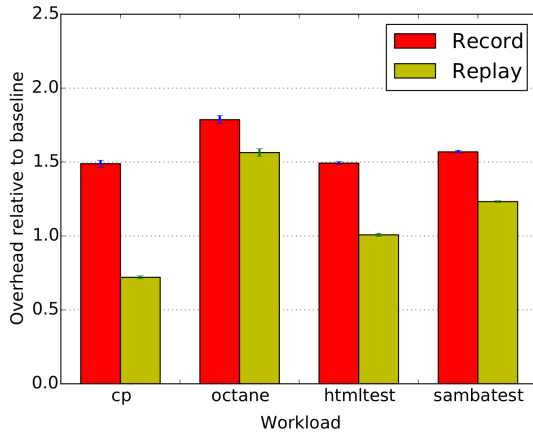


Figure 4: Run-time overhead excluding *make*

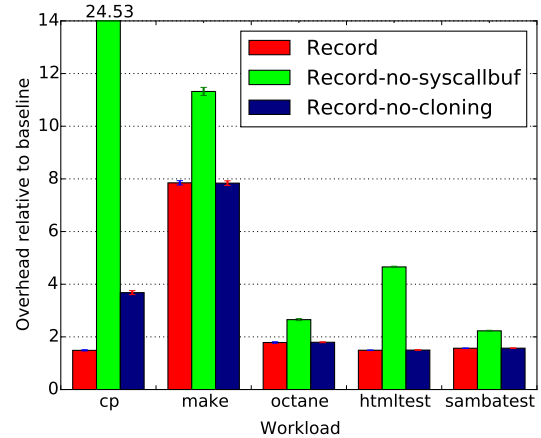


Figure 5: Impact of optimizations

Excluding *make*, RR’s recording slowdown is less than a factor of two. Excluding *make*, RR’s replay overhead is lower than its recording overhead. Replay can even be faster than normal execution, in *cp* because system calls do less work. For interactive applications, not represented here, replay can take much less time than the original execution because idle periods are eliminated.

octane is the only workload here other than *make* making significant use of multiple cores, and this accounts for the majority of RR’s overhead on *octane*.

Figure 5 shows the impact of system-call interception and blocking cloning on recording. The system-call interception optimization produces a large reduction in recording (and replay) overhead. Cloning file data blocks is a major improvement for *cp* recording but has essentially no effect on the other workloads.

Figure 6 compares RR recording overhead with DynamoRio’s “null tool”, which runs all code through the DynamoRio instrumentation engine but does not modify the code beyond whatever is necessary to maintain supervised execution; this represents a minimal-overhead code instrumentation configuration. DynamoRio crashed on *octane*¹. *cp* executes very little user-space code and

¹We reported DynamoRio’s crash on our “octane” workload to the developers at <https://github.com/DynamoRIO/dynamorio/issues/1930>.

DynamoRio’s overhead is low on that workload. On *make* and *sambatest* DynamoRio overhead is similar to RR recording, even though on *make* DynamoRio can utilize multiple cores. On *htmltest* DynamoRio’s overhead is very high, possibly because that test runs a lot of Javascript with dynamically generated and modified machine code. Implementing record-and-replay on top of dynamic instrumentation would incur significant additional overhead, so we would expect the resulting system to have significantly higher overhead than RR.

4.4 Storage Space Usage

RR traces contain three kinds of data: cloned (or hard-linked) files used for memory-map operations, cloned file blocks, and all other trace data, especially event metadata and the results of general system calls.

Memory-mapped files are mostly the executables and libraries loaded by traces. While the original files are not changed or removed, which is usually true in practice, their clones take no additional space and require no data writes. RR makes no attempt to consolidate duplicate file clones, so most traces contain many duplicates and reporting meaningful space usage for these files is both difficult and unimportant in practice. The same is true for cloned file blocks.

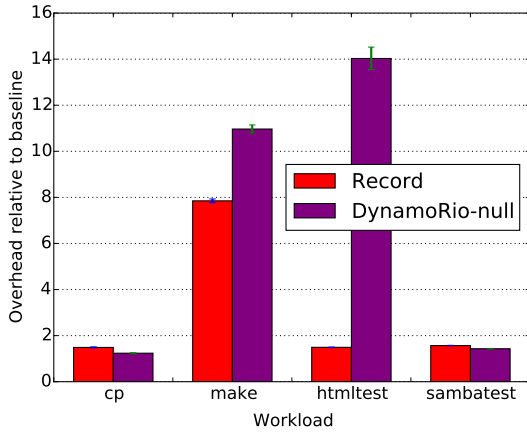


Figure 6: Comparison with DynamoRio-null

Workload	Compressed MB/s	deflate ratio	Cloned blocks MB/s
cp	19.03	4.87×	586.14
make	15.82	8.32×	5.50
octane	0.08	8.33×	0.00
htmltest	0.79	5.94×	0.00
sambatest	6.85	21.87×	0.00

Table 2: Storage space usage

Table 2 shows the storage usage of each workload, in MB/s, for general trace data and cloned file blocks. We compute the geometric mean of the data usage for each trace and divide by the run-time of the workload baseline configuration. Space consumption shows very little variation between runs.

Different workloads have highly varying space consumption rates, but several MB/s is easy for modern systems to handle. In real-world usage, trace storage has not been a concern.

4.5 Memory Usage

Table 3 shows the memory usage of each workload. Every 10ms we sum the proportional-set-size (“PSS”) values of all workload processes (including RR if running); we determine the peak values for each run and take their geometric mean. In Linux, each page of memory mapped into a process’s address space contributes $1/n$ pages to that process’s PSS, where n is the number of processes mapping the page; thus it is meaningful to sum PSS values over processes which share memory. The same data are shown in Figure 7. In the figure, the fraction of PSS used by the RR process is shown in orange. Memory usage data was gathered in separate runs from the timing data shown above, to ensure the overhead of gathering

Workload	Baseline	Record	Replay	Single core
cp	0.51	34.54	9.11	0.51
make	510.51	327.19	314.16	288.65
octane	419.47	610.48	588.01	392.95
htmltest	690.81	692.06	324.71	689.75
sambatest	298.68	400.49	428.79	303.03

Table 3: Memory usage (peak PSS MB)

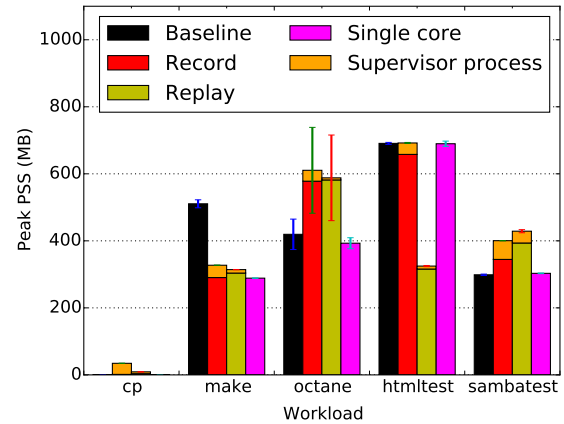


Figure 7: Memory usage

memory statistics did not impact those results.

Given these experiments ran on an otherwise unloaded machine with 16GB RAM and all data fits in cache, none of these workloads experienced any memory pressure. *cp* uses almost no memory. In *make*, just running on a single core reduces peak PSS significantly because not as many processes run simultaneously. In *octane* memory usage is volatile (highly sensitive to small changes in GC behavior) but recording significantly increases application memory usage; recording also increases application memory usage a small amount in *sambatest* but slightly decreases it in *htmltest*. (We expect a small increase in application memory usage due to system-call interception and scratch buffers.) These effects are difficult to explain due to the complexity of the applications, but could be due to changes in timing and/or effects on application or operating system memory management heuristics.

Replay memory usage is similar to recording except in *htmltest*, where it’s dramatically lower because we’re not replaying the test harness.

RR’s memory overhead is not an issue in practice.

5 Hardware/Software Design Constraints

We summarize the hardware and software features RR depends on, for system designers who may be interested in supporting RR-like record-and-replay.

5.1 Hardware

As discussed in Section 2.4.1, RR requires a “deterministic” hardware performance counter to measure application progress. The ideal performance counter for our purposes would count the exact number of instructions retired as observed in user-space (e.g., counting an interrupted-and-restarted instruction once). Virtual machines should support reliable performance-counter virtualization. Currently RR works under KVM and VMware, but VMware’s *VM exit clustering* optimization [4], as implemented, breaks the determinism of the RCB counter and must be manually disabled.

Some x86 CPU instructions are nondeterministic. Section 2.6 discusses our current workarounds for this. Exposing hardware support for trapping CPUID is important for long-term control over these instructions.

We would like to support record-and-replay of programs using hardware transactional memory (XBEGIN/XEND). It would suffice if hardware and the OS could be configured to raise a signal on any failed transaction.

Trapping on all other nondeterministic instructions (e.g. RDRAND) would be useful.

Porting RR to ARM failed because all ARM atomic memory operations use the “load-linked/store-conditional” approach, which is inherently nondeterministic. The conditional store can fail because of non-user-space-observable activity, e.g. hardware interrupts, so counts of retired instructions or conditional branches for code performing atomic memory operations are nondeterministic. These operations are inlined into very many code locations, so it appears patching them is not feasible except via pervasive code instrumentation or compiler changes. On x86(-64), atomic operations (e.g. compare-and-swap) are deterministic in terms of user-space state, so there is no such problem.

5.2 Software

As noted in Section 2.5, RR depends on configuring applications to avoid sharing memory with non-recorded processes.

We described how RR performance depends on modern Linux features: `seccomp-bpf` to selectively trap system calls, `PERF_COUNT_SW_CONTEXT_SWITCHES` performance events to handling blocking system calls, and copy-on-write file and block cloning APIs to reduce I/O overhead.

Efficient record-and-replay depends on clearly identifying a boundary within which code is replayed deterministically, and recording and replaying the timing and contents of all inputs into that boundary. In RR, that boundary is mostly the interface between the ker-

nel and user-space. This suits Linux: most of the Linux user/kernel interface is stable across OS versions, relatively simple and well-documented, and it’s easy to count hardware performance events occurring within the boundary (i.e. all user-space events for a specific process). This is less true in other operating systems. For example, in Windows, the user/kernel interface is not publicly documented, and is apparently more complex and less stable than in Linux. Implementing and maintaining the RR approach for Windows would be considerably more challenging than for Linux, at least for anyone other than the OS vendor.

6 Related Work

6.1 Whole-System Replay

ReVirt [17] was an early project that recorded and replayed the execution of an entire virtual machine. VMware [28] used the same approach to support record-and-replay debugging in VMware Workstation, for a time, but discontinued the feature. The full-system simulator Simics supports reverse-execution debugging via deterministic reexecution [20]. There have been efforts to add some record-and-replay support to QEMU [15, 16, 38] and Xen [18, 10]. Whole-system record-and-replay can be useful, but it is often inconvenient to hoist the application into a virtual machine. Many applications of record-and-replay require cheap checkpointing, and checkpointing a VM image is generally more expensive than checkpointing one or a few processes.

6.2 Replaying User-Space With Kernel Support

Scribe [26], dOS [6] and Arnold [14] replay a process or group of processes by extending the OS kernel with record-and-replay functionality. Kernel changes make maintenance and deployment more difficult — unless record-and-replay is integrated into the base OS. But adding invasive new features to the kernel has risks, so if record-and-replay can be well implemented outside the kernel, moving it into the kernel may not be desirable.

6.3 Pure User-Space Replay

Pure user-space record-and-replay systems have existed since at least MEC [11], and later Jockey [36] and liblog [21]. Those systems did not handle asynchronous event timing and other OS features. PinPlay [34], iDNA [7], UndoDB [1] and TotalView ReplayEngine [22] use code instrumentation to record and replay asynchronous event timing. Unlike UndoDB and RR, PinPlay and iDNA instrument all loads, thus supporting parallel recording in

the presence of data races and avoiding having to compute the effects of system calls, but this gives them higher overhead than the other systems. Compared to the other systems that support asynchronous events, RR achieves lower overhead by avoiding code instrumentation.

6.4 Higher-Level Replay

Record-and-replay features have been integrated into language-level virtual machines. DeJaVu [12] added record-and-replay capabilities to the Jalapeño Java VM. Microsoft IntelliTrace [2] instruments CLR bytecode to record high-level events and the parameters and results of function calls; it does not produce a full replay. Systems such as Chronon [13] for Java instrument bytecode to collect enough data to provide the appearance of replaying execution, without actually doing a replay. Dolos [9] provides record-and-replay for JS applications in Webkit by recording and replaying nondeterministic inputs to the browser. R2 [23] provides record-and-replay by instrumenting library interfaces; handling data races or asynchronous events requires user effort to isolate the nondeterminism. Such systems are all significantly narrower in scope than the ability to replay general user-space execution.

6.5 Parallel Replay

Recording application threads running concurrently on multiple cores, with the possibility of data races, with low overhead, is extremely challenging. PinPlay [34] and iDNA/Nirvana [7] instrument shared-memory loads and report high overhead. SMP-ReVirt [18] tracks page ownership using hardware page protection and reports high overhead on benchmarks with a lot of sharing. DoublePlay [39] runs two instances of the application and thus has high overhead when the application alone could saturate available cores. ODR [5] has low recording overhead but replay can be extremely expensive and is not guaranteed to reproduce the same program states. Castor [29] instruments synchronization code by modifying compilers and runtime systems, which creates barriers to easy deployment, and cannot replay reliably in the presence of data races.

The best hope for general, low-overhead parallel recording seems to be hardware support. Projects such as FDR [41], BugNet [31], Rerun [24], DeLorean [30] and QuickRec [35] have explored low-overhead parallel recording hardware.

7 Future Work

RR perturbs execution, especially by forcing all threads onto a single core, and therefore can fail to reproduce

bugs that manifest outside RR. We have addressed this problem by introducing a “chaos mode” that intelligently adds randomness to scheduling decisions, enabling us to reproduce many more bugs, but that work is beyond the scope of this paper. There are many more opportunities to enhance the recorder to find more bugs.

Putting record-and-replay support in the kernel has performance benefits, e.g. reducing the cost of recording context switches. We may be able to find reusable primitives that can be added to kernels to improve the performance of user-space record-and-replay while being less invasive than a full kernel implementation.

Recording multiple processes running in parallel on multiple cores seems feasible if they do not share memory — or, if they share memory, techniques inspired by SMP-ReVirt [18], dthreads [27] or Castor [29] may work for some workloads.

The applications of record-and-replay are perhaps more interesting and important than the base technology. For example, one can perform high-overhead dynamic analysis during replay [14, 15, 34], potentially parallelized over multiple segments of the execution. With RR’s no-instrumentation approach, one could collect performance data such as sampled stacks and performance counter values during recording, and correlate that data with rich analysis generated during replay (e.g. cache simulation). Always-on record-and-replay would make finding and fixing bugs in the field much easier. Demonstrating compelling applications for record-and-replay will build the case for building support into commodity hardware and software.

8 Conclusions

The current state of Linux on commodity x86 CPUs enables single-core user-space record-and-replay with low overhead, without pervasive code instrumentation — but only just. This is fortuitous; we use software and hardware features for purposes they were not designed to serve. It is also a recent development; five years ago `seccomp-bpf` and the Linux file cloning APIs did not exist, and commodity architectures with a deterministic hardware performance counter usable from user-space had only just appeared (Intel Westmere)². By identifying the utility of these features for record-and-replay, we hope that they will be supported by an increasingly broad range of future systems. By providing an open-source, easy-to-deploy, production-ready record-and-replay framework we hope to enable more compelling applications of this technology.

²Performance counters have been usable for kernel-implemented replay [17, 33] for longer, because kernel code can observe and compensate for events such as interrupts and page faults.

References

- [1] Reversible debugging tools for C/C++ on Linux & Android. <http://undo-software.com>. Accessed: 2016-04-16.
- [2] Understanding IntelliTrace part I: What the @\$% is IntelliTrace? <https://blogs.msdn.microsoft.com/zainnab/2013/02/12/understanding-intellitrace-part-i-what-the-is-intellitrace>. Accessed: 2016-04-16.
- [3] Wine windows-on-posix framework. <https://www.winehq.org>. Accessed: 2016-09-20.
- [4] O. Agesen, J. Mattson, R. Rugina, and J. Sheldon. Software techniques for avoiding hardware virtualization exits. In *Proceedings of the 2012 USENIX Annual Technical Conference*, June 2012.
- [5] G. Altekar and I. Stoica. ODR: Output-deterministic replay for multicore debugging. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*, October 2009.
- [6] T. Bergan, N. Hunt, L. Ceze, and S. D. Gribble. Deterministic process groups in dOS. In *Proceedings of the 9th USENIX Symposium on Operating Systems Design and Implementation*, October 2010.
- [7] S. Bhansali, W.-K. Chen, S. de Jong, A. Edwards, R. Murray, M. Drinić, D. Mihočka, and J. Chau. Framework for instruction-level tracing and analysis of program executions. In *Proceedings of the 2nd International Conference on Virtual Execution Environments*, June 2006.
- [8] D. Bruening, Q. Zhao, and S. Amarasinghe. Transparent dynamic instrumentation. In *Proceedings of the 8th International Conference on Virtual Execution Environments*, March 2012.
- [9] B. Burg, R. Bailey, A. Ko, and M. Ernst. Interactive record/replay for web application debugging. In *Proceedings of the 26th ACM Symposium on User Interface Software and Technology*, October 2013.
- [10] A. Burtsev, D. Johnson, M. Hibler, E. Eide, and J. Regehr. Abstractions for practical virtual machine replay. In *Proceedings of the 12th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, April 2016.
- [11] M. E. Chastain. <https://lwn.net/1999/0121/a/mec.html>, January 1999. Accessed: 2016-04-16.
- [12] J.-D. Choi, B. Alpern, T. Ngo, M. Sridharan, and J. Vlissides. A perturbation-free replay platform for cross-optimized multithreaded applications. In *Proceedings of the 15th International Parallel and Distributed Processing Symposium*, April 2001.
- [13] P. Deva. <http://chrononsystems.com/blog/design-and-architecture-of-the-chronon-record-0>, December 2010. Accessed: 2016-04-16.
- [14] D. Devecsery, M. Chow, X. Dou, J. Flinn, and P. M. Chen. Eidetic systems. In *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation*, October 2014.
- [15] B. Dolan-Gavitt, J. Hodosh, P. Hulin, T. Leek, and R. Whelan. Repeatable reverse engineering with PANDA. In *Proceedings of the 5th Program Protection and Reverse Engineering Workshop*, December 2015.
- [16] P. Dovgalyuk. Deterministic replay of systems execution with multi-target QEMU simulator for dynamic analysis and reverse debugging. 2012.
- [17] G. Dunlap, S. King, S. Cinar, M. A. Basrai, and P. M. Chen. ReVirt: Enabling intrusion analysis through virtual-machine logging and replay. In *Proceedings of the 5th USENIX Symposium on Operating Systems Design and Implementation*, December 2002.
- [18] G. W. Dunlap, D. G. Lucchetti, M. A. Fetterman, and P. M. Chen. Execution replay of multiprocessor virtual machines. In *Proceedings of the 4th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, March 2008.
- [19] J. Engblom. A review of reverse debugging. In *System, Software, SoC and Silicon Debug Conference*, September 2012.
- [20] J. Engblom, D. Aarno, and B. Werner. Full-system simulation from embedded to high-performance systems. In *Processor and System-on-Chip Simulation*, 2010.
- [21] D. Geels, G. Altekar, S. Shenker, and I. Stoica. Replay debugging for distributed applications. In *Proceedings of the 2006 USENIX Annual Technical Conference*, June 2006.
- [22] C. Gottbrath. Reverse debugging with the TotalView debugger. In *Cray User Group Conference*, May 2008.

- [23] Z. Guo, X. Wang, J. Tang, X. Liu, Z. Xu, M. Wu, M. F. Kaashoek, and Z. Zhang. R2: An application-level kernel for record and replay. In *Proceedings of the 8th USENIX Symposium on Operating Systems Design and Implementation*, December 2008.
- [24] D. Hower and M. Hill. Rerun: Exploiting episodes for lightweight memory race recording. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*, June 2008.
- [25] U. Hlzle, C. Chambers, and D. Ungar. Optimizing dynamically-typed object-oriented languages with polymorphic inline caches. In *Proceedings of the 1991 European Conference on Object-Oriented Programming*, July 1991.
- [26] O. Laadan, N. Viennot, and J. Nieh. Transparent, lightweight application execution replay on commodity multiprocessor operating systems. In *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems*, June 2010.
- [27] T. Liu, C. Curtsinger, and E. Berger. Dthreads: Efficient deterministic multithreading. In *Proceedings of the ACM SIGOPS 23rd Symposium on Operating Systems Principles*, October 2011.
- [28] V. Malyugin, J. Sheldon, G. Venkitachalam, B. Weissman, and M. Xu. ReTrace: Collecting execution trace with virtual machine deterministic replay. In *Proceedings of the Workshop on Modeling, Benchmarking and Simulation*, June 2007.
- [29] A. J. Mashtizadeh, T. Garfinkel, D. Terei, D. Mazières, and M. Rosenblum. Towards practical default-on multi-core record/replay. In *Proceedings of the 22nd International Conference on Architectural Support for Programming Languages and Operating Systems (to appear)*, April 2017.
- [30] P. Montesinos, L. Ceze, and J. Torrellas. DeLorean: Recording and deterministically replaying shared-memory multiprocessor execution efficiently. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*, June 2008.
- [31] S. Narayanasamy, G. Pokam, and B. Calder. Bugnet: Continuously recording program execution for deterministic replay debugging. In *Proceedings of the 32nd Annual International Symposium on Computer Architecture*, June 2005.
- [32] R. O’Callahan, C. Jones, N. Froyd, K. Huey, A. Noll, and N. Partush. Engineering record and replay for deployability: Extended technical report. <http://arxiv.org/abs/1705.05937>.
- [33] M. Olszewski, J. Ansel, and S. Amarasinghe. Kendo: Efficient deterministic multithreading in software. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems*, March 2009.
- [34] H. Patil, C. Pereira, M. Stallcup, G. Lueck, and J. Cownie. PinPlay: A framework for deterministic replay and reproducible analysis of parallel programs. In *Proceedings of the 8th Annual IEEE/ACM International Symposium on Code Generation and Optimization*, April 2010.
- [35] G. Pokam, K. Danne, C. Pereira, R. Kassa, T. Kranich, S. Hu, J. Gottschlich, N. Honarmand, N. Dautenhahn, S. King, and J. Torrellas. Quick-Rec: Prototyping an intel architecture extension for record and replay of multithreaded programs. In *Proceedings of the 40th Annual International Symposium on Computer Architecture*, June 2013.
- [36] Y. Saito. Jockey: A user-space library for record-replay debugging. In *Proceedings of the 6th International Symposium on Automated Analysis-driven Debugging*, September 2005.
- [37] K. Serebryany, D. Bruening, A. Potapenko, and D. Vyukov. Addresssanitizer: A fast address sanity checker. In *Proceedings of the 2012 USENIX Annual Technical Conference*, June 2012.
- [38] D. Srinivasan and X. Jiang. Time-traveling forensic analysis of VM-based high-interaction honeypots. In *Security and Privacy in Communication Networks: 7th International ICST Conference*, September 2011.
- [39] K. Veeraraghavan, D. Lee, B. Wester, J. Ouyang, P. M. Chen, J. Flinn, and S. Narayanasamy. Doubleplay: Parallelizing sequential logging and replay. In *Proceedings of the 16th International Conference on Architectural Support for Programming Languages and Operating Systems*, March 2011.
- [40] V. Weaver, D. Terpstra, and S. Moore. Non-determinism and overcount on modern hardware performance counter implementations. In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software*, April 2013.
- [41] M. Xu, R. Bodik, and M. D. Hill. A “flight data recorder” for enabling full-system multiprocessor deterministic replay. In *Proceedings of the 30th Annual International Symposium on Computer Architecture*, June 2003.

