

# Navigating the Concurrency Landscape: A Survey of Race Condition Vulnerability Detectors

Aishwarya Upadhyay, Vijay Laxmi, Smita Naval

Department of Computer Science and Engineering

Malaviya Naional Institue of Technology Jaipur

Email: 2022rcp9026@mnit.ac.in, vlaxmi@mnit.ac.in, smita.cse@mnit.ac.in

**Abstract**—As technology continues to advance and we usher in the era of Industry 5.0, there has been a profound paradigm shift in operating systems, file systems, web, and network applications. The conventional utilization of multiprocessing and multicore systems has made concurrent programming increasingly pervasive. However, this transformation has brought about a new set of issues known as concurrency bugs, which, due to their wide prevalence in concurrent programs, have led to severe failures and potential security exploits. Over the past two decades, numerous researchers have dedicated their efforts to unveiling, detecting, mitigating, and preventing these bugs, with the last decade witnessing a surge in research within this domain. Among the spectrum of concurrency bugs, data races or race condition vulnerabilities stand out as the most prevalent, accounting for a staggering 80% of all concurrency bugs. This survey paper is focused on the realm of race condition bug detectors. We systematically categorize these detectors based on the diverse methodologies they employ. Additionally, we delve into the techniques and algorithms associated with race detection, tracing the evolution of this field over time. Furthermore, we shed light on the application of fuzzing techniques in the detection of race condition vulnerabilities. By reviewing these detectors and their static analyses, we draw conclusions and outline potential future research directions, including enhancing accuracy, performance, applicability, and comprehensiveness in race condition vulnerability detection.

**Keywords:** Race condition, Vulnerability detection, Concurrent programming, Synchronization, Shared resources, Software security, Multi-threading

## I. INTRODUCTION

The current era witnessed the emergence of Industry 5.0, also known as the Fifth Industrial Revolution, characterized by the collaboration between humans and advanced technology. This industrial paradigm shift involves the integration of humans and AI-powered machines, such as robots, to enhance workplace processes[1]. As a result, industries like textiles, agriculture, healthcare, and transportation are leveraging information and communication technology to drive efficiency[2]. Therefore, in the contemporary era of multi-core processors, harnessing concurrency has become a key driver for enhancing the performance of system software. As OS kernels and file systems evolve, an array of programming paradigms, including asynchronous work queues, read-copy-update (RCU), and optimistic locking, have been introduced to leverage multi-core computing capabilities [3][4][5][6]. However, these performance enhancements have not come without a cost. Concurrency bugs, particularly race conditions, have infiltrated the codebase, significantly compromising the

reliability and security of file systems. These bugs can lead to deadlock situations, kernel panics, data inconsistencies, and privilege escalations [7][8][9][10][11][12][13][14]. Despite ongoing efforts to harden OS kernels against various attacks, such as kASLR, kCFI, and UniSan, these defenses primarily address memory errors (e.g., stack and buffer overflows) and have limited effectiveness in mitigating attacks that exploit concurrency bugs. Race conditions, in particular, represent a challenging class of concurrency bugs, wherein two threads erroneously access a shared memory location without proper synchronization or ordering, as depicted in Figure 1. Due to the inherent non-determinism in thread interleaving, race conditions are notoriously elusive, demanding precise timing and rare interleaving conditions to manifest. For instance, even widely used file systems like *ext4*[14] and *btfs*[12], with 50,000 and 130,000 lines of code respectively, witnessed an alarming number of reported bugs in a single year[15]. According to the Linux kernel bug tracker, there are currently 1,234 open bugs in Linux kernel 5.15.0. Of these, 627 are classified as critical, 307 are classified as serious, and 300 are classified as moderate.[16]. In extreme cases, the

Count = 0 ;	
for(i=0; i<100; i++) { count++; }	for(i=0; i<100; i++) { count++; }
Thread 1	Thread 2

Fig. 1. A simple depiction of data race condition

failure to protect the system from concurrency vulnerability has led to a fatality. The Therac-25, a radiation therapy machine developed by Atomic Energy of Canada Limited (AECL), serves as a striking example of a significant threat posed by concurrent programming errors, commonly known as race conditions. Tragically, these errors occasionally led to patients receiving radiation doses hundreds of times higher than intended, resulting in fatalities or severe injuries[17].

## A. Motivation

Recent hardware architectures have seen a significant increase in the prevalence of multithreading support, with processors like the Intel Xeon Platinum 9282 having up to 56 cores per socket, each capable of running 2 threads, and the Intel Xeon Phi Processor 7290F, an accelerator with up to 72 cores, each supporting 4 threads. Systems with six to 260 cores per socket were included in the 500 list of supercomputers.[18] Because these systems use simultaneous multithreading (SMT), many threads can be executed concurrently on operating systems that support SMT and/or Symmetric Multi-Processing (SMP). All threads in a node share the memory (not the cache), which allows for shared memory multithreading. The memory can be accessed uniformly or non-uniformly (UMA/NUMA). Language extensions or application programming interfaces (APIs) were once used by the scientific community to execute parallel programs written in languages like FORTRAN and C/C++ across several nodes in a cluster or grid. As multi-core computers became more common, the emphasis moved to shared memory programming models. This meant combining ordinary languages like C/C++ or specialized parallel languages like array-FORTRAN or HPF with tools like the pthreads library/runtime system.[19] However, the adoption of multiprocessing and parallel architectures requires careful consideration of potential bugs and challenges associated with concurrent software execution. Addressing these issues is crucial to harness the full potential of the rapidly evolving technological landscape. The increasing prevalence of concurrent vulnerabilities and cyberattacks has underscored the need for secure IT infrastructure [20][21]. The graph in Figure 2 clearly shows how concurrency vulnerabilities have evolved from 1999 to 2021. Over this time frame, we see a consistent increase in the number of these vulnerabilities that have been reported[22]. This rise in vulnerability disclosures is closely tied to the growth of multi-core hardware and the increasing complexity of software that runs concurrently. It serves as a clear indicator of just how important it is to continue studying and researching concurrency vulnerabilities in the world of modern computing.

Understanding the nuances of concurrent applications and the security implications they pose is critical [23]. Researchers in the field have developed various tools and techniques to detect and mitigate vulnerabilities, particularly race conditions [24]. These vulnerabilities are prevalent in applications that utilize shared resources and often arise due to mistakes made by users and developers [25]. Efforts to address these challenges have led to advancements in detecting and preventing race condition vulnerabilities [26].

## B. Research Questions

This survey seeks to conduct a comprehensive and longitudinal analysis of race condition vulnerabilities in concurrent programming. Researchers in the field of multiprocessing and parallel computing have proposed various tools and techniques to detect, prevent, and mitigate these vulnerabilities. The review process involved a meticulous examination of available




Fig. 2. Increasing trend of concurrency vulnerability in last 20 years

literature, achieved by systematically searching through journal articles and conference papers that tackled various aspects related to concurrency bugs. This encompassed methodologies, tools, techniques, empirical assessments, and surveys. Major internet repositories were searched for pertinent papers. Initially, searches were conducted using the DBLP database, Google Scholar, and prestigious academic publishers including Wiley, ACM, IEEE, Springer, and Elsevier. Only those papers with certain keywords in their names, abstracts, or keywords—such as “concurrency bug,” “concurrent program,” “multithreaded program,” “data race,” “multiprocessing,” “race condition,” or “concurrency vulnerability”—were submitted. In addition, a period of time spanning from 1993 to 2023 was set aside for the inclusion of articles, and those deemed to be of inconsequential importance or with defective ideas were purposefully excluded. Additional articles of noteworthy value that were not found in the previously listed databases were hand-picked and added during the evaluation process. Figure 3 shows the spread of the selected research and their publishing category. This paper surveys the concurrency bugs and the most up-to-date and well-known race condition detectors. We categorize the existing detectors based on the types of techniques used to detect race condition vulnerability. Also, we shed special light on different types of fuzzers used in the detection of concurrency bugs. This paper aims to present an overview of the state-of-the-art in race condition vulnerability assessment and address the following research questions.

**RQ1: What causes concurrency vulnerabilities to arise and which type of concurrency bug makes up the majority of concurrency bugs**

Before going to the detection methodologies of concurrent bugs we have made an in-depth study of the various concurrency bugs and the reason for their occurrence. We have also included in our study the importance of lock-free synchronization and the concurrency issues arising from them.

**RQ2: Study the different methods and techniques for the exposure and detection of data race bugs.**

We have provided an overview of the current state-of-the-art race condition vulnerability assessment and analysis in hyper-

threaded and concurrent applications. We have reviewed the methods and techniques that have been developed to identify these vulnerabilities, and we have discussed their strengths and limitations. We have created our own taxonomy of the techniques used by the detectors and on that basis classified them into three categories namely static, dynamic, and hybrid. Specifically, we have reviewed the fuzzers that have been developed to identify these vulnerabilities and discussed their strengths and limitations.

### RQ3: What can be the future research aspect in the area of detection of race condition vulnerability?

Reviewing the most significant and current data race bug detectors which allowed to draw various conclusions about future research directions, including those related to accuracy, performance, application, and integrity.

We have also provided details of CWEs related to race condition vulnerabilities and described how they were discovered and remediated. Finally, we have outlined best practices that can be used to avoid or mitigate race condition vulnerabilities in hyperthreaded and concurrent applications.




Fig. 3. Division of included papers

### C. Review Organisation

This literature review is structured into three primary sections to address research questions. In Section 2, we offer background information, including concepts and notation related to concurrent programming. Section 3 outlines the issues in concurrent programming, discussing the characteristics of each bug and its severity. In Section 4, we emphasize understanding race condition vulnerability, with a particular focus on data race bugs. Section 5 explores application analysis for vulnerability detection, examining three analysis techniques: static, dynamic, and hybrid analysis. Section 6 discusses tools and repositories for application analysis and vulnerability detection, including studies on code vulnerability detection using machine learning and traditional approaches. We categorize experimental studies into three core sections: application/code analysis, code vulnerability detection, and supporting tools. Section 7 addresses potential research directions, and we provide a concluding summary of our findings.

## II. BACKGROUND

Prior to conducting a security analysis, it's crucial to understand how multiprocessing systems operate. This section offers an overview of the evolution of concurrent applications, their operation, the significance of concurrent and hyperthreaded applications, and key technical concepts used in this paper. It also explores the use of multithreading and synchronization methods in concurrent applications, aiming to improve overall reader comprehension.

### A. Notations and Concepts of Concurrent Programming

In the 1990s, C programming incorporated the POSIX thread library, known as pthreads, enabling multi-threaded software development. Today, modern computers and operating systems can run multiple programs concurrently, where concurrency means executing multiple instruction sequences simultaneously. This involves process threads that communicate through shared memory or message passing. However, concurrent resource sharing can lead to issues like resource scarcity and deadlocks. To improve efficiency, strategies like process coordination, memory allocation, and execution scheduling are employed. This section briefly introduces processes, threads, and computer architecture advancement, which are crucial for understanding vulnerabilities stemming from shared resources and concurrency.

1) *Processes and Threads*: Processes and threads are fundamental to concurrent computing, each serving distinct roles. Processes are encapsulated execution units that provide isolation by maintaining separate memory space and resources [27]. This isolation ensures robustness and fault tolerance in modern operating systems. In contrast, threads are lightweight execution units within a process, sharing the same memory space, which facilitates efficient communication and coordination among concurrent tasks [28]. This shared-memory model allows for parallelism exploitation, resulting in improved application performance and responsiveness. The interaction between processes and threads is crucial in designing concurrent software systems. Processes offer high isolation levels, making them suitable for running independent tasks or safeguarding critical system resources [29]. Threads, with their shared-memory model, enable efficient collaboration and parallelism within a process [30]. Striking a balance between process and thread usage is a critical aspect of concurrent software design, ensuring resource utilization and efficient execution in modern computing environments. Figure 4 depicts the multithreading in a single processor systems. For a mutlicore or multiprocessor mutlithreading the system will have multiple copies of code , data and file with each core acting as a independent processor.

2) *Synchronisation mechanism*: OS employs sophisticated strategies to enable concurrent execution of multiple programs (processes) and concurrent sub-tasks within these processes (threads) [31]. These mechanisms play a pivotal role in ensuring efficient system performance and enabling users to perform tasks simultaneously, contributing to a seamless and responsive computing experience.




Fig. 4. Multithreaded applications on a single core

- Locks** Mutex locks are commonly used to synchronize access to shared resources in multithreaded programs, allowing one thread at a time to access the shared resource. When a thread needs access, it must acquire the lock, use the resource, and then release the lock for other threads. Early work in operating systems identified two key synchronization types: mutual exclusion and condition synchronization, leading to the development of various synchronization primitives, such as busy-waiting semaphores, structured semaphores (e.g., critical regions, monitors, path expressions), message-passing primitives, and remote procedure calls. Busy-waiting semaphores, the earliest primitives, were found to be challenging and inefficient[29]. Structured semaphores were introduced to address these issues, offering a more organized approach. Message-passing primitives enhance synchronization with data, and remote procedure calls combine procedural interfaces with message passing. This historical and conceptual relationship among these primitives is illustrated in Figure 5.
- Lockless synchronization** Lockless synchronization, or lock-free synchronization, is a programming technique that manages shared resources without traditional locks or mutexes. It enhances the scalability and performance of multi-threaded or multi-core applications by reducing contention and minimizing blocking operations[32]. This approach allows multiple threads to access shared resources concurrently, preventing data races and deadlocks. It's particularly useful in scenarios where locks would introduce significant overhead due to thread contention[33]. Lockless synchronization offers increased application parallelism, optimizing multi-core processor utilization. However, designing and implementing lockless algorithms can be intricate, necessitating deep knowledge of low-level memory operations and platform-specific intricacies[34]. The concept of Compare and Swap (CAS) is often integral to lockless synchronization, providing a foundation for atomic operations and ensuring thread safety. CAS allows threads to compare a value in memory to an expected value and update it if the comparison succeeds, all in a single, uninterruptible operation [32].

This concept is crucial for implementing lock-free data structures and algorithms in lockless synchronization paradigms.

**Algorithm 1** Compare and Swap (CAS)

```

1: procedure CAS(address, expectedValue, newValue)
2:   currentValue ← address
3:   if currentValue = expectedValue then
4:     address ← newValue
5:   end if
6:   return currentValue
7: end procedure

```

The Compare and Swap (CAS) algorithm is a vital tool in concurrent programming, providing a means to atomically update a shared variable while ensuring thread safety. CAS takes three inputs: the memory location *address* to be updated, the expected value *expectedValue* that the address should have, and the new value *newValue* to set if the address matches the expected value. CAS reads the current value at the address, compares it to the expected value, and if they match, updates the address with the new value in an uninterruptible operation. CAS returns the current value regardless of whether the update occurs, making it useful for building lock-free data structures and synchronization primitives in multi-threaded applications, where only one thread at a time can successfully perform the update, ensuring coordinated and thread-safe modifications of shared variables.




Fig. 5. Synchronization techniques and language classes using locks.

**B. The Advancement of Computer Architecture**

The concept of "Moore's Law" coined by Gordon Moore in 1965 predicted that the number of transistors on a single chip would double every year or two, resulting in increased computing speed[35]. However, as depicted in Figure 6, the proliferation of transistors also led to overheating issues. This

compelled a shift in technology towards multi-processing, driving significant advancements in computer architecture. These advancements aimed to create faster, more efficient, and highly parallel computing systems, fueled by factors like the demand for enhanced computational power, energy efficiency, and the emergence of new technologies[36].




Fig. 6. Evolution of Intel Platform over the years [37]

One notable development is the transition from single-core to multi-core processors, enabling simultaneous execution of multiple tasks and enhancing system performance. Multi-core architectures have become commonplace in modern computing, facilitating parallel execution and improved throughput.[38]. Another crucial advancement is the introduction of vector processing and SIMD (Single Instruction, Multiple Data) architectures. These technologies enable the simultaneous execution of a single instruction on multiple data elements, accelerating data-intensive computations. SIMD architectures find applications in scientific simulations, multimedia processing, and machine learning. [39]. Moreover, specialized accelerators like GPUs and FPGAs have revolutionized computer architecture. GPUs, originally designed for graphics, have evolved into powerful parallel processors for various computations. FPGAs provide hardware reconfigurability, allowing customization for specific workloads. [23]. Advances in memory technologies, including DDR4 and HBM, have alleviated memory bandwidth bottlenecks and facilitated efficient data access for high-performance computing.[40]. In summary, computer architecture has evolved to focus on parallelism and concurrency, as the era of increasing clock speeds has come to an end. Harnessing this computing power necessitates deep understanding and efficient software development.

1) *Multiprocessing Architecture*: Parallelism for program acceleration is a core objective of multiprocessor systems. Research in this area explores parallelism identification in sequential programs, resource allocation among competing processes, synchronization of cooperative processes, and the verification of parallel programs. The term "multiprocessing" often encompasses parallel processing systems, including multiprocessor systems. There are two types of multiprocessing:

- **Symmetric Multiprocessing** Symmetric multiprocessing (SMP) is a computing architecture where multiple processors or cores are identical in capabilities and roles. All processors share the same memory and operate under a single operating system, enabling parallel task execution and efficient resource utilization. SMP systems provide load balancing, scalability, and fault tolerance, making them suitable for various computing environments, including servers, workstations, and high-performance computing clusters. Challenges in SMP systems include memory consistency and synchronization, necessitating synchronization mechanisms like locks and semaphores to ensure data integrity [41][42][43][44].
- **Asymmetric Multiprocessing** Asymmetric multiprocessing (AMP) is a computing architecture where multiple processors or cores have distinct capabilities and roles. This diversity allows for optimized resource allocation and enhanced system efficiency. Unlike symmetric multiprocessing (SMP), where processors have similar capabilities, AMP systems categorize processors into roles, such as application processors for general tasks and auxiliary processors for specific functions like graphics rendering or I/O operations. AMP is commonly used in mobile devices, where heterogeneous multi-core processors optimize power consumption and balance performance. It is also applied in high-performance computing (HPC) environments, leveraging processor capabilities for parallelism and specific workload acceleration. Effective task scheduling and load balancing are critical in AMP systems, given varying processor capabilities [45][46][47].

Figure7 shows the underlying difference between symmetric and asymmetric multiprocessing systems, where the symmetric multiprocessing architecture works on a shared memory principle whereas the asymmetric multiprocessing system works in the master-slave model.

2) *Hyperthreading*: Various multithreading strategies have been employed over the years to address memory latency issues. One approach is Simultaneous Multi-Threading (SMT), which exposes greater parallelism to the CPU by handling instructions from multiple streams, thus enhancing processor utilization. Notably, SMT, such as Intel's Hyper-Threading (HT), is a cost-effective design choice that does not require extensive additional hardware [48]. Hyper-Threading, patented by Intel, is a technology present in Intel® Xeon™ and Intel Pentium® 4 processors. It enables a single physical processor to appear as two logical processors to the operating system. By duplicating each processor's architectural state while sharing execution resources, it allows a single physical processor to execute instructions from multiple threads in parallel, potentially improving overall performance [49]. Figure8 gives the detailed working advantage of a hyperthreaded system over a single-core processor, where the throughput is increased due to hyperthreading technology. Research has explored the performance impact of Hyper-Threading. NASA conducted an analysis of HT's effect on processor resource utilization in scientific applications, concluding that while HT

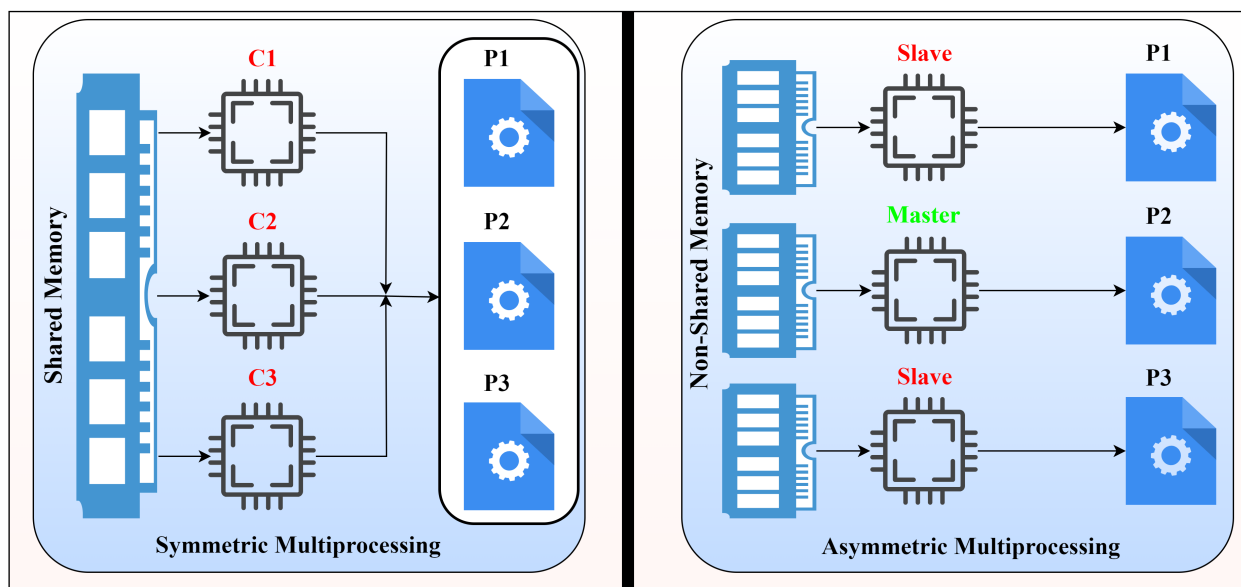


Fig. 7. Symmetric and Asymmetric Multiprocessing

enhances processor resource utilization efficiency, it does not consistently result in improved application performance [48]. Another study found that Hyper-Threading on specific processors provided minimal performance gains and even degraded the effectiveness of certain applications, particularly those reliant on vectorization and facing communication bottlenecks [50]. Conversely, a study focusing on database management systems (DBMS) reported performance improvements with Hyper-Threading, specifically in TPC-C-equivalent and TPC-H-equivalent queries [51]. Additionally, research has explored the impact of Hyper-Threading on resource utilization, concluding that while it enhances processor resource utilization efficiency, it may not universally improve application performance. Efficient code optimization and leveraging vector units are vital for enhancing code efficiency and maximizing processor resource utilization, ultimately leading to improved application performance [52].

### III. ISSUES IN CONCURRENCY

In the previous section, we examined the integration of multiprocessing systems and the advancements in hyperthreading technology, both of which have contributed to the widespread adoption of concurrent programming. This section will address the complexities and challenges associated with concurrent programming, along with an exploration of the distinct characteristics of concurrency-related software bugs. As understood in previous sections the concurrent programs use different synchronization primitives to coordinate between the threads and how the program will execute. The failure of these synchronization mechanisms can lead to various bugs such as data races, deadlock, atomicity violation, live lock, ABA problems & spurious wakeup. Broadly, concurrency-related bugs can be categorized into four major types: data races

(race conditions), order violations, violations of atomicity, and deadlocks, we include the challenges related to the ABA problem (arriving from lockless synchronization) and spurious wake-ups as well in this section.

#### A. Data Races

Data races occur when multiple threads simultaneously access the same part of computer memory (like a shared piece of information) without the proper coordination, and at least one of them is trying to change it. This can cause memory operations to happen unexpectedly, leading to different threads seeing inconsistent data and values. Surprisingly, sometimes, programmers intentionally allow data races for performance reasons, although it's important to note that data races and race conditions, despite often being used interchangeably, refer to different things. Data races relate to improper coordination in memory access, while race conditions are about timing-dependent issues during execution.

#### B. Atomicity Violation

Atomicity violation is a type of problem in multithreading, and it accounts for nearly 70 percent of all reported issues in this category. It happens when one thread's actions disrupt the order of operations in another thread, particularly when both threads are using the same shared resource. This disruption can cause the program to behave unpredictably.

#### C. Order Violation

Order violation is another issue after race conditions and atomicity violations. It arises due to the reordering of operations involving different parts of memory. In simpler terms, when two operations, let's call them A and B, are supposed to happen in a specific order (like A always before B), but

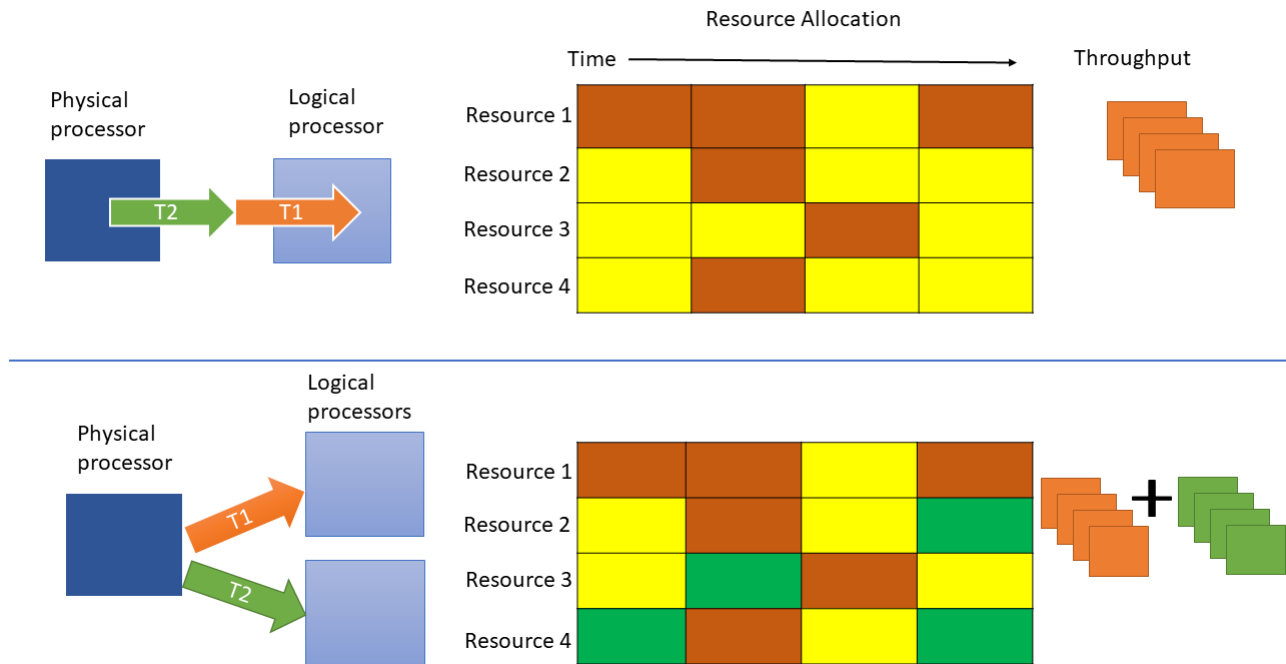


Fig. 8. Performance Improvement using Hyperthreading Technology

during execution, that order gets mixed up. Fixing an order violation often leads to a deadlock, a situation where threads get stuck waiting for each other. Interestingly, order violation is the least studied among these issues, often confused with atomicity violation. Research shows that both programmers and testers often struggle to identify the correct sequence of thread execution, making it a challenging problem to solve. It's also noteworthy that fixing errors in multithreaded applications usually takes longer compared to single-threaded ones.

#### D. Deadlock

Deadlock is an error in multithreaded programming when threads become stuck, unable to gain control over a mutex, which is vital for preventing issues like race conditions. Deadlocks typically arise from improper sequencing of operations or failure to release the mutex correctly. They can be classified into resource and communication deadlocks, with the former being more frequent. Four conditions must be met for a deadlock to occur: mutual exclusion (resources used by only one thread at a time), holding and waiting (a thread holds one resource and waits for another held by a different thread), no expropriation (only the holding thread can release a resource), and recurring wait (threads forming a cycle while waiting for each other's resources). All four conditions must be satisfied for a deadlock to manifest. Common scenarios for deadlock errors include mutexes mutually excluding each other, missing mutex release operations, attempting to create a mutex in a loop, and repeated calls to a function that generates a mutex.

#### E. ABA Problem

The ABA problem is a concurrency bug that occurs in lock-free and wait-free algorithms. It happens when one thread reads a shared variable, another thread alters the variable's value and then reverts it to the original value before the first thread writes to it. This can result in concurrency issues as the first thread remains unaware of the variable's change, leading to incorrect assumptions, unexpected behavior, or crashes. For instance, in a lock status scenario, if one thread acquires the lock (setting the variable to 1), and another thread waits for it to return to 0, an ABA problem can arise. If the first thread reacquires the lock, the variable changes back to 1, leaving the second thread blocked indefinitely. Detecting and resolving the ABA problem can be complex due to intricate interactions between threads, but mitigation techniques like using compare-and-swap (CAS) operations or version numbers can help prevent it.

#### F. Spurious Wakeup

A spurious wakeup is like an unexpected event in a multi-tasking program. It happens when a part of the program wakes up for no clear reason, usually when it's waiting for something to happen. This unexpected wakeup can lead to problems like strange behavior or even crashes in the program. For example, think of a program that uses a special signal to control when different parts can use a shared resource. If a part of the program wakes up without getting the right signal, it might try to use the shared resource at the wrong time, causing problems like data errors and other unexpected issues. Detecting and fixing spurious wakeups can be tricky because they often

TABLE I  
CONCURRENCY VULNERABILITY CAUSES AND MITIGATION

Concurrency Vulnerability	Causes	Mitigation
Race Condition	Arises when several threads or processes use the same data at the same time and at least one of them writes something.	<ul style="list-style-type: none"> <li>✓ Access to shared data can be coordinated by using synchronization techniques like locks, semaphores, or atomic actions.</li> <li>✓ Employ thread-safe programming practices and avoid sharing mutable state between threads without proper synchronization.</li> </ul>
Deadlock	Occurs when a circular dependency is created between two or more threads or processes that are blocked as they wait for resources that they jointly own.	<ul style="list-style-type: none"> <li>✓ Employ resource allocation and release strategies to ensure proper resource management and prevent circular dependencies.</li> <li>✓ Use techniques like resource ordering or resource pre-emption to break potential deadlocks.</li> </ul>
Starvation	Happens when a thread or process is consistently refused access to resources or execution because other threads or processes are given a higher priority.	<ul style="list-style-type: none"> <li>✓ Implement fair scheduling policies to prevent indefinite starvation.</li> <li>✓ Employ strategies to prevent priority inversion, such as priority ceiling protocols or priority inheritance.</li> </ul>
Data Inconsistency	Arises when multiple threads or processes perform concurrent read and write operations on shared data without proper synchronization, leading to inconsistent or incorrect results.	<ul style="list-style-type: none"> <li>✓ Use atomic operations, locks, or transactions as synchronization techniques to guarantee correct consistency and coordination.</li> <li>✓ Apply proper data access patterns and avoid unsynchronized concurrent modifications.</li> </ul>
ABA problem	Two or more threads modify a shared variable, leaving the variable with the same value as it had before the modifications.	<ul style="list-style-type: none"> <li>✓ In order to guarantee that only one thread is able to change the shared variable at once, use a compare-and-swap (CAS) procedure.</li> <li>✓ Use a version stamp to track the number of times the shared variable has been modified.</li> </ul>
Spurious wakeup	A thread is awakened from a wait state even though no other thread has signaled it.	<ul style="list-style-type: none"> <li>✓ Use a wait-free or lock-free algorithm to avoid spurious wakeups.</li> <li>✓ Use a retry loop to repeatedly attempt the wait operation until it succeeds.</li> </ul>

happen in complex situations with many parts of the program working together. However, there are ways to make spurious wakeups less likely, like using signals that can handle these unexpected events and using special tools to protect shared data.

A summarisation of the concurrency bug with their possible causes and mitigation is explained in Table I. From the definitions above and the Figure 9 as well we can infer that all the concurrency vulnerabilities are interrelated and there is a likelihood that a program containing a concurrency bug is very likely to trigger at most one other concurrency bug. In their extensive 2022 study, Lilibo et al. meticulously analyzed 839 concurrency vulnerabilities sourced from the National Vulnerability Database (NVD), revealing crucial insights into this prevalent issue. These vulnerabilities have surged in response to the growing scale and complexity of software systems, posing significant security threats. The most common among them are race conditions, particularly notorious for their potential to lead to severe remote attacks across diverse services. These concurrency vulnerabilities primarily fall into the medium-risk category, making up 67.2% of cases, although the distribution is relatively even among high and low-risk instances. Notably, both local and remote network access are used for exploitation, with many attackers favoring concurrency vulnerabilities due to their ease of exploitation. The study identified *race conditions* as the primary culprits across varying severity levels, often resulting in denial of service

issues. Furthermore, these vulnerabilities can lead to system file leaks, notably affecting system availability, with a higher risk of complete loss after exploitation[53].

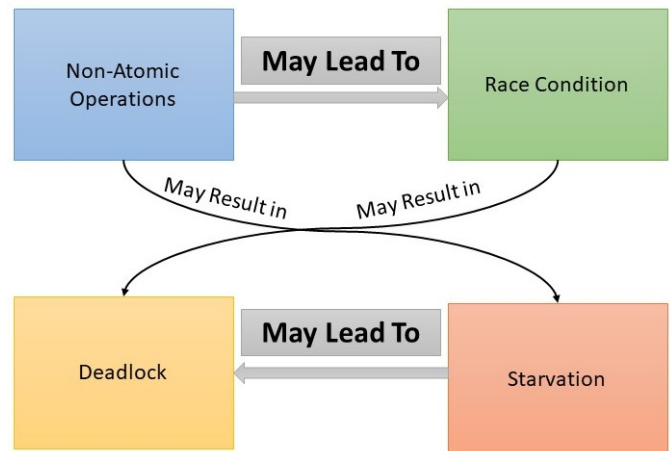


Fig. 9. Relationship between the concurrency bugs

#### IV. RACE-CONDITION VULNERABILITY

The analysis of Lilibo et al. study[53] reveals the critical role of race conditions in security concerns. therefore, in this and the preceding sections, we discuss everything keeping in view race condition vulnerability. In this section, our primary

focus is on race condition vulnerabilities, a type of software flaw that arises when multiple execution contexts, such as threads or processes, share a shared resource and can modify it concurrently, disregarding the need for mutual exclusion. An incorrect assumption often made is that a sequence of instructions or system calls will execute atomically, preventing any interference from other threads or processes. Regrettably, some developers tend to downplay the significance of addressing this issue, even when presented with evidence. In reality, most system calls involve the execution of thousands, or even millions, of instructions and often don't complete until another process or thread is given a chance to run.

The race condition vulnerability can occur in different parts of a computer system including :

- **Multithreaded Applications:** When multiple threads in a program access shared resources concurrently and perform read and write operations without proper synchronization, race conditions can arise.
- **Operating Systems:** When many processes or threads access shared kernel resources, like system data structures or device drivers, without the proper synchronization, race conditions can arise within the kernel or operating system components.
- **Network Communication:** Race conditions can also occur in network communication protocols when multiple processes or systems try to access and modify shared network resources concurrently.
- **Database Systems:** When several transactions try to access and modify the same data concurrently in database systems without the necessary separation mechanisms—such as locks or transaction management—race situations may occur.
- **File Systems:** Concurrent access to shared files or directories by multiple processes or threads can lead to race conditions if proper file locking or synchronization mechanisms are not employed.

All the above-listed sections where race condition can occur in a computer system can be mapped with the different CWEs related to race condition as depicted in Figure 10. Based on data association and operation block, Hong et al. have illustrated four types of race bugs in [54]: data race bug, block race bug, multi-data race bug, and multi-data block race bug. All these four classes can be related to the CWEs listed in TableII. Figure10 gives a relational structure of CWE-362 "Concurrent Execution using Shared Resource with Improper Synchronization" also known as race condition. An improper synchronization of shared resources is a cause of multiple CWEs which are said to be the child of the CWE-362.

To understand better, here is an example of a code snippet that intentionally contains a race condition:

```

1 def race_condition(x):
2     # Acquire a lock on the variable x.
3     lock.acquire()
4
5     # Modify the variable x.
```




Fig. 10. Relationship chart for CWE-362

```

6     x = x + 1
7
8     # Release the lock on the variable x.
9     lock.release()
10
11 # Create a lock object.
12 lock = threading.Lock()
13
14 # Create two threads that will both call the
15 # race_condition() function.
16 thread1 = threading.Thread(target=
17 # race_condition, args=(1,))
18 thread2 = threading.Thread(target=
19 # race_condition, args=(2,))
20
21 # Start the threads.
22 thread1.start()
23 thread2.start()
24
25 # Wait for the threads to finish.
26 thread1.join()
27 thread2.join()
28
29 # Print the value of x.
30 print(x)
```

The value of x will be printed to the console using this code. Nevertheless, if we execute this procedure more than once, we might observe various values for x. This is a result of the two threads simultaneously changing the variable x. When this happens, there may be a race condition and x does not end up as expected. We can use a lock to make sure that only one thread can access the variable x at a time in order to avoid race situations. Before changing the variable x in the code above, we first constructed a lock object and obtained the lock. We changed the variable x and then released the lock. This guarantees that the variable x cannot be modified simultaneously by the two threads.

#### A. Data Race Bugs

When two or more threads visit the same shared variable concurrently and at least one of those accesses is a write, it might lead to a specific kind of race condition known as data race. Data corruption and deadlocks are two examples

TABLE II  
CWEs RELATED TO RACE CONDITION VULNERABILITIES IN VARIOUS APPLICATIONS/COMPONENTS

Name	CWE -ID	Description
Insufficient Control Flow Management	691	The code lacks adequate control flow management during execution, resulting in potential scenarios where the control flow can be unexpectedly altered.
Concurrent Execution using Shared Resource with Improper Synchronization ('Race Condition')	362	The product incorporates a code sequence that can execute simultaneously with other code. This code sequence necessitates temporary, exclusive control over a shared resource. However, there is a specific time interval during which the shared resource may be altered by another concurrently running code sequence.
Signal Handler Race Condition	364	The product utilises a signal handler that creates a race condition.
Race Condition within a Thread	366	When two threads of execution concurrently access a resource, there is a potential risk of using the resource in an invalid state, which can result in an undefined state of execution.
Time-of-check Time-of-use (TOCTOU) Race Condition	367	The product performs a pre-use verification of a resource's state. However, it is possible for the resource's state to be altered after the verification and before its utilisation, thereby rendering the verification results invalid. The occurrence of this issue may result in the product executing actions that are not valid, particularly when the resource is in an unforeseen or abnormal state.
Context Switching Race Condition	368	The product executes a sequence of non-atomic operations to transition between contexts that span privilege or security boundaries. However, a race condition exists, which enables an attacker to manipulate or distort the product's behavior during the transition.
Race Condition During Access to Alternate Channel	421	The product provides an additional communication channel for authorized users, however, it is important to note that this channel may also be accessible to other actors.
Permission Race Condition During Resource Copy	689	The product does not establish the permissions or access control for a resource during the process of copying or cloning. As a result, the resource remains vulnerable to external entities until the copying process is finished.
Race Condition for Write-Once Attributes	1223	The hardware design includes a write-once register that can be programmed by an untrusted software component before the trusted software component. This sequence of events leads to the occurrence of a race condition issue.
Hardware Logic Contains Race Conditions	1298	The presence of a race condition within the hardware logic has the potential to compromise the security guarantees provided by the system.
Improper Synchronization	662	The product employs multiple threads or processes to enable temporary access to a shared resource that can only be accessed exclusively by one process at a time. However, it lacks proper synchronization of these actions, potentially resulting in simultaneous access to this resource by multiple threads or processes.

of unexpected or improper behavior that can result from data races. This survey's primary goal is to evaluate the approaches used to identify data races, both solely and inclusively. When we look at the elucidation of data race though Figure 11 where two threads, Thread 1 and Thread2, update the values of shared resources, namely sharedvar1 and sharedvar2. Later on, these tasks retrieve the values of these shared resources using the functions `do_sth_with_shared_resources1()` and `do_sth_with_shared_resources2()`, without any protective measures in place for these operations. In the absence of protective measures, it is important to note that the expected value of a shared resource, such as sharedvar1, may not necessarily match the value that was written in Thread 1 just before the function call. This discrepancy arises due to the concurrent execution of Thread 1 and Thread 2, which can lead to unexpected outcomes, such as reading a value of 21 or, in some cases, even a corrupt, random value. To elucidate this, consider the two possible sequences of events Sequence 1: Thread 1 writes sharedvar1 as 11 → Thread 1 invokes the `do_sth_with_shared_resources1()` function. Sequence 2: Thread 1 writes sharedvar1 as 11 → Thread 2 overwrites sharedvar1 with a value of 21 → Thread 1 subsequently calls the `do_sth_with_shared_resources1()` function. In the absence of protective mechanisms, any code implemented within the `do_sth_with_shared_resources1()` function cannot assume a specific sequence of events or a predeter-

mined value for sharedvar1. Therefore, if your code relies on a particular value of sharedvar1, it can lead to unintended consequences, rendering the data race a software bug. Data races manifest when shared resources are accessed concurrently by multiple tasks, resulting in unpredictable outcomes. Understanding data races can be challenging since the execution of instructions may not follow the order in which they are written, and the outcome can vary between different test runs, making data races elusive and difficult to reproduce and rectify. These findings emphasize the critical importance of implementing protective mechanisms to detect and mitigate data races and ensure the reliability and integrity of concurrent software systems. Therefore our proceeding sections will concentrate on the **various detection techniques and tools for data race bugs**.

## V. ANALYSIS FOR VULNERABILITY

The Program under test (PUT) should be analyzed for vulnerability as a first step to the detection of vulnerability. Based on the analysis techniques used by the researchers of all the papers that were considered for review, the papers were categorized into three basic detection techniques namely static, dynamic, and hybrid. Figure 12 depicts the division of the various techniques into these three categories. The three categories can be defined as :

- 1) Static Analysis: Potential vulnerabilities can be found before the application is even compiled by using static




Fig. 11. Simultaneous access to shared resources by two tasks without specific protection which can lead to data race

analysis, which examines an application's code without running it.

- 2) **Dynamic Analysis:** In dynamic analysis, the program is run and its behavior is observed for indicators of vulnerability, including race situations or deadlocks.
- 3) **Hybrid Analysis:** The benefits of both static and dynamic are combined in hybrid analysis. Static analysis is usually the first step, in which the code is examined for possible problems before running. Certain vulnerabilities can be found and problems can be detected early in the development phase thanks to static analysis. Next, the software is run and the behavior is observed under real-world situations during the dynamic analysis phase. Runtime-specific bugs, performance difficulties, and security vulnerabilities that might not be visible with static analysis alone can be found during this phase.

#### A. Algorithms Detection of Race condition

Detection algorithms are used to identify races in concurrent programming. To identify these races, numerous algorithms have been created, each with a different strategy and set of trade-offs. The following are some popular techniques (which include lock and lockless synchronization) for race detection:

1) *The lockset algorithm:* This approach, which is used in both static and dynamic analysis tools, detects a possible race problem when several threads access shared memory without any of them having a shared lock. Essentially, the technique determines that for every shared memory variable,  $v$ , there is a non-empty set of locks,  $C(v)$ , that each thread that wants to access the variable has to hold. At first, all of the locks that are available are placed in  $C(v)$ . Every thread has two sets of locks:  $writeLocks(t)$  indicates the write locks held, and  $locks(t)$  indicates all the locks held by that thread. The following is how the algorithm functions:

- 1) For every shared memory variable  $v$ , initialize  $C(v)$  with the complete list of locks.
- 2) Each thread keeps track of its held locks ( $locks(t)$ ) and its write locks ( $writeLocks(t)$ ) during execution.
- 3) Whenever a thread accesses a shared memory variable  $v$ :
  - It checks whether it holds all the locks in  $C(v)$  (i.e.,  $locks(t)$  contains all the locks in  $C(v)$ ).
  - It also verifies whether it holds the write lock for  $v$  (i.e.,  $writeLocks(t)$  contains  $v$ ).
- 4) If both conditions are met for a thread when accessing  $v$ , there is no potential race condition detected. However, if any of the conditions are not satisfied, it signifies a potential race condition, as shared memory  $v$  is being accessed without the necessary locks.

Static and dynamic analysis tools employ this approach to identify possible race situations when many threads access shared shared memory without the necessary locks. It maintains sets of locks for each shared variable, ensuring that threads adhere to the locking protocol when accessing shared memory to prevent data races. Regrettably, not all of the races identified by a lockset algorithm represent real races. It is possible to write code that is free from data races, either by employing clever programming techniques or by utilizing alternative synchronization methods such as signaling. This poses a significant challenge when it comes to distinguishing genuine bugs from false positives. Annotations and specific suppressions can be employed as strategies to mitigate this issue. An alternative algorithm for race detection is the "happens-before" algorithm.

2) *Happen Before Algorithm:* The Happens-Before algorithm is designed to determine the partial ordering of events in distributed systems, particularly in the context of identifying data races. Here's an overview of how it operates A single

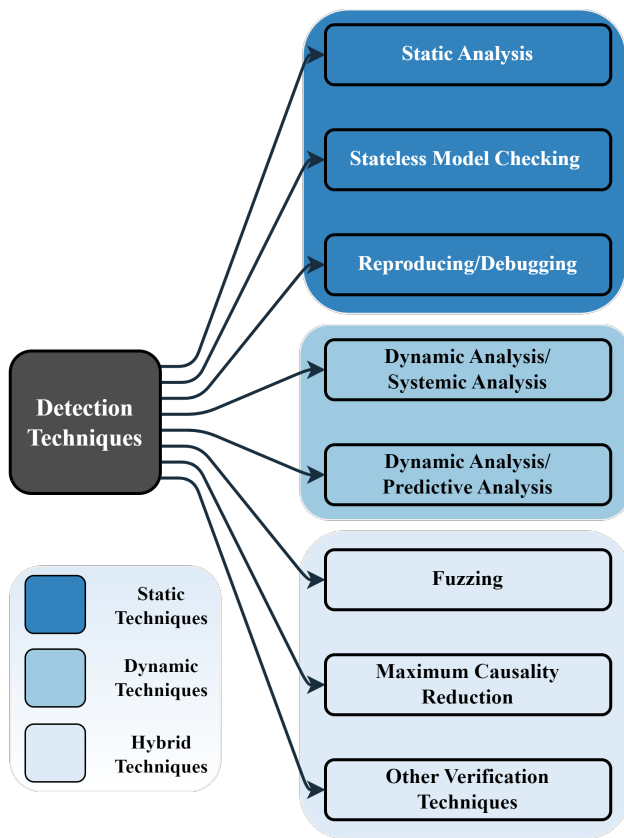


Fig. 12. Categorisation of different detection techniques used in the reviewed papers

thread's events are naturally arranged according to the order in which they occur. Events are arranged among threads based on synchronization fundamental features. In the event that two threads are vying for the same lock ('lock(a)'), for example, the unlocking of one thread is seen as having occurred before the locking of another thread. A possible race issue is indicated when many threads visit the same variable and their accesses are not deterministically ordered by the "happens-before" relationship.

However, there are some limitations to this algorithm: It can be computationally expensive to monitor these relationships. The order in which events are scheduled has a significant impact on this algorithm's performance. The partial order constructed is specific to a particular scheduling instance and may not detect the same bugs if the scheduling is different on another day. While some executions (Fig13(2)) may not report any races, others (Figure13(3)) may detect races. This inconsistency can lead to missed race conditions, with some only becoming apparent years after a product's release.

Comparably, although the Lockset algorithm is effective, it frequently produces a large number of false positives. Attempts have been made to integrate these algorithms in order to take advantage of their individual advantages and minimize their disadvantages. It's worth noting that race detection algorithms

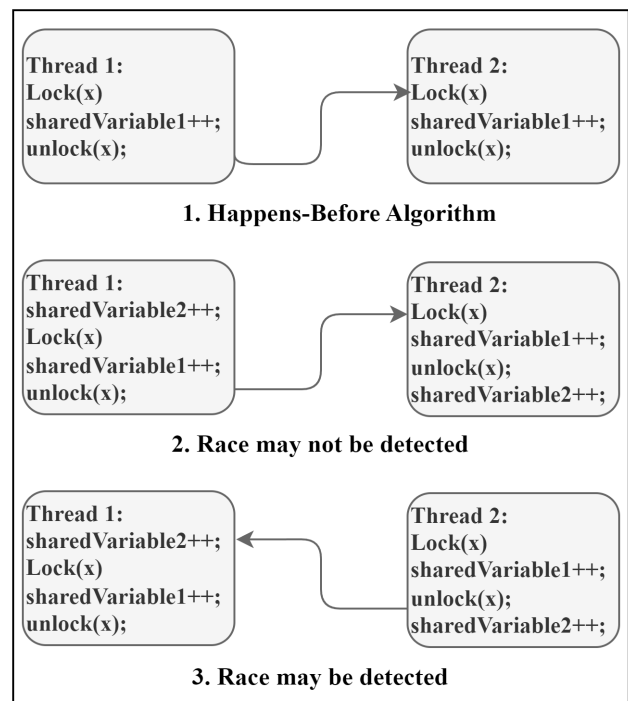


Fig. 13. Happen Before Algorithm

have undergone evolution over several years with the usage of a combination of other techniques such as *point analysis*, *CFG(Control flow graphs)*, *reachability analysis* and *lock analysis* etc.

### B. Fuzzing

Fuzzing is an approach that begins by identifying possible concurrency issues using bug detectors. It then takes control of thread scheduling and execution based on these bug reports, aiming to execute specific interleaving patterns in the program's operation to uncover concurrency bugs intentionally. To extract traits that indicate the existence of concurrency flaws in real-world scenarios, several fuzzing approaches rely on static analysis. When fuzzing is done intentionally, it exposes concurrency flaws more effectively than other techniques. Fuzzing falls under the hybrid analysis category since it employs both static and dynamic techniques to find concurrency issues.

### C. Machine learning

Machine learning or specific deep learning methods are also applied by researchers for the detection of concurrency bugs. There are several methods that deep learning can be applied to identify concurrency issues in concurrent programs. Using a dataset of acknowledged concurrency bugs to train a deep learning model is one such method. The model can then be used to identify new concurrency bugs in other programs by comparing them to the known bugs in the dataset. Another approach is to use deep learning to generate synthetic execution traces of concurrent programs and then train a model to identify concurrency bugs in the traces. This approach can

be more effective than traditional static analysis tools because it can detect concurrency bugs that are only exposed during runtime. Deep learning can also be used to develop dynamic analysis tools that can detect concurrency bugs in real-time. These tools can be used to monitor concurrent programs while they are running and to identify concurrency bugs as they occur. DeepRace[55] is one such method that leverages deep learning techniques to automatically detect data races, eliminating the need for manual creation of data race detectors. To find data races in code approaches, it uses a convolutional neural network (CNN) that has one layer and several window sizes. To extract the weights of the final convolutional layer, it also integrates the class activation map function with global average pooling. The lines of code linked to data races are then identified by backpropagating these weights across the input source code. As a result, it is claimed that the DeepRace model may identify data race issues in both files and lines of code.

## VI. TOOLS AND TECHNIQUES FOR DETECTION

To Detecting concurrent application vulnerabilities can be challenging, as they often involve complex interactions between multiple threads or processes. Figure 14 illustrates the comprehensive process of identifying concurrency bugs in a concurrent program (PUT) using an input seed. The various methods for detecting these concurrency bugs can be categorized into three distinct groups based on whether they involve program execution. These categories are static analysis, dynamic analysis, and hybrid analysis.

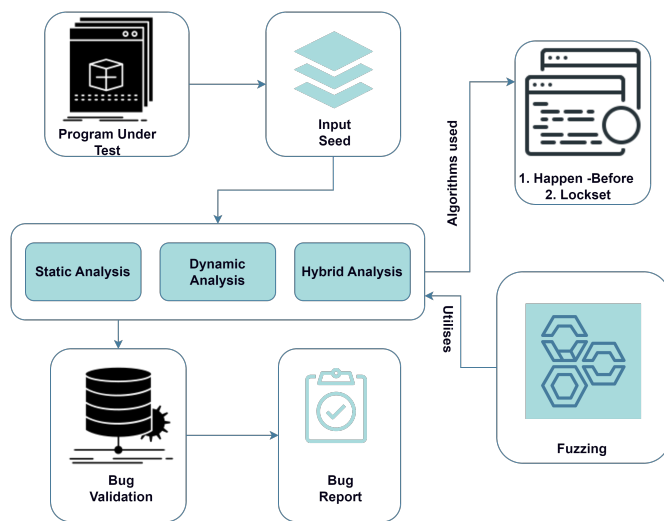


Fig. 14. Flow of Concurrency Detection

In handling potential race conditions, hangs, and all the other problems that arise when processing simultaneous transactions, a variety of concurrency testing tools are available on the market. Every tool we examine here is useful in a specific capacity.

### A. Static Analysis Tools

Because concurrency issues in software are intimately related to the way the program's scheduler manages threads, they are difficult to find using standard compile-time checks and testing techniques. To address this, various static analysis techniques have been developed. A static race detection analysis for concurrent Java programs was introduced by Flanagan and Freund et al.[56]. Their system focuses on lock-based synchronization rules for avoiding races. It checks to see if locks are appropriately held while using shared variables. But using this approach necessitates adding more type annotations, which might add to the complexity. To address this, Abadi and colleagues presented a static race detection study for Java programs running in high quantities concurrently[57]. There are no restrictions on this type-based analysis concerning test coverage issues. It works well with classes that have client-side or internal synchronization, as well as thread-local synchronization. Using 4000 lines of Java code, they tested this technique and were able to find race problems in other test programs as well as common Java libraries. Engler presented RacerX[58], a program that can identify deadlocks and race conditions without the need for annotations. Facebook created the static analysis tool RacerX. By examining the code, keeping track of shared variable accesses, and looking for order violations, it finds race situations in Java programs. RacerX uses flow-sensitive analysis and concurrent inter-procedural program analysis to efficiently find problems in complex systems. Large applications can be handled by it, but its false-negative rate is somewhat high. Naik et al. suggested a novel static method for Java program race detection in order to improve RacerX. Their four stage algorithm (reachable pairs, aliasing pairs, escaping pairs, and unlocked pairs) is context-sensitive. Using context sensitivity, it increases scalability even in large, extensively used programs and finds more flaws than earlier static race detection methods, such as RacerX. Naik and Aiken et al. proposed an algorithm to overcome RacerX's shortcomings in demonstrating racial freedom. Their main goal is to show that if two locks are different, then the memory regions they guard have to be different as well. In order to statically analyze multi-threaded programs and identify races, they use conditional must-not-aliasing. In order to counter code-centric techniques, Vaziri et al. presented a data-centric approach[59]. As a new criterion for data races, they identified 11 suspicious interleaving patterns and statically examined the code parts that required alteration to stop data races. Voung et al. introduced RELAY[60], a static and scalable method for identifying data races in the Linux kernel, in the context of C programs. In contrast to RacerX, RELAY detects data races through a context-sensitive, bottom-up analysis. In order to examine function behavior without regard to the calling context, it presents the idea of relative locksets. In 4.5 million lines of C code, RELAY detected 53 races; however, its false-positive rate was more than 70%. To identify concurrency in Java programs, several tools have been developed [61][62]. A commercial static analysis tool

with race situation detection features is called Coverity[61]. It tracks variable accesses and thread synchronization actions to perform inter-procedural analysis and find possible data races. A static analysis tool called ThreadSafe [62] is intended to find concurrency problems in Java programs, including race conditions. In order to find potential data races, deadlocks, and thread-safety violations, it examines the code. A commercial static analysis tool with race condition detection capabilities is Polyspace [63]. Formal verification and abstract interpretation approaches are employed to identify possible concurrency-related problems, such as data races, in C/C++ systems. On the other hand, CogniCrypt [64] is a static analysis tool designed to find data races in cryptographic code in Java systems. It finds possible cryptography-related race situations by combining symbolic execution and program slicing. These methods encompass various approaches to static and dynamic analysis for detecting concurrency bugs in both Java and C programs, each with its strengths and limitations. Table III depicts the comparative analysis of all the multiple techniques used by the static analysis tools for data race detection. Of all the tools majority of them use Happen Before and Point Analysis algorithms for the detection methodology. The use of CFG(Control flow graph) is also seen for the visualization of the program flow which makes it easier to statically analyze the pieces of code where the bug can occur. But, as the complexity increases static analysis cannot be scaled and can lead to many false positives therefore there arises a need for dynamic and hybrid analysis

### B. Dynamic Analysis Tools

Several dynamic analysis tools have been developed to detect data races and concurrency issues in various programming languages. These dynamic methods offer certain advantages, primarily the ability to examine observed feasible execution paths and accurately assess variable values and thread interleaving modes. This flexibility allows dynamic methods to effectively detect a range of issues, including plagiarism detection, making them versatile in problem identification. One notable tool in this category is Eraser[81], introduced by Savage et al. in 1997. Eraser employs a dynamic approach using the lockset algorithm to detect data races in multi-threaded production servers. It ensures that shared-memory accesses adhere to programming policies, safeguarding programs from data races. However, Eraser may have limitations in error detection due to incomplete test coverage. Smaragdakis et al. introduced the concept of “causally precedes” (CP), a generalization of the happens-before algorithm, enabling the observation of more races while maintaining accuracy and completeness. This dynamic technique enhances the detection of concurrency issues[82] Choi et al. presented an approach that combines elements of both lockset and happens-before algorithms to dynamically detect data races in object-oriented programs[83] Their experimental results demonstrated significant improvements in detection efficiency and reduced overhead compared to existing techniques. Yu et al. presented RaceTrack [84], a useful runtime race detection tool for

object-oriented applications, in opposition to Eraser and Choi’s method. RaceTrack, which was created especially for Microsoft’s Common Language Runtime, functions at the virtual machine level. It uses a hybrid detection approach to improve accuracy, reports questionable memory access patterns, and tracks program execution traces using instrumentation data. Xu et al. created the Serializability Violation Detector (SVD), a dynamic tool that offers root causes for debugging and employs backward error recovery (BER) to safeguard erroneous concurrent programs from errors[85]. With SVD, one can examine something after the fact without having to first annotate the program. In order to identify error messages, Flanagan et al. underlined the significance of thorough and accurate dynamic analysis for atomicity violations[56]. They did this by examining precise relationships between memory accesses in certain code areas. Ratanaworabhan et al.’s ToleRace[86] addressed asymmetric races and used a transaction-like technique to detect and tolerate races, greatly lowering the overhead associated with dynamic race detection. Jin et al. invented Cooperative Crug Isolation (CCI), a technique that may detect a wide range of concurrency flaws with minimal overhead and scalability, in contrast to many previous tools that target certain types of concurrency bugs[87]. Falcon[88] by Park et al. 2010 was introduced as a dynamic fault localization tool that directly corresponds to faults, detects various types of concurrency bugs, and effectively captures both order violations and atomicity violations using patternbased analysis. UNICORN [89] enhanced Falcon by monitoring pairs of memory access in C++ programs, allowing detection of both single-variable and multi-variable violations. A non-pattern-based tool called Recon[90] from Lucia et al. is used to handle single- and multivariable mistakes. It offers small segments of failure-inducing execution schedules to help with bug comprehension in addition to problem identification. Zhang et al. introduced Anticipating Invariant (AI) as a program invariant to detect and tolerate various concurrency bugs, exposing order violations and generating emergency patches when necessary[91] Some tools, like MagicFuzzer [92] by Cai et al. 2012 and Magiclock [93] from Cai and Chan 2014, were developed exclusively for deadlock detection, identifying hidden deadlock cycles and efficiently reducing the overhead of dynamic deadlock detection. ConMem [94] focused on bugs that can lead to program crashes caused by incorrect thread interleavings and memory problems. To efficiently identify such problems, it keeps an eye on how programs are being executed and examines memory accesses and synchronizations. ThreadSanitizer [95], developed by Google, is used for C/C++ and Go programs. It employs runtime instrumentation and happensbefore analysis to identify potential data races. Inspector, designed for Java programs, utilizes dynamic analysis techniques, including happens-before relations and synchronization tracking, to identify race conditions during program execution. Microsoft Research’s CHES [96] is a dynamic analysis tool that systematically explores thread schedules to detect potential race conditions, deadlocks, and other concurrency issues in concurrent programs. FastTrack [97], tailored for C/C++ concurrent

TABLE III  
COMPARATIVE ANALYSIS OF TECHNIQUES USED BY STATIC DETECTORS

Name of Tool	Intended Use	Methodology Employed					Description		
		Happen Before Analysis	Point to Analysis	Reachability Analysis	Model Checker	DAG/CFG Lock Set Analysis			
RacerX [58]	Linux, Free BSD	✓			✓	✓	✓	Flow sensitive, interprocedural analysis	
RELAY[60]	Linux Kernel	✓			✓		✓	Relative Lockset	
ERIGONE[65]	C language	✓			✓		✓	A partial reimplementa-tion of spin model checker	
RaceView[66]	C language						✓	DAG-based data race visualization for investigation and classification.	
DR-Frame[67]	Network Applications	✓					✓	Based on dataflow analysis technique. Uses happen before algo-rithm.	
SWORD[68]	Java Program written in eclipse IDE	✓	✓					Points-to and happens-before analysis fusion for efficient and precise static data race detection.	
Concurrent CFG[69]	C language	✓	✓					Points-to and happens-before analysis for static data race detection.	
Chord[70]	Java Program		✓	✓			✓	The proposed technique is based on a combination of points-to analy-sis, lock analysis, and reachability analysis	
LOCKSMITH[71]	C language		✓				✓	It uses a combination of interproce-dural analysis and constraint solv-ing to detect race conditions	
KISS[72]	Device Driver	✓			✓		✓	Sequential data race detection on a simplified concurrent program	
EPAJ[73]	Program Wirtten in JAVA						✓	✓	Partial discovered types for static data race detection and atomic-ity checking with runtime perfor-mance optimization.
COBE[74]	Program with asynchronous call	✓			✓		✓	✓	Happens-before inference for asyn-chronous data race detection in concurrent programs.
IteRace[75]	JAVA Program	✓	✓	✓					Points-to, happens-before, and reachability analysis for precise and efficient static data race detection.
codorrect[76]	C/C++, JAVA /ANDROID APP		✓				✓	✓	Origin-based static race detection for precise and efficient reasoning about shared memory and pointer aliases.
Open Race[77]	C/C++	✓	✓	✓					Points-to, happens-before, and reachability analysis for sound and complete static data race detection.
LLOV[78]	OpenMP	✓	✓	✓			✓		OpenMP control flow analysis-based data race detection for OpenMP programs using points-to, happens-before, and reachabil-ity analysis.
OMPRacer[79]	OpenMP	✓		✓			✓		OpenMP control flow analysis to construct a CFG of the OpenMP program and identify shared ob-jects, followed by happens-before analysis and reachability analysis for data race detection.
Goblint[80]	Device Driver (Written in C,C++)	✓	✓	✓					Goblint uses points-to, happens-before, and reachability analysis for data race detection in device drivers.

programs, combines happens-before analysis and lockset-based techniques for dynamic race detection. Lastly, ConTest [98] focuses on multi-threaded Java programs, employing dynamic analysis and formal methods to offer precise race detection and in-depth analysis of execution schedules. These tools provide valuable assistance in identifying and mitigating concurrency-related problems during runtime. In summary, these dynamic analysis tools offer diverse approaches to detect and address concurrency bugs, ranging from data races and atomicity violations to deadlocks and program crashes, providing valuable assistance to programmers in improving software reliability and robustness.

### C. Hybrid Analysis Tools

In addition to the two primary types of concurrent bug detectors mentioned above, there exist additional strategies, as shown in Table IV, that either blend static and dynamic methods together or combine one of the two ways with other approaches. Scholars have acknowledged the shortcomings of previous methods in identifying concurrent issues, especially those that depend on particular synchronization semantics, which may result in false reports of flaws. To address these limitations, Lu et al. introduced AVIO [99], a tool for detecting atomicity violations, which combines static and dynamic approaches. AVIO identifies AccessInterleaving invariants (AI invariants) that represent code sections expected to execute atomically. Concurrency problems are identified if memory access interleavings break these invariants during runtime. Similarly, Shi et al. introduced definition-use invariants (DefUse invariants) [100] extracted from training runs to dynamically detect violations and different types of program bugs. Zhang et al. proposed a consequence-oriented approach, focusing on a bug's lifecycle stages, leading to improved bug detection accuracy and coverage [101]. Kasikci et al. presented RaceMob [102], a data race detector that guarantees low runtime overhead and good accuracy by combining static and dynamic discovery. Deng et al. introduced Concurrent Function Pair (CFP) [103], an interleaving-coverage metric, to improve bug-detection efficiency in dynamic analysis. Additionally, innovative tools like MUVI [99], which detects inconsistent updates and multi-variable concurrency bugs, and Portend+ [104], which have been designed to improve concurrency bug identification and analysis. It not only detects data races but also evaluates possible repercussions to classify them depending on severity. Helgrind [105] is a versatile race detection tool within the Valgrind framework, employing a hybrid approach that combines both dynamic and static analysis methods for C/C++ programs. It effectively identifies race conditions by utilizing lockset-based analysis to uncover potential data races. Similarly, Eraser [81] is another hybrid tool designed for multithreaded C/C++ programs, employing a combination of static analysis and dynamic instrumentation to detect races. It relies on lockset-based analysis and thread preemption techniques to identify potential data race issues. RoadRunner [106] focuses on concurrent Java programs and employs a hybrid dynamic static analysis approach to uncover

race conditions and other concurrency-related problems. However, Saturn [107] is designed specifically for multithreaded C/C++ programs. It uses a combination of dynamic symbolic execution and static analysis to identify any data races and provide accurate execution traces. Finally, SWORD [68] operates in the realm of multithreaded Java programs, employing a hybrid approach that integrates static and dynamic analysis techniques. It conducts static analysis to identify potential races and then employ dynamic analysis to validate and refine the results, ensuring effective race condition detection. These tools collectively leverage both static and dynamic methods to overcome the limitations of each and provide more effective bug detection and classification.

1) *Fuzzers*: As discussed earlier one of the hybrid technique used for bug detection which is quite prominently used nowadays both for sequential and concurrent applications is fuzzing. Most fuzzing techniques for identifying concurrency bugs rely on static analysis to extract bug manifestation features from real-world concurrency issues. Fuzzing, as a method for deliberately exposing concurrency bugs, is considered more efficient compared to other approaches. Race-Fuzzer [113], developed by Sen et al. in 2008, is designed to efficiently reproduce data races while keeping overhead to a minimum. The approach initially identifies potential races using a detection technique. It then employs a randomized thread scheduler to manipulate thread execution in a way that triggers actual data races from the previously identified potential ones. Race-Fuzzer is effective at distinguishing real data races, which have the potential to cause program exceptions, from other potential races. However, its effectiveness in exposing bugs depends heavily on the underlying data race detection tools, potentially missing some bugs due to limited coverage. For detecting deadlocks, DeadlockFuzzer [114], introduced by Joshi et al. in 2009, identifies real deadlocks in multithreaded programs. It begins by locating suspicious deadlocks using dynamic analysis. DeadlockFuzzer then uses a randomized scheduler to regulate thread scheduling in order to replicate these suspicious deadlocks, hence raising the probability of exposing them. It has been noted, meanwhile, that DeadlockFuzzer might have trouble accurately confirming the presence of a true deadlock. ConLock [115], a mechanism that dynamically tests for deadlocks using constraint-based approaches, was invented by Cai et al. (2014) in order to address this problem. ConLock begins by examining a possible deadlock situation in order to derive a set of thread scheduling constraints. The rules guiding the relevant thread pairs' acquisition and release of locks during the possible deadlock are specified by these limitations. After then, the program is run within the limitations in an effort to produce a deadlock. ConLock labels a stalemate as false positive if it is not reproducible. Linux kernels using both lockless and lock-based synchronization methods for concurrency are more susceptible to bugs. Therefore, there have been multiple fuzzers developed for detecting concurrency bug in kernel programs. [116][117][118][119]. Table V gives a brief overview of the fuzzers developed for concurrency bug detection.

TABLE IV  
HYBRID DETECTION TECHNIQUES

Tool	Year	Author	PUT	Techniques used	Key Benefits
PACER[108]	2010	Bond et.al	Real world applications	Proportional Analysis	Efficiency, precision and versatility
RaceMob[102]	2013	Kasicki et.al	Real world applications	Crowd sourcing + statistical significance	Always-on in production, Accuracy, Low overhead
ConSeq[101]	2011	Zhang et.al	C/C++ applications	Sequential replay	Increased bug detection coverage, lower false positive
MultiRace[109]	2007	Pozniansky et.al	C++ applications	Djit and Lockset	Lower overhead, Scalability
ColFinder[110]	2015	Wu. et. al	Program, written in C language	static program slicing +Thread scheduling	Reduce time of bug manifestation and overhead
SDRacer[111]	2020	Wang et.al	Embedded programs written in C language	Static analysis+ Symbolic Execution	Automation, Accutacy and efficiency
HistLock+[112]	2018	Yang.et.al	C and C++ program	Lockset-based + History-based	Completeness, Precision, Efficiency,Versatility
SVD[85]	2005	Xu et.al	C programs	Lockset analysis + Happens-before analysis+ Atomic region analysis	Improved Reliability and lower false positive
Helgrind[105]	2005	Valgrind	C/C++ programs	Part of Valgrind tool suite+instrumenting program code	Ease of use, Accuracy

## VII. DISCUSSION OF DETECTION METHODS AND MITIGATION TECHNIQUES

The application analysis method employed in 32.3% and 45.4% of the analyzed studies, respectively, was static analysis and dynamic analysis. The method of hybrid analysis was employed by the remaining 22.3%. A diagram showing this is shown in Fig15(a).As demonstrated in Figure15(b), the adoption of detecting techniques has also changed, impacting conventional procedures. Static procedures were one of the main methods for detection in the early years, but as the 20th century approaches, dynamic and hybrid techniques become more prominent. Because static analysis concentrates more on code features than code-level analysis, it may be more commonly used than code-level analysis. In addition, the cost of static analysis is less than that of the other two techniques. To run the source code, dynamic needs extra resources like emulators or actual hardware. However, as technology advanced, programs grew more complicated and were harder to solve using static analysis. Therefore, in order to lower the complexity and subsequently the overheads, formal and dynamic methods of symbolic execution were employed.

Based on the review study, we also infer that the use of machine learning has been lower in the detection of data race bugs which can be a good direction to explore.

### A. Mitigation Techniques

Once a vulnerability has been detected, several strategies can be used to mitigate it. These include locking and synchronization mechanisms, such as semaphores and mutexes, which can prevent race conditions and deadlocks. Other mitigation strategies include input validation, error checking, and memory protection mechanisms such as ASLR and DEP. Software developers can also use programming languages and libraries designed to reduce the risk of concurrency-related vulnerabilities, such as Rust and Go.

## VIII. CONCLUSION & FUTURE DIRECTIONS

We have conducted a comprehensive assessment of the most current and critical tools for detecting data race bugs. We have organized these tools into categories based on the specific types of techniques they have used for bug identification. This categorization can help researchers understand which types of techniques have been the focus of previous detectors and provide insights for developing new algorithms that can complement existing tools. Furthermore, we have briefly discussed various techniques that have been used in recent times to enhance the effectiveness and efficiency of data race detection. These techniques include methods such as symbolic execution, fuzzing, guided coverage, proportional analysis, and static slicing. Researchers can leverage these techniques when designing new concurrency bug detectors. A statistical study of several data race bug detectors is also included in our analysis, taking into account the different types of approaches that each one uses. Notably, the discovery of data race bugs frequently makes use of dynamic analytic approaches. We have proposed future study avenues based on our findings, which can help researchers working on data race detection. Additionally, we can conclude that the majority of currently available detectors experience false negatives if we take into account the accuracy, precision, performance, and application of the devices. These results might direct future research efforts aimed at reducing the drawbacks of data race detection. Overall, we expect to see significant advancement in the area of concurrency bug detection. Our analysis has yielded valuable insights into these limitations and therefore our view of the present scenario and future scope can be seen as :

- Concurrency bugs are rarely discovered due to two primary factors: the immense and exponentially expanding interleaving possibilities as code size increases, and the fact that these bugs are typically hidden within uncommon and unique memory access interleaving patterns.

TABLE V  
CONCURRENCY BUG DETECTION FUZZERS

Fuzzer	Author	Year	PUT	Contribution	Limitation
ConFuzz[115]	Padhiyar et.al	2021	OCaml Programs	Builds on AFL to generate inputs that maximize coverage of concurrent event-driven programs' non-deterministic state space.	Computationally expensive and can generate false positives
CalFuzzer[120]	Koushik sen et.al	2007	C Programs	Uses the RAPOS algorithm to reduce sampling non-uniformity, leading to more efficient and effective testing.	needs to explicitly control the scheduler of the concurrent program to trigger concurrency bugs
Atomfuzzer[121]	Park et.al	2008	JAVA	Real-world flaws are effectively found by Java's randomized dynamic analysis technique, which alters the thread scheduler to discover atomicity violations with a high likelihood.	Computationally expensive, False Positive
Race Fuzzer[113]	Koushik sen et.al	2008	JAVA	Propose a new approach that is based on the idea of directing the random testing process towards areas of the program that are more likely to contain race conditions.	False Positive, Scalibility, Overhead
Deadlockfuzzer[114]	Pallavi Joshi et.al	2009	JAVA	Uses the LDG to identify potential deadlocks and then use a random thread scheduler to create a potential deadlock	False Negative, False Positive, Overhead
Assetfuzzer[122]	Zhifeng et.al	2010	C programs	The method prunes those interleavings that are free of violations and infers possible violations that do not show up in a particular execution.	False Negative, Overhead
Magicfuzzer[123]	Yan Cai et.al	2012	C programs	Founded on a unique strategy that prunes lock dependencies repeatedly until none of them has an incoming or outgoing edge.	Overhead and False Positive
RAZZER[116]	Dae R. Jeong et. al	2019	Linux Kernel	Combines the strengths of both stateless and stateful fuzzing to find kernel race bugs	False Positive, Scalibility, Overhead
DDRace[117]	Yuan et.al	2023	Linux Kernel	Uses directed fuzzing to detect inputs that cause concurrency UAF vulnerabilities after it has identified target driver interfaces and race pairs.	False Negative, Overhead
KRACE[118]	Meng Xu et.al	2020	Linux Kernel	Combines the strengths of both stateless and stateful fuzzing to find data race bugs in kernel file systems	False Negative, False Positive, Overhead
CONZZER[119]	Zu-Ming Jiang et.al	2022	User and Kernel	Technique that uses knowledge of the program's context and direction of execution to generate more effective inputs for detecting data-races.	Complex, Computationally Expensive, cannot detect all data races

- Recent research has made progress in detecting and mitigating race conditions in concurrent execution but challenges remain. Balancing performance and security is crucial, and researchers are exploring dynamic analysis techniques while seeking more automated methods for identifying race conditions. Future research should focus on improved techniques, formal methods for correctness, and secure programming languages.
- Static detectors produce few false negatives but numerous false positives, posing challenges for debugging. Dynamic detectors can miss some concurrency bugs due to unexecuted code and may be inefficient. Research directions include reducing false negatives in dynamic

detectors, addressing the accuracy-efficiency trade-off, and designing detectors to find real source code bugs and library-related issues

- Detectors often require manual annotations, limiting their use in complex programs. Researchers aim to automate these detectors. Some detectors generate excessive false positives, hindering bug identification and resolution. Improving accuracy enhances their applicability. A two-step detection process, akin to hybrid detectors, could be employed, with a focus on improving convenience. Most detectors are designed for C/C++ and Java, while languages like Objective-C and Python are also prevalent. There's a need for detectors compatible with a broader




Fig. 15. (a) Shows the distribution of techniques used in data race detection (b) Shows the statistical analysis of different detection techniques used over the years.

range of programming languages.

- Some dynamic detectors slow down program execution due to excessive instrumentation. Researchers can mitigate this by employing strategies like sampling or using virtual machines for monitoring. Dynamic detectors often execute programs redundantly, impacting efficiency. Strategies should be developed to eliminate redundant dynamic analyses. To make detectors suitable for production environments, researchers should work on hardware and algorithms to minimize runtime overhead. Deterministic execution techniques reduce non-determinism in large programs, improving software reliability and aiding data race bug detection. Predictive analysis can identify safe code, reducing the number of interleavings to be analyzed.
- Use of Formal methods such as model checking and Symbolic execution reduces the complexity of the programs and thus needs to be explored more in the area of guided symbolic execution.

- A promising research direction in the field of data race bugs, can be a detection solution that focuses on the collaboration between software and hardware. While software-only solutions are versatile and can be used on different platforms, they often slow down program performance. On the other hand, some systems try to make hardware changes to improve efficiency, but these modifications might not work well in a wide range of situations.

#### REFERENCES

- [1] Wolfgang Bauer et al. "Humans in Industry 5.0—A Paradigm Shift". In: *Procedia CIRP* 110 (2022), pp. 4–9.
- [2] Jiewu Leng et al. "Industry 5.0: Prospect and retrospect". In: *Journal of Manufacturing Systems* 65 (2022), pp. 279–295. ISSN: 0278-6125. DOI: <https://doi.org/10.1016/j.jmsy.2022.09.017>. URL: <https://doi.org/10.1016/j.jmsy.2022.09.017>.

- <https://www.sciencedirect.com/science/article/pii/S0278612522001662>.
- [3] Lanyue Lu et al. “A study of Linux file system evolution”. In: *Transactions on Storage* 10.1 (Jan. 2014), 3:1–3:32.
  - [4] C. Min et al. “Understanding Manycore Scalability of File Systems”. In: *Proceedings of the 2016 USENIX Annual Technical Conference (ATC)*. Denver, CO, June 2016.
  - [5] J. Huang, M. K. Qureshi, and K. Schwan. “An Evolutionary Study of Linux Memory Management for Fun and Profit”. In: *Proceedings of the 2016 USENIX Annual Technical Conference (ATC)*. Berkeley, CA, USA, June 2016, pp. 465–478. ISBN: 978-1-931971-30-0.
  - [6] A. Aghayev et al. “File Systems Unfit As Distributed Storage Backends: Lessons from 10 Years of Ceph Evolution”. In: *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP)*. Ontario, Canada, Oct. 2019.
  - [7] MITRE Corporation. *CVE-2009-1235*. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2009-1235> [Accessed: (22 august 2023)]. 2009.
  - [8] J. Corbet. *Unprivileged filesystem mounts, 2018 edition*. 2022. URL: <https://lwn.net/Articles/755593>.
  - [9] P. Fonseca, R. Rodrigues, and B. B. Brandenburg. “SKI: Exposing Kernel Concurrency Bugs Through Systematic Schedule Exploration”. In: *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. Broomfield, Colorado, Oct. 2014.
  - [10] W. Xu et al. “Fuzzing File Systems via Two-Dimensional Input Space Exploration”. In: *Proceedings of the 40th IEEE Symposium on Security and Privacy (Oakland)*. San Francisco, CA, May 2019.
  - [11] S. Kim et al. “Finding Semantic Bugs in File Systems with an Extensible Fuzzing Framework”. In: *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP)*. Ontario, Canada, Oct. 2019.
  - [12] bugzilla.kernel. *Kernel.org Bugzilla - Btrfs bug entries*. <https://bugzilla.kernel.org/buglist.cgi?component=btrfs> [Accessed: (22 august 2023)]. 2022.
  - [13] MITRE Corporation. *F2FS CVE entries*. <http://cve.mitre.org/cgi-bin/cvekey.cgi?keyword=f2fs> [Accessed: (22 august 2023)]. 2022.
  - [14] bugzilla.kernel. *Kernel.org Bugzilla - ext4 bug entries*. <https://bugzilla.kernel.org/buglist.cgi?component=ext4> [Accessed: (22 august 2023)]. 2022.
  - [15] Ohad Rodeh, Josef Bacik, and Chris Mason. “BTRFS: The Linux B-tree filesystem”. In: *ACM Transactions on Storage (TOS)* (2013).
  - [16] bugzilla.kernel. *Kernel.org Bugzilla*. <https://bugzilla.kernel.org> [Accessed: (22 august 2023)]. 2022.
  - [17] T. Gallagher. *Archived webpage: Concurrent Programming (CPE9001) - Monash University*. [https://web.archive.org/web/20071212183729/http://neptune.netcomp.monash.edu.au/cpe9001/assets/readings/www\\_uguelph\\_ca\\_tgallagh\\_tgallagh.html](https://web.archive.org/web/20071212183729/http://neptune.netcomp.monash.edu.au/cpe9001/assets/readings/www_uguelph_ca_tgallagh_tgallagh.html). Accessed on: July 31, 2023. 2007.
  - [18] TOP500.Org. *Top 500 Supercomputer Sites*. <http://www.top500.org/>. Retrieved May 8, 2019. 2019.
  - [19] Working Group for POSIX. *IEEE Standard for Information Technology–Portable Operating System Interface (POSIX(R)) Base Specifications, Issue 7. IEEE Std 1003.1-2017 (Rev. of IEEE Std 1003.1-2008)*. Tech. rep. IEEE, Jan. 2018, pp. 1–3951. DOI: 10.1109/IEEESTD.2018.8277153.
  - [20] Guanjun Lin et al. “Software vulnerability detection using deep neural networks: a survey”. In: *Proceedings of the IEEE* 108.10 (2020), pp. 1825–1848.
  - [21] Peng Zeng et al. “Software Vulnerability Analysis and Discovery Using Deep Learning Techniques: A Survey”. In: *IEEE Access* 8 (2020), pp. 197158–197172. DOI: 10.1109/ACCESS.2020.3034766.
  - [22] CVE details. *Vulnerabilities as Threats to IT Infrastructure*. <https://www.cvedetails.com/browse-by-date.php> [Accessed: (June 28, 2023)]. 2022.
  - [23] Zhen Li et al. “A survey of FPGA-based accelerator systems”. In: *ACM Computing Surveys (CSUR)* 50.3 (2017), p. 34.
  - [24] Peter M. Chen, Wen-Tsuen Hwang, and David J. Farber. “DERT: A Software Tool for Detecting and Eliminating Racy Variables”. In: *Proceedings of the 1990 International Conference on Parallel Processing*. Vol. 1. 1990, pp. 375–378.
  - [25] Manuel Castro. “An Empirical Study of Operating Systems Errors”. In: *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*. 2001, pp. 73–88.
  - [26] Sheng Yang et al. “Finding and Understanding Bugs in C Compilers”. In: *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. 2011, pp. 283–294.
  - [27] Andrew S. Tanenbaum and Herbert Bos. *Modern Operating Systems*. 4th. Pearson, 2014.
  - [28] David R. Butenhof. *Programming with POSIX Threads*. Addison-Wesley, 1997.
  - [29] Abraham Silberschatz, Peter B. Galvin, and Greg Gagne. *Operating System Concepts*. 10th. Wiley, 2018.
  - [30] J. P. Lewis and Mitchell Berg. *Threads Primer: A Guide to Multithreaded Programming*. Prentice Hall, 1998.
  - [31] Abraham Silberschatz, Peter B. Galvin, and Greg Gagne. *Operating System Concepts*. 9th. Wiley, 2014.
  - [32] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming*. Morgan Kaufmann, 2008.
  - [33] Mark Moir, Daniel Nussbaum, and Ori Shalev. “Using elimination to implement scalable and lock-free FIFO queues”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 29.3 (2007), pp. 1–34.

- [34] Timothy L. Harris. “A pragmatic implementation of non-blocking linked-lists”. In: *Distributed Computing* 14.3 (2001), pp. 195–212.
- [35] Intel. *Cramming More Components onto Integrated Circuits Moore’s Law*. Accessed on July 10, 2023. URL: <https://www.intel.com/content/www/us/en/history/virtual-vault/articles/moores-law.html>.
- [36] Ethan Mollick. “Establishing Moore’s Law”. In: *Annals of the History of Computing, IEEE* 28 (Aug. 2006), pp. 62–75. DOI: 10.1109/MAHC.2006.45.
- [37] Yevgen Voronenko. “Library Generation For Linear Transforms Dissertation”. In: (May 2012).
- [38] John L. Hennessy and David A. Patterson. *Computer Architecture: A Quantitative Approach*. Elsevier, 2011.
- [39] Wen-mei W. Hwu. “SIMD architecture and vectorization”. In: *ACM SIGARCH Computer Architecture News* 39.1 (2011), pp. 2–13.
- [40] Bruce Jacob, Spencer Ng, and David T. Wang. *Memory systems: cache, DRAM, disk*. Morgan Kaufmann, 2011.
- [41] R. Dutta and N. K. Jha. “Symmetric Multiprocessing”. In: *Handbook of Hardware/Software Codesign*. Boston, MA: Springer, 2017.
- [42] J. Bosch. “Symmetric Multiprocessing (SMP)”. In: *Encyclopedia of Parallel Computing*. Boston, MA: Springer, 2011.
- [43] A. Silberschatz, P. B. Galvin, and G. Gagne. *Symmetric Multiprocessing*. Wiley, 2018.
- [44] S. P. Dandamudi. *Symmetric Multiprocessing*. Boston, MA: Springer, 2005.
- [45] R. K. Gupta. “Heterogeneous Multiprocessor Architectures for Embedded Systems”. In: *Proceedings of the IEEE* 89.3 (Mar. 2001), pp. 420–437.
- [46] S. Kumar, D. Raghavendra, and L. Kallianpur. “A Survey of Asymmetric Multiprocessing in Modern SoCs”. In: *IEEE Design & Test* 33.1 (Feb. 2016), pp. 45–55.
- [47] D. Lebel, J. Levinsky, and D. Morin. “Asymmetric Multiprocessing in Mobile and Embedded Systems”. In: *2013 IEEE 19th International Symposium on Asynchronous Circuits and Systems*. June 2013, pp. 49–56.
- [48] Subhash Saini et al. “The impact of hyper-threading on processor resource utilization in production applications”. In: *2011 18th International Conference on High Performance Computing*. 2011, pp. 1–10. DOI: 10.1109/HiPC.2011.6152743.
- [49] Rizwan Ali et al. “Intel Hyper-Threading Technology to Achieve Computational Efficiency HIGH-PERFORMANCE COMPUTING”. In: *dell high performace computing* (Dec. 2003).
- [50] Subhash Saini and Robert T. Hood. “Performance Evaluation of Intel Broadwell Nodes Based Super-computer Using Computational Fluid Dynamics and Climate Applications”. In: *2017 IEEE 19th International Conference on High Performance Computing and Communications Workshops (HPCCWS)* (2017), pp. 58–65.
- [51] Wessam Hassanein, Layali Rashid, and Moustafa Hammad. “Analyzing the Effects of Hyperthreading on the Performance of Data Management Systems”. In: *International Journal of Parallel Programming* 36 (Apr. 2008), pp. 206–225. DOI: 10.1007/s10766-007-0066-x.
- [52] Okonta Emmanuel and D Ajani. “Performance Evaluation of Hyper Threading Technology Architecture Using Microsoft Operating System Platform”. In: *West African Journal of Industrial and Academic Research* (July 2019).
- [53] Lili Bo et al. “A Comprehensive Analysis of NVD Concurrency Vulnerabilities”. In: *2022 IEEE 22nd International Conference on Software Quality, Reliability and Security (QRS)*. 2022, pp. 9–18. DOI: 10.1109/QRS57517.2022.00012.
- [54] Shin Hong and Moonzoo Kim. “A survey of race bug detection techniques for multithreaded programmes”. In: *Software Testing, Verification and Reliability* 25 (Dec. 2014). DOI: 10.1002/stvr.1564.
- [55] Ali Tehrani Jamsaz et al. “DeepRace: Finding Data Race Bugs via Deep Learning”. In: *CoRR* abs/1907.07110 (2019). arXiv: 1907.07110. URL: <http://arxiv.org/abs/1907.07110>.
- [56] Cormac Flanagan, Stephen N Freund, and Joshua B Saxe. “ThreadSafe: Static analysis for thread-locality”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 32.4 (2010), p. 13.
- [57] Martin Abadi, Cormac Flanagan, and Stephen N. Freund. “Types for Safe Locking: Static Race Detection for Java”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 28 (2006), pp. 207–255.
- [58] Dawson Engler and Kathryn Ashcraft. “RacerX: Effective, Static Detection of Race Conditions and Deadlocks”. In: *ACM SIGOPS Operating Systems Review* 37.5 (2003), pp. 237–252.
- [59] Mandana Vaziri, Frank Tip, and Julian Dolby. “Associating Synchronization Constraints with Data in an Object-Oriented Language”. In: *ACM Sigplan Notices* 41 (2006), pp. 334–345.
- [60] Jason W. Vounq, Ranjit Jhala, and Sorin Lerner. “Relay: Static Race Detection on Millions of Lines of Code”. In: *Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering*. Dubrovnik, Croatia, 2007, pp. 205–214.
- [61] Satish Chandra et al. “Using static analysis to find bugs”. In: *IEEE Software* 25.5 (2008), pp. 22–29.
- [62] Cormac Flanagan, Stephen N Freund, and Joshua B Saxe. “ThreadSafe: Static analysis for thread-locality”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 32.4 (2010), p. 13.

- [63] Joshua Berdine et al. “Static detection of concurrency defects”. In: *ACM SIGPLAN Notices* 43.1 (2008), pp. 205–216.
- [64] Juraj Somorovsky et al. “CogniCrypt: supporting developers in using cryptography”. In: *IEEE Security & Privacy* 14.4 (2016), pp. 68–77.
- [65] M. Ben-Ari. “Tool Presentation: Teaching Concurrency and Model Checking”. In: *Proceedings of the International SPIN Workshop on Model Checking Software*. Cham, Switzerland: Springer, 2009, pp. 6–11. DOI: 10.1007/978-3-642-02652-2\\_5.
- [66] Nikolaos Koutsopoulos et al. “Advancing data race investigation and classification through visualization”. In: *2015 IEEE 3rd Working Conference on Software Visualization (VISOFT)* (2015), pp. 200–204. URL: <https://api.semanticscholar.org/CorpusID:14574499>.
- [67] D. Qi, N. Gu, and J. Su. “Detecting Data Race in Network Applications Using Static Analysis”. In: *Proceedings of the International Conference on Networking and Network Applications (NaNA)*. Oct. 2019, pp. 313–318. DOI: 10.1109/NaNA.2019.00061.
- [68] Chunhua Yang and Zhendong Su. “SWORD: a developer-oriented tool for detecting concurrency bugs”. In: *Proceedings of the 2008 ACM SIGPLAN conference on Programming language design and implementation*. ACM. 2008, pp. 45–56.
- [69] V. Kahlon et al. “Static Data Race Detection for Concurrent Programs with Asynchronous Calls”. In: *Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE)*. 2009, pp. 13–22. DOI: 10.1145/1595696.1595701.
- [70] M. Naik, A. Aiken, and J. Whaley. “Effective Static Race Detection for Java”. In: *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. New York, NY, USA: Association for Computing Machinery, 2006, pp. 308–319. DOI: 10.1145/1133981.1134018.
- [71] P. Pratikakis, J. S. Foster, and M. Hicks. “LOCKSMITH: Practical Static Race Detection for C”. In: *ACM Transactions on Programming Languages and Systems (TOPLAS)* 33.1 (Jan. 2011), pp. 1–55. DOI: 10.1145/1889997.1890000.
- [72] S. Qadeer and D. Wu. “KISS: Keep It Simple and Sequential”. In: *Proceedings of the 25th Annual ACM SIGPLAN Conference on Programming Language Design and Implementation*. New York: ACM, 2004, pp. 14–24.
- [73] R. Agarwal et al. “Optimized Run-time Race Detection and Atomicity Checking Using Partial Discovered Types”. In: *Proceedings of the 20th IEEE/ACM International Conference on Automated Software Engineering*. New York: ACM, 2005, pp. 233–242.
- [74] V. Kahlon et al. “Static Data Race Detection for Concurrent Programs with Asynchronous Calls”. In: *Proceedings of the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering*. New York: ACM, 2009, pp. 13–22.
- [75] C. Radoi and D. Dig. “Practical Static Race Detection for Java Parallel Loops”. In: *Proceedings of the International Symposium on Software Testing and Analysis*. New York: ACM, 2013, pp. 178–190.
- [76] Bozhen Liu\* et al. “When Threads Meet Events: Efficient and Precise Static Race Detection with Origins”. In: *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '21)*. Virtual, Canada: ACM, June 2021, 15 pages. DOI: 10.1145/3453483.3454073. URL: <https://doi.org/10.1145/3453483.3454073>.
- [77] Bradley Swain et al. “OpenRace: An Open Source Framework for Statically Detecting Data Races”. In: *Coderrect Inc.* (2021). Emails: brad@coderrect.com, bozhen.liu@tamu.edu, peiming@coderrect.com, yanze@coderrect.com, addison.crump@coderrect.com, rohan.khera@coderrect.com, jeff@coderrect.com”.
- [78] Utpal Bora et al. “LLOV: A Fast Static Data-Race Checker for OpenMP Programs”. In: *ACM Trans. Archit. Code Optim.* 17.4 (Dec. 2020). ISSN: 1544-3566. DOI: 10.1145/3418597. URL: <https://doi.org/10.1145/3418597>.
- [79] Bradley Swain et al. “OMPRacer: A Scalable and Precise Static Race Detector for OpenMP Programs”. In: *SC20, November 9-19, 2020, Is Everywhere We Are*. Affiliations and emails: Computer Science and Engineering, Texas A&M University, College Station, TX, brad, yanzeli, peiming@tamu.edu, jeff@cse.tamu.edu; Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA, ilaguna, georgakoudis1@llnl.gov; 0 Coderrect Inc, College Station, TX, brad@coderrect.com. IEEE. College Station, TX and Livermore, CA, 2020. ISBN: 978-1-7281-9998-6.
- [80] Vesal Vojdani et al. “Static Race Detection for Device Drivers: The Goblin Approach”. In: *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering. ASE '16*. Singapore, Singapore: Association for Computing Machinery, 2016, pp. 391–402. ISBN: 9781450338455. DOI: 10.1145/2970276.2970337. URL: <https://doi.org/10.1145/2970276.2970337>.
- [81] Stefan Savage et al. “Eraser: a dynamic data race detector for multithreaded programs”. In: *ACM SIGPLAN Notices*. Vol. 29. 6. ACM. 1994, pp. 27–37.
- [82] Yannis Smaragdakis et al. “Sound Predictive Race Detection in Polynomial Time”. In: *ACM Sigplan Notices* 47 (2012), pp. 387–400.
- [83] Jong-Deok Choi et al. “Efficient and Precise Data Race Detection for Multithreaded Object-Oriented

- Programs”. In: *ACM Sigplan Notices* 37 (2002), pp. 258–269.
- [84] Yuan Yu, Thomas Rodeheffer, and Wei Chen. “Race-track: Efficient Detection of Data Race Conditions via Adaptive Tracking”. In: *ACM SIGOPS Operating Systems Review*. Vol. 39. ACM. 2005, pp. 221–234.
- [85] Min Xu, Rastislav Bodík, and Mark D. Hill. “A Serializability Violation Detector for Shared-Memory Server Programs”. In: *ACM Sigplan Notices* 40 (2005), pp. 1–14.
- [86] Paruj Ratanaworabhan et al. “Detecting and Tolerating Asymmetric Races”. In: *ACM SIGPLAN Notices*. Vol. 44. ACM. 2009, pp. 173–184.
- [87] Guoliang Jin et al. “Instrumentation and Sampling Strategies for Cooperative Concurrency Bug Isolation”. In: *ACM Sigplan Notices* 45 (2010), pp. 241–255.
- [88] Soyeon Park, Richard W. Vuduc, and Mary Jean Harrold. “Falcon: Fault Localization in Concurrent Programs”. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering*. Vol. 1. 2010, pp. 245–254.
- [89] Soyeon Park, Richard Vuduc, and Mary Jean Harrold. “A Unified Approach for Localizing Non-Deadlock Concurrency Bugs”. In: *2012 IEEE 5th International Conference on Software Testing, Verification and Validation*. IEEE. 2012, pp. 51–60.
- [90] Brandon Lucia, Brian P. Wood, and Luis Ceze. “Isolating and Understanding Concurrency Errors Using Reconstructed Execution Fragments”. In: *ACM Sigplan Notices* 46 (2011), pp. 378–388.
- [91] Min Zhang et al. “A Lightweight System for Detecting and Tolerating Concurrency Bugs”. In: *IEEE Transactions on Software Engineering* 42.10 (2016), pp. 899–917.
- [92] Yifei Cai and Wing-Kwong Chan. “MagicFuzzer: Scalable Deadlock Detection for Large-Scale Applications”. In: *International Conference on Software Engineering*. 2012, pp. 606–616.
- [93] Yifei Cai and Wing-Kwong Chan. “MagicLock: Scalable Detection of Potential Deadlocks in Large-Scale Multithreaded Programs”. In: *IEEE Transactions on Software Engineering* 40 (2014), pp. 266–281.
- [94] Wei Zhang, Chang Sun, and Shan Lu. “ConMem: Detecting Severe Concurrency Bugs Through an Effect-Oriented Approach”. In: *ACM SIGARCH Computer Architecture News* 38 (2010), pp. 179–192.
- [95] Konstantin Serebryany et al. “ThreadSanitizer: data race detection in practice”. In: *Proceedings of the 2012 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. ACM. 2012, pp. 265–284.
- [96] Madanlal Musuvathi, Satish Narayanasamy, and Sebastian Burckhardt. “CHESS: a systematic testing tool for concurrent programs”. In: *Proceedings of the ACM SIGPLAN 2008 conference on Programming language design and implementation*. ACM. 2008, pp. 335–348.
- [97] Cormac Flanagan and Stephen N Freund. “FastTrack: efficient and precise dynamic race detection”. In: *ACM SIGPLAN Notices* 39.6 (2004), pp. 121–133.
- [98] Patrice Godefroid, Nils Klarlund, and Koushik Sen. “ConTest: Speculative testing for concurrent programs”. In: *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*. ACM. 2005, pp. 52–61.
- [99] Shan Lu et al. “Avio: Detecting Atomicity Violations via Access-Interleaving Invariants”. In: *IEEE Micro* 27 (2007), pp. 26–35.
- [100] Ying Shi et al. “Do I Use the Wrong Definition?: Defuse: Definition-Use Invariants for Detecting Concurrency and Sequential Bugs”. In: *ACM Sigplan Notices* 45 (2010), pp. 160–174.
- [101] Wenguang Zhang et al. “Conseq: Detecting Concurrency Bugs Through Sequential Errors”. In: *ACM Sigplan Notices* 39 (2011), pp. 251–264.
- [102] Baris Kasikci, Cristian Zamfir, and George Candea. “Racemob: Crowdsourced Data Race Detection”. In: *Twenty-Fourth ACM Symposium on Operating Systems Principles*. 2013, pp. 406–422.
- [103] Dongdong Deng, Wenguang Zhang, and Shan Lu. “Efficient Concurrency-Bug Detection Across Inputs”. In: *ACM Sigplan Notices* 48 (2013), pp. 785–802.
- [104] Baris Kasikci, Cristian Zamfir, and George Candea. “Automated Classification of Data Races Under Both Strong and Weak Memory Models”. In: *ACM Transactions on Programming Languages & Systems* 37 (2015), pp. 1–44.
- [105] Julian Seward, Nicholas Nethercote, and Josef Weidendorfer. “Valgrind: A framework for heavyweight dynamic binary instrumentation”. In: *Proceedings of the 2005 ACM SIGPLAN conference on Programming language design and implementation*. ACM. 2005, pp. 89–100.
- [106] Zheng Cui, Xuejun Yang, and Yiying Zhou. “Road-Runner: Towards Automatic Checkpointing of Whole-System Computations”. In: *Proceedings of the 2013 USENIX Annual Technical Conference (USENIX ATC)*. USENIX Association. 2013, pp. 179–190.
- [107] Wenguang Chen, Ying Lin, and Jingling Xue. “Saturn: efficient software error detection using relaxed dependencies”. In: *IEEE Transactions on Software Engineering* 32.11 (2006), pp. 873–890.
- [108] M. D. Bond, K. E. Coons, and K. S. McKinley. “PACER: Proportional Detection of Data Races”. In: *Proceedings of the 31st Annual ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM. New York, 2010, pp. 255–268.
- [109] E. Pozniansky and A. Schuster. “MultiRace: Efficient On-the-Fly Data Race Detection in Multithreaded C++ Programs”. In: *Concurrent Computation: Practice and Experience* 19 (2007), pp. 327–340.

- [110] Z. D. Wu et al. “Collaborative Technique for Concurrency Bug Detection”. In: *International Journal of Parallel Programming* 43 (2015), pp. 260–285.
- [111] Y. Wang et al. “Automatic Detection, Validation, and Repair of Race Conditions in Interrupt-Driven Embedded Software”. In: *IEEE Transactions on Software Engineering* (Apr. 2020). Early Access. DOI: 10.1109/TSE.2020.2989171.
- [112] J. Yang, B. Jiang, and W. K. Chan. “HistLockC: Precise memory access maintenance without lockset comparison for complete hybrid data race detection”. In: *IEEE Transactions on Reliability* 67.3 (Sept. 2018), pp. 786–801. DOI: 10.1109/TR.2018.2832226.
- [113] Koushik Sen. “Race Directed Random Testing of Concurrent Programs”. In: *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI ’08. Tucson, AZ, USA: Association for Computing Machinery, 2008, pp. 11–21. ISBN: 9781595938602. DOI: 10.1145/1375581.1375584. URL: <https://doi.org/10.1145/1375581.1375584>.
- [114] Pallavi Joshi et al. “A Randomized Dynamic Program Analysis Technique for Detecting Real Deadlocks”. In: *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation*. PLDI ’09. Dublin, Ireland: Association for Computing Machinery, 2009, pp. 110–120. ISBN: 9781605583921. DOI: 10.1145/1542476.1542489. URL: <https://doi.org/10.1145/1542476.1542489>.
- [115] Sumit Padhiyar and K. C. Sivaramkrishnan. “ConFuzz: Coverage-Guided Property Fuzzing for Event-Driven Programs”. In: *Practical Aspects of Declarative Languages*. Ed. by José F. Morales and Dominic Orchard. Cham: Springer International Publishing, 2021, pp. 127–144. ISBN: 978-3-030-67438-0.
- [116] Dae R. Jeong et al. “Razzer: Finding Kernel Race Bugs through Fuzzing”. In: *2019 IEEE Symposium on Security and Privacy (SP)*. 2019, pp. 754–768. DOI: 10.1109/SP.2019.00017.
- [117] Ming Yuan et al. “DDRace: Finding Concurrency UAF Vulnerabilities in Linux Drivers with Directed Fuzzing”. In: *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*. Ed. by Joseph A. Calandrino and Carmela Troncoso. USENIX Association, 2023. URL: <https://www.usenix.org/conference/usenixsecurity23/presentation/yuan-ming>.
- [118] Meng Xu et al. “Krace: Data Race Fuzzing for Kernel File Systems”. In: *2020 IEEE Symposium on Security and Privacy (SP)*. 2020, pp. 1643–1660. DOI: 10.1109/SP40000.2020.00078.
- [119] Zu-Ming Jiang et al. “Context-Sensitive and Directional Concurrency Fuzzing for Data-Race Detection”. In: *29th Annual Network and Distributed System Security Symposium, NDSS 2022, San Diego, California, USA, April 24-28, 2022*. The Internet Society, 2022. URL: <https://www.ndss-symposium.org/ndss-paper/auto-draft-198/>.
- [120] Koushik Sen. “Effective Random Testing of Concurrent Programs”. In: *Proceedings of the 22nd IEEE/ACM International Conference on Automated Software Engineering*. ASE ’07. Atlanta, Georgia, USA: Association for Computing Machinery, 2007, pp. 323–332. ISBN: 9781595938824. DOI: 10.1145/1321631.1321679. URL: <https://doi.org/10.1145/1321631.1321679>.
- [121] Chang-Seo Park and Koushik Sen. “Randomized Active Atomicity Violation Detection in Concurrent Programs”. In: *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. SIGSOFT ’08/FSE-16. Atlanta, Georgia: Association for Computing Machinery, 2008, pp. 135–145. ISBN: 9781595939951. DOI: 10.1145/1453101.1453121. URL: <https://doi.org/10.1145/1453101.1453121>.
- [122] Zhifeng Lai, S. C. Cheung, and W. K. Chan. “Detecting Atomic-Set Serializability Violations in Multi-threaded Programs through Active Randomized Testing”. In: *Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1*. ICSE ’10. Cape Town, South Africa: Association for Computing Machinery, 2010, pp. 235–244. ISBN: 9781605587196. DOI: 10.1145/1806799.1806836. URL: <https://doi.org/10.1145/1806799.1806836>.
- [123] Yan Cai and W. K. Chan. “MagicFuzzer: Scalable Deadlock Detection for Large-Scale Applications”. In: *Proceedings of the 34th International Conference on Software Engineering*. ICSE ’12. Zurich, Switzerland: IEEE Press, 2012, pp. 606–616. ISBN: 9781467310673.