

Dynamic Race Detection with LLVM Compiler

Compile-time instrumentation for ThreadSanitizer

Konstantin Serebryany, Alexander Potapenko, Timur Iskhodzhanov, and
Dmitriy Vyukov

Google LLC, 7 Balchug st., Moscow, 115035, Russia
{kcc,glider,timurrrr,dvyukov}@google.com

Abstract. Data races are among the most difficult to detect and costly bugs. Race detection has been studied widely, but none of the existing tools satisfies the requirements of high speed, detailed reports and wide availability at the same time. We describe our attempt to create a tool that works fast, has detailed and understandable reports and is available on a variety of platforms. The race detector is based on our previous work, ThreadSanitizer [1], and the instrumentation is done using the LLVM compiler. We show that applying compiler instrumentation and sampling reduces the slowdown to less than 1.5x, fast enough to use instrumented programs interactively.

1 Introduction

Recently the growth of CPU frequencies has transformed into the growth of the number of cores per CPU. As a result, multithreaded code became more popular on desktops, and concurrency bugs, especially *data races*, became more frequent. The classical approach to dynamic race detection assumes that program code is instrumented and program events are passed to an analysis algorithm [8, 11]. Some of the publicly available race detectors for native code [7, 1, 12] use *run-time instrumentation*. There are also tools that use *compiler instrumentation* [3, 6, 10], but none is publicly available on most popular operating systems.

In [1] we described ThreadSanitizer (TSan-Valgrind), a dynamic race detector for native code based on *run-time instrumentation*. The tool has found hundreds of harmful races in a number of C++ programs at Google, including some in the Chromium browser [4]. Significant slowdown remains the largest problem of ThreadSanitizer: for many tests we observed 5x–30x slowdown due to the complex race detection algorithm; on heavy web applications the slowdowns were even greater (50x and more) because of the underlying translation system (Valgrind, [12])¹. Another problem with Valgrind is that it serializes all threads; with multicore machines this becomes a serious limitation. Finally, Valgrind is not available on some platforms we are interested in (entirely unavailable on Windows, hard to deploy on ChromiumOS).

¹ Mainly because Valgrind had to execute much single-threaded JavaScript code.

In this paper we present TSan-LLVM, a dynamic race detector that uses *compile-time instrumentation* based on a widely available LLVM compiler² [9]. The new tool shares the race detection logic with ThreadSanitizer, but has greater speed and portability. Our work resembles LiteRace [10] (both use compiler instrumentation and sampling, the performance figures are comparable), but the significant advantages of our tool are the more precise race detection algorithm [1], the granularity of sampling and public availability.

2 Compiler instrumentation

The compiler instrumentation is implemented as a pass for the LLVM compiler. The resulting object files are linked against our runtime library.

2.1 Runtime library

As opposed to a number of popular race detection algorithms [11, 12, 10], ThreadSanitizer [1] tracks both locksets and the happens-before relation. This allows it to switch between the pure happens-before mode, which reports no false positives, but may miss potential bugs, and the hybrid mode, which finds more potential races, but may give false reports. In both modes the tool reports the call stacks of all the accesses constituting the race, along with the locks taken and the origin of memory involved. This is vital in order to give all the necessary information to the tool users.

The algorithm is basically a state machine – it receives program events, updates the internal state and, when appropriate, reports a potential race. The major events handled by the state machine are: READ, WRITE (memory accesses); SIGNAL, WAIT (happens-before events); LOCK, UNLOCK (locking events).

The runtime library provides entry points for the instrumented code, keeps all the information about the running program (e.g. the location and size of thread stacks and thread-local storage) and generates the events by wrapping the functions that are of interest for the race detector: synchronization primitives and thread manipulation routines, memory allocation routines, other functions that imply happens-before relations in the real world programs (e.g. `read()/write()`), and dynamic annotations [1].

2.2 Instrumentation

The instrumentation is done at the LLVM IR level. For each translation unit the following steps are done:

Call stack instrumentation. In order to report nearly precise contexts for all memory accesses that constitute a race, ThreadSanitizer has to maintain a correct call stack for every thread at all times. We keep a per-thread stack with

² We have also made an instrumentation plugin for GCC, but do not describe it here due to the limited space.

a pointer to its top; the stack is updated at every function entry and exit, as well as at every basic block start³.

To keep the call stack consistent, the tool also needs to intercept `setjmp()` and instrument the LLVM `invoke` instruction to roll back the stack pointer when necessary. This is not done yet, because these features are rarely used at Google.

Memory access instrumentation. Each memory access event is a tuple of 5 attributes: *thread id*, *ADDR*, *PC*, *isWrite*, *size*. The last three are statically known. Memory accesses that happen in one basic block⁴ are grouped together; for each block the compiler module creates a *passport* – an array of tuples representing each memory access. Every memory access is instrumented with the code that records the effective address of the access into a thread-local buffer. The buffer contents are processed by the ThreadSanitizer state machine [1] at the end of each block.

2.3 Sampling

In order to decrease the runtime overhead even more, we’ve experimented with sampling the memory accesses. We exploit the *cold-region hypothesis* [10]: data races are more likely to occur in cold regions of well-tested programs, because the races in hot regions either have been already found and fixed or are benign.

The technique we use for sampling is similar to that suggested in LiteRace [10]: ThreadSanitizer adapts the thread-local sampling rate per code region such that the sampling rate decreases logarithmically with the total number of executions of a particular region. Unlike in LiteRace, the instrumented code is always executed and the memory access addresses are put into the buffer, which is then either processed or ignored depending on the value of the execution counter. Another difference from LiteRace is that we apply sampling to smaller regions (basic blocks or superblocks, as opposed to whole functions), which allows to find races in cold regions of hot functions with higher probability.

2.4 Limitations and further improvements

The compiler-based instrumentation has some disadvantages over the run-time instrumentation: the races in the code which was not re-compiled with the instrumentation enabled (system libraries, JIT-ed code) will be missed, the tool usage is less convenient since it requires a custom build⁵. As we show in the next section, the benefit of much higher speed outweighs these limitations for our use cases.

Much could be done to decrease the overhead even further by reducing the number of instrumented memory accesses without losing races. A promising direction is to use compiler’s static analysis to skip accesses that never escape the current thread. Another optimization is to instrument only one of the accesses to the same memory location on the same path.

³ Optimizations may apply.

⁴ We also extend this approach to handle larger acyclic regions of code (superblocks).

⁵ Valgrind-based tools also usually require a custom build to avoid false positives.

3 Results

To estimate the performance of our tool, we ran it on two Chromium tests and a synthetic microbenchmark. We’ve already used TSan-Valgrind to test Chromium (see [1]) and were able to compare the results and assess the benefits of the compile-time instrumentation approach for a real-world application. `cross_fuzz` [5] is a cross-document DOM binding fuzzer that is known to stress the browser and reveal complex bugs, including races. `net_unittests` [4] is a set of nearly 2000 test cases that test various networking features and create many threads. The third test we ran just calls a simple non-inlined function⁶ many times:

```
void IncrementMe(int *x) { (*x)++; }
```

One variant of the test is single-threaded, the other variant spawns 4 threads that access separate memory regions. The measurements were done on an HP Z600 machine (2 quad-core Intel Xeon E5620 CPUs, 12G RAM).

Table 1 contains execution times for uninstrumented binaries run natively and under TSan-Valgrind compared to the instrumented binaries tested in two modes: with full memory access analysis (TSan-LLVM, sampling disabled) and with race detection disabled (TSan-LLVM-null, an empty stub is called at the end of each block). We’ve also measured run times under Intel Inspector XE [7], Memcheck⁷ and Helgrind version 3.6.1 [12]. The comparison shows that TSan-LLVM outperforms TSan-Valgrind by 1.7x–2.9x on the big tests. TSan-LLVM does not instrument `libc` and other system libraries, but we estimate their performance impact to be within 2%–3%.

Table 1. TSan-LLVM compared to other tools. Time in seconds.

tool	<code>cross_fuzz</code>	<code>net_unittests</code>	synthetic, 1 thread	synthetic, 4 threads
native run	71.6	87	0.9	0.9
Memcheck	1275	991	33	133
Inspector XE	failed	1064	130	480
Helgrind	failed	2529	40	154
TSan-Valgrind	325.2	592	49	191
TSan-LLVM	190.9	206	15.5	17
TSan-LLVM-null	78.6	119	2	2.1

Table 2 shows how the performance depends on the sampling parameter (a number k which means that the tool starts ignoring some memory accesses after executing the region 2^{32-k} times). Using the sampling value of 20 is 1.5x–2x faster than without sampling on the chosen benchmarks. In this mode the slowdown compared to the native run is less than 1.5x, and the tool is still capable

⁶ Part of `racecheck_unittest` [2], a test suite for data race detectors.

⁷ Memcheck, the Valgrind memory error detector, does different kind of instrumentation and can not ignore JavaScript, but its figures may still serve as a data point.

of finding a number of known races. We found over 15 races in Chromium while running `cross_fuzz` with TSan-Valgrind; these races (except one, which happens in a system library) are also detectable with TSan-LLVM, without sampling and even with sampling value 20.

Table 2. TSan-LLVM performance with various sampling values.

test name	sampling parameter	-	10	20	30
<code>cross_fuzz</code>	time, sec	190.9	142.3	94.5	78.1
	accesses analyzed, %	100.0	77.8	16.2	3.6
<code>net_unittests</code>	time, sec	206	190	134	117
	accesses analyzed, %	100.0	33.7	14.1	13.4

4 Conclusions

We present a dynamic race detector based on low-level compiler instrumentation. This detector has a large speed advantage (1.7x–2.9x on the real-world applications) over our previous Valgrind-based tool, and a slowdown factor of 2.5x (less than 1.5x, if sampling is used), which is fast enough to run interactive UI tests on the instrumented Chromium browser. The achieved speedup can be improved even further if additional compile-time static analysis is employed.

References

1. Serebryany, K., Iskhodzhanov, T.: ThreadSanitizer: data race detection in practice. WBIA (2009)
2. ThreadSanitizer project: documentation, source code, dynamic annotations, unit tests, <http://code.google.com/p/data-race-test>
3. Sun Studio, <http://developers.sun.com/sunstudio>
4. Chromium browser, <http://dev.chromium.org>
5. Cross Fuzz, http://lcamtuf.coredump.cx/cross_fuzz
6. Duggal, A.: Stopping Data Races Using Redflag. Master’s thesis, Stony Brook University (May 2010), technical Report FSL-10-02
7. Intel Inspector XE, <http://software.intel.com/en/articles/intel-parallel-studio-xe>
8. Lamport, L.: Time, clocks, and the ordering of events in a distributed system. *Commun. ACM* (1978)
9. The LLVM Compiler Infrastructure, <http://llvm.org>
10. Marino, D., Musuvathi, M., Narayanasamy, S.: Literace: effective sampling for lightweight data-race detection. In: PLDI (2009)
11. Savage, S., Burrows, M., et al.: Eraser: a dynamic data race detector for multi-threaded programs. *ACM TOCS* 15(4), 391–411 (1997)
12. Valgrind, Helgrind., <http://www.valgrind.org>