

KUDA: GPU Accelerated Split Race Checker

Position paper

U. Can Bekar

Koç University, İstanbul, Turkey
ucbekar@ku.edu.tr

Tayfun Elmas

University of California, Berkeley,
USA
elmas@eecs.berkeley.edu

Semih Okur

University of Illinois at Urbana
Champaign, Urbana, USA
okur2@illinois.edu

Serdar Tasiran

Koç University, İstanbul, Turkey
stasiran@ku.edu.tr

Abstract

We propose a novel approach for runtime verification on computers with a large number of computation cores, without any hardware extension to mainstream PC environment. The goal of the approach is making use of all hardware resources to decouple the computational overhead of traditional race checkers via parallelizing the runtime verification. We distinguish between two kinds of computational overhead: (i) overhead caused by monitoring, and (ii) overhead due to the verification algorithm(s). So far, runtime verification algorithms have been designed to run on the same threads as the code being monitored and both (i) and (ii) contribute to the slowdown of the program being monitored. The framework we propose allows us to carry out (ii) on separate, dedicated cores and threads. As a result, the program being monitored only experiences slowdown due to (i) and plus the communication of the events to be monitored. There exists related work showing that with some inexpensive hardware support (i) can be reduced to negligible levels. By parallelizing analyses our experiments show that they run as fast as the program being monitored, but on separate computational resources, one can potentially use this approach for monitoring, error detection, containment, and recovery from errors. As a demonstration of concept, we investigate runtime monitoring for concurrency bugs, in particular, data race detection. We use a few CPU threads and a large number of cores on a GPU to minimize the slowdown of the application on which race detection is being run.

1. Introduction

We propose an approach to make use of some of the computation cores and other hardware resources in a computer to monitor the programs running on other cores for concurrency errors, to contain and/or recover from these errors, if not immediately, shortly after they take place. While exploring such an approach, we had two goals: (i) to have minimal, tolerable impact on the threads being monitored, and (ii) to

have the monitoring algorithms work at the same speed as the program, while possibly lagging behind by a bounded amount. The rationale behind the first goal is to enable efficient, even post-deployment use of the monitoring and bug-detection algorithms for safety-critical systems. The rationale behind the second goal is to make it possible to contain concurrency errors, notify the threads that have experienced the errors, and gracefully shut down the program or to recover from the error. One result of the second goal is to force the monitoring framework to parallelize the event logging and analysis algorithms as much as possible.

In our approach, the application being monitored and the actual monitoring code run on separate processing units/resources. They communicate with each other using some shared memory or message passing. The instrumented application code only has the additional responsibility of communicating relevant events to the monitoring code. The monitoring and runtime analysis code can be quite complex, but runs on separate processors and is parallelized, thus, the application performance is not affected by the runtime analyses being performed. We conjecture that the performance penalty on the application being monitored due to instrumentation and communication of relevant events can be reduced to negligible levels, for example, using inexpensive hardware support such as hardware-assisted message passing [4]. The goal is for the monitoring code to run at least the same speed as the application being monitored, but lag behind by a very small delay due to event communication. (In our experiments this delay was in milliseconds.) This makes possible scenarios in which, in response to errors detected, the application has shut down gracefully, or a previous valid checkpoint is restored or the application is restarted.

As a demonstration of concept, we investigate runtime monitoring for concurrency bugs, in particular, data race detection. Since systems with hundreds of cores are not yet available as mainstream, to investigate the feasibility of our proposal, we use a few CPU threads/cores to carry out the

efficient concurrent transfer of logged events from the CPU cores to a Graphics Processing Unit (GPU), and we use the GPU to run our race detection algorithm. Today’s GPUs provide a highly parallel, multithreaded, computation environment with hundreds of processor cores and a higher memory bandwidth than CPUs. Thus, our framework allows us to investigate opportunities for efficiently performing various kinds of runtime analyses on highly parallel computing environments.

We provide a framework that instruments binaries so that the application threads log the interesting events in a central event list. The analysis threads then work off of this event list to perform possibly expensive but parallelized analyses. For this, we use carefully designed non-blocking algorithms and block-based handling of the event list for efficient recording of the events in this list. In particular, we communicate the events to the GPU for processing in fixed-size segments called *frames*. We accomplish fast, highly parallelized runtime analysis on a GPU with hundreds of cores by exploring algorithms that can check each event frame independently from other frames. Our experience was that this can be done without significantly affecting the soundness of the checking. Long-enough frames allow the analyses to catch all errors that can be caught by analyzing the entire execution. Since the computational cost of the analysis threads does not affect application performance, in this highly parallel setting, one can achieve a lower performance impact while still not sacrificing from precision.

For demonstration, we adapted the well known ERASER and GOLDLOCKS algorithms for data race checking, so that they can be parallelized to run on a large number of threads and cores on the GPU. Surprisingly, this high parallelism has a simplifying effect on the algorithm implementation. Since we have many threads/cores, the algorithm can be written to make each thread or core to perform a local and independent check (for a single memory access) without having to worry about sharing or interaction with other threads. Each thread creates only the necessary data structures for the check and discards/reuses them after the check completes; this avoids the need for memory management and sharing of complicated algorithm-specific data structures.

We implemented our proposed system in a tool called KUDA. KUDA is open source and available at <http://kuda.codeplex.com>. We use the Pin [5] library to instrument binaries for monitoring and the CUDA [7] library to run our analysis algorithms on the GPU. We applied KUDA to a number of multithreaded programs from the PARSEC and SPLASH-2 benchmark suites. We performed experiments using CPU and GPU implementations of the ERASER and GOLDLOCKS algorithms. We chose ERASER to represent a cheap (although imprecise) algorithm, while GOLDLOCKS served as a representative precise, higher-complexity algorithm. We contrasted two approaches: (i) a straightforward implementation of ERASER running on the same threads and cores as the application, and (ii) an implementation of GOLDLOCKS using our framework, where the checking threads are decoupled and run on the GPU. Over-

all, our early experimental results indicate that our approach is promising. Using a cheaper race detection algorithm using the traditional approach as exemplified by (i) causes about twenty times more slowdown compared to a more complex race-detection algorithm implemented in our approach!

2. Challenges in runtime monitoring

The purpose of this section is to point out the key challenges one is likely to face when building a runtime verification tool for concurrency-related errors. In order to make our ideas concrete we will work out this section using a well-known data-race detection algorithm.

2.1 Example: Eraser algorithm for data-race detection

ERASER is a well-known lockset-based algorithm for detecting race conditions dynamically [9]. For simplicity of presentation, we focus on the core algorithm using only locksets without distinguishing between read and write accesses.

A *race condition* occurs if two different threads perform conflicting accesses (i.e., at least one of them is a write) on a shared (global) variable and there is no proper synchronization between these accesses. In order to detect race conditions, ERASER enforces the locking discipline that every shared variable x is protected by a common lock throughout the execution.

The Eraser algorithm maintains a lockset $LS(x)$ representing the set of the algorithm’s guess of the locks protecting x . It also maintains for each thread t , a lockset $LH(t)$ representing the set of locks held by thread t at a given point in an execution. $LH(t)$ is updated appropriately when thread t acquires and releases a lock. The algorithm attempts to infer the actual protecting locks for each data variable x by initializing $LS(x)$ to the set of all locks in the program and then updating $LS(x)$ to be the intersection $LH(t) \cap LS(x)$ at each access to x by a thread t .

2.2 Cost of runtime monitoring on the CPU

The implementation of the ERASER algorithm requires 1) to monitor the events in an execution that the algorithm needs to keep track of, which is usually done by instrumenting the program’s source or binary code, and 2) to perform some computation to update the algorithm-specific data structures, i.e., the maps LH and LS , and check some conditions, i.e., “Does $LS(x)$ become empty?”. In 1), events are either immediately communicated to the algorithm by running a callback function to perform 2), or saved to a temporary buffer to be processed later (e.g., in a linked list of events as in [2]). There are two main sources of runtime cost, which combined together contribute highly to the overhead of the monitoring on the application:

The instrumentation cost, i.e., the cost of monitoring and communicating events to the algorithm for processing. In ERASER, every shared memory operation and synchronization (locking) operations has to be monitored. As the number and variety of events monitored by the algorithm increases, the frequency of interrupting the execution with callbacks to the algorithm increases and this becomes a bottleneck even

though the actions of the algorithm are simple and cheap. Our experimental results in Sec. 5 show that only instrumenting the program (without performing any computation at instrumentation points) can generate 1.6X to 7.1X overhead on the uninstrumented program.

The analysis cost, i.e., the cost of processing the events and updating the algorithm’s state accordingly. Runtime verification algorithms for concurrent programs usually maintain data structures shared among the threads participating in the algorithm. For example, the ERASER implementation maintains a lockset for each thread and for each shared variable. Other algorithms maintain, e.g., pointers to the last accessing thread or an additional virtual-clock vector to internal locks for synchronizing accesses by different threads. Fetching and manipulating these data structures at high frequencies creates a considerable overhead and may cause big divergences in the timing behavior of threads.

For ERASER, accessing the map LS (or $LS(x)$, if the map is distributed) may require to use a common lock to avoid two threads both accessing x to manipulate $LS(x)$ simultaneously. This creates large critical sections (code segments to be executed atomically) throughout the execution and is a significant source of runtime overhead. In summary, while ERASER is one of the simplest, cheapest algorithms for race detection, its implementations can cause considerable runtime and memory overhead and this makes the race checking hard to apply at the post-deployment.

In order to reduce the runtime overhead of the checking, we distribute the responsibilities for the algorithm to worker (checker) threads separate from the application threads. Checker threads run on separate cores, and do not slow down the application being monitored. In this paper, we investigate this idea by running the runtime analyses on GPUs. Our novel approach has the side benefit of simplifying the implementation of runtime verification algorithms, which are often forced to make use of tricky data structures and optimizations when run on the same threads as the applications. When run on separate cores, simpler but parallelized implementations of these algorithms provide the required performance.

The programmer has to spend a high amount of effort to make use of nontrivial and often error-prone mechanisms. For example, a naive ERASER implementation requires a large memory space to store locksets when there are a high number of shared variables. Moreover, it has to manage the locksets in the case of dynamic and frequent allocation/deallocation of threads and variables. In order to reduce the runtime and memory cost of the algorithm, the programmer has to develop a highly optimized implementation. For example, to reduce the number of memory allocations for locksets, locksets of deallocated variables are reused for newly allocated variables, requiring a memory pool of locksets. As another example, some algorithms adjust the granularity of shared variables between collection of variables (objects, arrays) to individual memory cells [10, 12]. These and similar extra management tasks are usually nontrivial and error-prone to implement and if not implemented carefully, may




Figure 1. Components of our runtime monitoring system.

create extra overhead on the core algorithm. This results in highly complicated implementations for very simple algorithms such as ERASER. Our separation of instrumentation and analysis allows one to focus on not tricky optimizations but a simple implementation of the core algorithm on highly efficient cores.


3. Our approach I: Overall system

Motivated by the challenges given in the previous section, our main goal is to design a runtime verification framework that will have the minimum negative impact on the program’s running time and concurrency. Our key design decision is to carry out the checking algorithm (i.e., data-race detection) on physically separate multi-processors, in our case the GPU cores. The application threads running on the CPU are only responsible for recording their events in a shared data structure and communicating events to the GPU for further processing. Fig. 1 illustrates this separation of responsibilities between the CPU and GPU threads. In this section, we present our techniques for observing an execution trace, i.e., recording events and communicating them to the GPU. The following section complements this description by giving GPU-based algorithms for data-race detection.

3.1 Observing the execution trace

Our technique is based on logging the execution as a linear sequence of events and running the analysis in a very efficient way. In order to enable efficient handling of the event log, we only process a fixed-size segment of this log, called *frame*, at a time. In our experiments we fixed this size as 1024-events and refer to it by the `FRAME_SIZE` constant. We treat each event frame a *unit of input* for the analysis implemented in the GPU. Each frame is checked independently from other frames and minimal information is kept between frames, e.g., racy variables to omit accesses to those variables. When a frame is completely checked, the events in it is discarded and it is reused to store later events.

While splitting the execution into independent frames may cause unsound results due to pairs of events from separate frames, our framework allows to adjust the precision to increase the chance of finding bugs. We have chosen to defer the soundness issue, since the goal of this study was to show the feasibility of highly parallel, at-speed runtime verification. Empirical evidence by other researchers indicates that this is a minor source of unsoundness: Many concurrency errors involve a small number of threads, and can be detected by focusing on a short portion of the execution [6].



Algorithm **RecordEvent**(*e*)

(Executed by application threads)

```

// Find the first frame to insert the event
1 frame := Head
2 index := AtomicGetAndIncrement(frame.size)
3 while (index ≥ FrameSize) {
4   if (frame = Tail) { goto line 1 } // restart
5   frame := frame.next
6   index := AtomicGetAndIncrement(frame.size)
7 }
// Insert event to the frame at index
8 frame[index] := e
// Shift Head, if the frame becomes full
9 if (index = FrameSize - 1) { Head := Head.next }

```

Algorithm **CheckFrames**()

(Executed by worker thread)

```

1 while (program is running) {
2   wait until Head ≠ Tail
3   frame := Tail
4   // Check frame at GPU
5   Copy frame to GPU device memory
6   Async-Call GPU kernel for race checking
7   // Shift Tail to reuse the frame
8   frame.size := 0
9   Tail := frame.next
10  wait until GPU kernel finishes
11  Copy result of the checking from GPU
12 }

```

Figure 2. The cyclic linked list of event frames and related algorithms.

Fig. 2 shows our main data structure for keeping event frames: a circular linked list. At any time this list contains a fixed number of frames, where each frame is a memory buffer to store `FRAME_SIZE` events. As explained below, the circular linked list allows us to reuse the frames in an efficient way throughout the execution. Fig. 2 shows pseudo code to record events (**RecordEvent**) and to process full event frames, i.e., communicating them to the GPU for the analysis (**CheckFrames**). While the former is performed by application threads, we dedicate a separate worker thread (running on the CPU) for the latter.

The event list is managed by non-blocking algorithms with few atomic instructions; no lock is required to record an event and process an event frame. At any point in the execution, we keep two pointers to frames in our event list: *Head* and *Tail*. The part of the list between *Head* and *Tail* (both inclusively) contains the frames (shown in white color in Fig. 2) that are being filled by application threads. The rest of the list between *Tail* and *Head* contains the frames that have become full and waiting to be checked (shown in gray color). At the initial state of the event list *Head* and *Tail* points to the same frame. While frames become full, *Head* is shifted, and as the full frames are checked, *Tail* is shifted (as shown in Fig. 2). In order to prevent data races on *Head* and *Tail*, we read from and write to these variables using atomic-reference operations.

Recording events (RecordEvent in Fig. 2). Each event frame has a field called *size*, which stores the number of events in the frame. When an application thread wants to

record an event, it traverses the list starting from *Head* (lines 1-7). At each step it performs an atomic operation that reads the current *size* of the frame being visited and increments its *size* by one (lines 2 and 6). If *size* of the last visited frame before incrementing was less than `FRAME_SIZE`, then the thread uses that value as *index* of the frame to record the event (line 8). Otherwise, the thread tries following frames in the list in a loop (lines 3-7). If a thread reaches *Tail* while traversing the list, it restarts as this indicates that the current frame is full and subject to checking by the worker thread. In our experiments we observed that, because the checking of the frames runs at speed very close to the program, such restarts were quite rare, i.e., there is always at least one empty slot to insert an event between *Head* and *Tail*.

After adding the event to the right frame, if the current application thread finds out that the current frame is *Head* and has just become full, it shifts the *Head* pointer to the next non-empty frame in the list (lines 9).

Processing full event frames (CheckFrames in Fig. 2). Our worker thread takes a full frame a time and sends it to the GPU for the checking (in Fig. 2 this is the rightmost frame in gray). For this, the worker thread continuously executes the loop until the program finishes (We omit the code that processes the non-empty frames after the program terminates). At each iteration of the loop, the thread first waits until *Head* and *Tail* do not point to the same frame, i.e., the list contains full frames (line 2). When the condition holds, the thread locates the frame pointed by *Tail* (line 3) and checks it at the GPU. See Sec. 4.1 for explanation of procedure (lines 4-5 and 8-9) for running the analysis on the GPU. As the analysis on the GPU runs asynchronously with the CPU, the worker thread spends the time to wait until the GPU computation terminates to mark the currently checked frame empty (line 6) and to shift *Tail* forward and make the frame available to be reused to record new events (line 7). Upon completion of the kernel call (line 8), the worker thread copies the result of the checking, i.e., racy accesses, from the GPU’s memory back to the CPU’s memory (line 9), in an algorithm-specific memory space. While our system reports all the errors at the end of the execution, it can be modified to report the errors as soon as it gets the response from the GPU.

4. Our approach II: Checking frames on the GPU

Having explained the CPU part of our runtime monitoring system, we present parallel algorithms for the data-race checking on the GPU cores. We first brief on GPU computing using the CUDA model and bring in some challenges in that model, which affected our system design. Then, we present our adaptation of ERASER and GOLDLOCKS algorithms to run on parallel GPU threads.

4.1 Background on GPU computing using CUDA

The CUDA model allows programmers to write code in an extension of the C language that will be run on GPU

in a highly parallel manner. The mapping of the code to physical processing units on the GPU is transparent to the programmer, and this enables one to write parallel code that can scale for devices with different parallel processing capabilities.

Each code portion to be run on GPU is written as a C function called *kernel* and can be called from C/C++ code executing on the CPU. Thus, in our framework, each analysis algorithm is written as a C function. The CPU and GPU threads operate on memory modules physically isolated from each other. As a result, we have to maintain a separate memory space on the GPU’s own device memory. For this, at the beginning of the execution, we pre-allocate a memory region, as large to fit a full event frame, on the GPU’s own device memory at the beginning of the execution. Additional space is also allocated to hold the intermediate results and outputs of the kernel’s computation. The pointers to these memory regions are given as arguments when to the kernel call. Our worker thread (running on the CPU) must follow the following steps to run an analysis on a full event frame:

1. The worker thread first copies the contents of the event frame to the pre-allocated region on the GPU device memory.
2. It calls the kernel function of an available checker algorithm. That kernel function is executed by the GPU cores in parallel and asynchronously with the CPU. When calling the kernel function, it passes as arguments the pointer to device memory region storing the current frame as well as additional values, such as the number of events in the frame. Each kernel is executed in a SIMD (single instruction, multiple data) style on multiple cores and threads.
3. The worker thread uses CUDA routines to synchronize with the kernel execution for further processing. Upon completion of the kernel call, the worker thread copies the result of the checking, i.e., pairs of racy accesses in the case of data-race detection, from the GPU’s device memory back to the CPU’s memory to be reported later.

Given the challenges in writing kernels, we wrote parallel kernels for the ERASER and GOLDLOCKS algorithms. The challenge in writing the kernels is to trade the challenges given above with the large number of cores available on the GPU. In our algorithms, each thread checks a unique variable access in the given event frame, creating the data structures, i.e. locksets, necessary for the check locally (in its stack) and discarding them after the check completes. The implementations of the kernels is also available at <http://kuda.codeplex.com>.

Due to the limited GPU memory space and the requirement to pre-allocate the memory used by the kernel, using bounded sized representations for data structures in the kernel is essential. For this, we represent locksets in both the ERASER and GOLDLOCKS kernels with bloom filters, which can represent a collection of addresses (of locks, variables, etc.) in constant-size bitsets. As the data required for a single check is of finite size and used locally and temporarily,

dynamically created objects or threads do not create extra memory space or management.

5. Experimental evaluation

We aim to evaluate two claims we referred to in Sec. 1: First, our separation of monitoring and analysis to CPU and GPU significantly reduces the overhead of the traditional approach in which both are performed on the same threads/cores. Second, our analysis code runs at a similar speed as the program and finishes soon after the program terminates. For this, we implemented our proposed system in a prototype tool called KUDA and applied KUDA on a collection of multithreaded benchmarks. KUDA consists of two parts:

1. A dynamic library containing the core functionality including the routines for recording events, managing event frames, and running the race detection kernels on the GPU. We use the CUDA 4.0 library [7] to write and call kernels for analyzing frames and to manage the GPU resources (e.g., transferring data to/from the GPU device memory). While our experiments are performed using the global memory, our system can use constant and texture memory. The fact that event frames are only read by the kernel enables us to make use of the constant and texture memory, which are cached for fast read-only access.
2. A Pin [5] tool to dynamically instrument x86 binaries in order to callback the routines in our dynamic library on certain events (shared memory read/write, thread creation/join, and inter-thread synchronization). Our Pin tool supports multithreaded programs written using the pthreads library (for thread creation and join, and synchronization primitives including mutex and readers/writer locks).

5.1 Experiments

Benchmarks. We applied our tool KUDA on a collection of multithreaded programs from PARSEC [1] and SPLASH-2 [11] benchmark suites. Due to space restrictions our experiment results table is not presented in this paper. But note that, in a typical execution, our benchmarks generate a few hundreds of millions of events and hundreds of thousands of frames, each of which is checked on the GPU.

Hardware. We performed our experiments on a HP xw9300 Workstation running Ubuntu Linux 10.10 32-bit kernel. Our machine has two (single-core) AMD Opteron processors with 2600 MHz clock frequency, 128 KB L1 cache, 1 MB L2 cache, and 8 GB memory (400 MHz). We used a GeForce GTX 465 GPU card with Fermi chipset. Our card provides 352 cores (11 processors with 32 cores each) with 1.21GHz clock rate, 1.23 GB of memory space with 1.4 GB/sec host-to-device memory bandwidth and 71.3 GB/sec in-device memory bandwidth.

Configuration parameters. For the experiments, we chose the following parameters that gave the best results in terms of runtime and memory overhead. We selected the event frame size (FRAME_SIZE) to be **1024** events. We initialize the cyclic list in Fig. 2 with **2048** frames. Thus, our system requires

only 2048 frames * 1024 events (each frame) * 8 bytes (each event) = **16 MBytes** of memory space to store the events for the CPU. We run **128** GPU threads over each event frame. In order to get the maximum benefit from the GPU device's concurrent computing functionality, we collect and send to the GPU **128** consecutive event frames at a time. In this way we aim to utilize the high parallelism on the GPU to analyze multiple frames simultaneously.

5.2 Results

The results indicate that the instrumentation even without executing any extra code incurs overhead that ranges between 1.6X and 7.1X.

In order to compare the runtime cost of our approach and the traditional approach in which the race detection runs on the same cores as the application, we implemented the ERASER, and two vector clock-based algorithms DJIT⁺ [8] and FASTTRACK [3] (available in our code base). For these algorithms, we used the same Pin instrumentation, but applied the algorithm's rules on the application threads immediately when a relevant event occurs. Our implementations are not perfectly optimized as in the original implementations, but still provide a rough estimate for the overhead of checking on the CPU.

We observed that the overhead of the DJIT⁺ and FASTTRACK implementations on the CPU are much higher than ERASER. Thus, the ERASER algorithm provides lower bounds for the runtime and slowdowns for these algorithms. Note that, the slowdown when running such a simple algorithm starts from 31.4X. Overall, running ERASER on the same cores as the application incurs a very high overhead and a few hundreds of times slowdown.

While our system contains GPU kernels for both the ERASER and GOLDILOCKS algorithms, we observed that the overhead when using ERASER gives only slightly lower overhead. GOLDILOCKS is a precise race-detection algorithm, and is the most expensive and complex one of the algorithms we investigated.

In fact, we observed that the analysis terminates shortly after the program terminates. The difference ranges between 1-3 milliseconds (on average 2.5 milliseconds). In addition, we observed that our system does not need to allocate new event frames; it simply reuses the initially allocated 2048 frames. This result, together with the small difference between the execution times of the program and analysis, indicates that the analysis runs at speed very close to the program, following the program behind only in milliseconds.

Our results clearly indicate that performing the checking on separate cores in a highly parallelized way dramatically reduces the overhead of the runtime verification. The ratio of the slowdown of the race checking on the CPU to that of the race checking on the GPU is between 3.3 (bodytrack) and 14.7 (fmm). Only for streamcluster the CPU-based implementation beats our system and gives less slowdown. Moreover, for raytrace benchmarks in both PARSEC and SPLASH-2, the execution took more than our specified upper time limit, 30 minutes; thus when we also consider these

benchmarks, the ratio of the slowdown of the CPU-based race checking to that of the GPU-based checking reaches at least 17 and 20 times, respectively.

The ratio of the overall slowdown to that of only managing the events goes only up to 1.4 (e.g., blackscholes). While the overhead of recording events is still high (e.g., for post-deployment purposes), this small difference between enabling and disabling on-GPU checking gives a promising evidence that the parallel processing events on the GPU gives negligible overhead.

References

- [1] Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh, and Kai Li. The parsec benchmark suite: Characterization and architectural implications. Technical Report TR-811-08, Princeton University, January 2008.
- [2] Tayfun Elmas, Shaz Qadeer, and Serdar Tasiran. Goldilocks: a race and transaction-aware java runtime. In *PLDI*, pages 245–255, 2007.
- [3] Cormac Flanagan and Stephen N. Freund. Fasttrack: efficient and precise dynamic race detection. In *PLDI*, pages 121–133, 2009.
- [4] Poletti Francesco, Poggiali Antonio, and Paul Marchal. Flexible hardware/software support for message passing on a distributed shared memory architecture. In *DATE*, pages 736–741, Washington, DC, USA, 2005.
- [5] Chi-Keung Luk, Robert Cohn, Robert Muth, Harish Patil, Artur Klauser, Geoff Lowney, Steven Wallace, Vijay Janapa Reddi, and Kim Hazelwood. Pin: building customized program analysis tools with dynamic instrumentation. In *PLDI*, pages 190–200, 2005.
- [6] Madanlal Musuvathi, Shaz Qadeer, Thomas Ball, Gerard Basler, Piramanayagam Arumuga Nainar, and Iulian Neamtii. Finding and reproducing heisenbugs in concurrent programs. In *OSDI*, pages 267–280, 2008.
- [7] NVIDIA Corporation. *NVIDIA CUDA Programming Guide v4.0*. NVIDIA Corporation, 2011.
- [8] Eli Pozniansky and Assaf Schuster. Multirace: Efficient on-the-fly data race detection in multithreaded c++ programs. *Concurr. Comput. : Pract. Exper.*, 19(3):327–340, 2007.
- [9] Stefan Savage, Michael Burrows, Greg Nelson, Patrick Sobalvarro, and Thomas Anderson. Eraser: a dynamic data race detector for multithreaded programs. *ACM Trans. Comput. Syst.*, 15(4):391–411, 1997.
- [10] Christoph von Praun and Thomas R. Gross. Object race detection. In *OOPSLA*, pages 70–82, New York, NY, USA, 2001.
- [11] Steven Cameron Woo, Moriyoshi Ohara, Evan Torrie, Jaswinder Pal Singh, and Anoop Gupta. The splash-2 programs: characterization and methodological considerations. In *ISCA*, pages 24–36, 1995.
- [12] Yuan Yu, Tom Rodeheffer, and Wei Chen. Racetrack: efficient detection of data race conditions via adaptive tracking. In *SOSP*, pages 221–234, 2005.