Bridging the Gap between Hardware Fuzzing and Industrial Verification

Ruiyang Ma*, Tianhao Wei*, Jiaxi Zhang, Chun Yang, Jiangfang Yi, Guojie Luo†
School of Computer Science, Peking University, Beijing, China
National Key Laboratory for Multimedia Information Processing, Peking University, Beijing, China
{ruiyang, weitianhao}@stu.pku.edu.cn, {zhangjiaxi, chunyang, yijiangfang, gluo}@pku.edu.cn

ABSTRACT

As hardware design complexity increases, hardware fuzzing emerges as a promising tool for automating the verification process. However, a significant gap still exists before it can be applied in industry. This paper aims to summarize the current progress of hardware fuzzing from an industry-use perspective and propose solutions to bridge the gap between hardware fuzzing and industrial verification. First, we review recent hardware fuzzing methods and analyze their compatibilities with industrial verification. We establish criteria to assess whether a hardware fuzzing approach is compatible. Second, we examine whether current verification tools can efficiently support hardware fuzzing. We identify the bottlenecks in hardware fuzzing performance caused by insufficient support from the industrial environment. To overcome the bottlenecks, we propose a prototype, HwFuzzEnv, providing the necessary support for hardware fuzzing. With this prototype, the previous hardware fuzzing method can achieve a several hundred times speedup in industrial settings. Our work could serve as a reference for EDA companies, encouraging them to enhance their tools to support hardware fuzzing efficiently in industrial verification.

CCS CONCEPTS

• **Hardware → Functional verification**; **Simulation and emulation**; **Software tools for EDA**.

ACM Reference Format:
Ruiyang Ma, Tianhao Wei, Jiaxi Zhang, Chun Yang, Jiangfang Yi, Guojie Luo. 2025. Bridging the Gap between Hardware Fuzzing and Industrial Verification. In *Great Lakes Symposium on VLSI 2025 (GLSVLSI '25), June 30-July 2, 2025, New Orleans, LA, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3716368.3735190

1 INTRODUCTION

The growing complexity of hardware designs becomes a significant challenge for verification engineers. It is increasingly difficult for them to identify potential vulnerabilities during the hardware development process [2, 14]. Extensive effort is dedicated to creating test cases of the design under test (DUT), which is crucial for achieving DUT coverage closure and discovering bugs in corner cases.

*Both authors contributed equally to this research.
†Corresponding author.

GLSVLSI '25, June 30-July 2, 2025, New Orleans, LA, USA

ACM ISBN 979-8-4007-1496-2/2025/06.
https://doi.org/10.1145/3716368.3735190

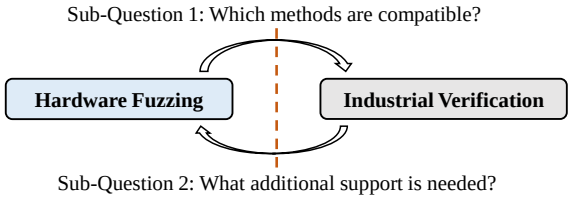

Figure 1: Our sub-questions about the current gap between hardware fuzzing and industrial verification.

Hardware fuzzing has recently emerged as a promising technique for reducing the substantial human effort required in hardware verification [2, 3, 6, 9, 10, 13–15, 17, 28]. This innovative approach employs a gray-box test generation strategy, using coverage as feedback to heuristically mutate previous high-quality inputs and generate new stimuli. It has demonstrated a remarkable ability to automatically complete test plans, achieving both code and functional coverage closures [2, 14, 28]. Furthermore, fuzzing methods have enabled researchers to uncover previously hidden vulnerabilities [3, 10, 13], some even within industrial hardware designs [27]. There is significant potential to apply hardware fuzzing in real-world industrial hardware development.

However, a key question remains: What is the current gap between hardware fuzzing and its practical application in industrial verification? In this paper, we aim to address this topic by summarizing and exploring two sub-questions, as illustrated in Figure 1.

- What types of hardware fuzzing methods are compatible with industrial verification?
- What additional support is needed from the industrial verification environment to facilitate hardware fuzzing?

For the first sub-question, we review recent influential research on hardware fuzzing and find that many of these works are actually not compatible with industrial verification. We analyze their compatibilities from three dimensions: simulation platform, coverage metric, and mutation strategy, which will be detailed in Section 3. In each dimension, we identify the most appropriate approach for industry. Firstly, the framework should utilize industrial simulators rather than acceleration platforms. Secondly, it should focus on achieving traditional coverage metrics that align with real-world test plans. Finally, it should eliminate the effort of creating design-specific mutation strategies for each individual design.

After summarizing the first sub-question, we turn to the second: Can current industrial hardware verification tools efficiently support hardware fuzzing? We select a hardware fuzzing work that aligns with our proposed standards and conduct a quantitative analysis, which will be detailed in Section 4. Our findings indicate that existing industrial verification tools are not adequately equipped to support hardware fuzzing, resulting in significant inefficiencies. In fact, over 90% of the performance loss is attributed to the lack of necessary support from the industrial verification environment.

We identify three key limitations in current industrial tools that hinder the efficiency of hardware fuzzing in practical applications. Firstly, existing industrial simulators lack support for efficient coverage collection. In fuzzing iterations, coverage is collected frequently as feedback and an inefficient coverage interface wastes a large portion of time. Secondly, industrial tools do not facilitate efficient mutation operations. Current hardware fuzzing works mainly rely on software fuzzers, which are time-consuming due to complex mutation processes and high communication overheads. Thirdly, industrial tools do not incorporate efficient parallel mechanisms for hardware fuzzing, leaving significant room for optimization.

To overcome the limitations listed above, we develop a prototype hardware fuzzing environment called HwFuzzEnv, with three simple but effective enhancements. HwFuzzEnv incorporates essential support for hardware fuzzing within the hardware verification environment and associated tools. The experiments demonstrate that our enhancements provide a significant speedup for previous hardware fuzzing practices in industrial applications.

In summary, our main contributions are as follows:

- We assess the compatibilities of current hardware fuzzing methods for industrial verification and summarize our criteria that determine their suitabilities for industry use.
- We identify the limitations in industrial tools that do not support hardware fuzzing efficiently. To the best of our knowledge, we are the first to address the lack of support for hardware fuzzing in industrial verification environments.
- We develop a prototype industrial hardware fuzzing environment named HwFuzzEnv[1], which incorporates simple but effective enhancements for each stage of hardware fuzzing, including coverage collection, mutation, and simulation.
- We evaluate the efficiency and utility of HwFuzzEnv. With the support of HwFuzzEnv, RtlFuzzLab [6] achieves a speedup of up to 21.5× (on average 12.7×) with a single thread, and up to 621× (on average 361×) with 64 threads.

2 BACKGROUND

2.1 Industrial Hardware Verification

In the industrial hardware development lifecycle [5, 19], engineers use hardware description languages (HDLs) at the register-transfer level (RTL) to describe microarchitectural modules. Electronic design automation (EDA) tools synthesize these RTL modules into gate-level designs, which are then mapped to the transistor level and eventually to physical layout for manufacturing.

During hardware development, writing HDL code at the RTL is error-prone. As a result, an overwhelming 60% of a project's total effort is dedicated to verification [7]. Simulation-based verification is a key technique in industrial hardware verification, which necessitates writing input test cases and uses hardware simulators like Synopsys VCS [25] or Siemens Questa [23] to compare the design's outputs with expected results. Hardware coverage is crucial in this process, as it indicates how thoroughly the design has been tested.

When striving to achieve code and functional coverage closures, random testing fails to reach numerous corner cases within a limited time. Manual test case creation, while precise, is labor-intensive and inefficient. As a result, the research community has sought to automate this process through coverage-directed test generation methods [12]. Among these, hardware fuzzing has emerged as a particularly promising technique.

[1]The project is available at https://github.com/magicYang1573/fast-hw-fuzz.

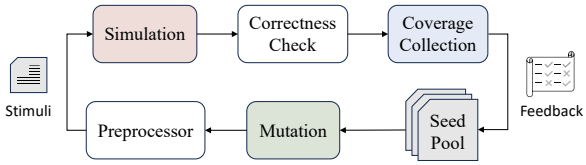

Figure 2: Basic hardware fuzzing workflow.

2.2 Hardware Fuzzing

Fuzzing, originally a concept derived from software testing, denotes the process of testing with minimal or no internal knowledge during the verification phase [18]. Among the various fuzzing techniques, coverage-guided mutational fuzzing stands out for effectively identifying correctness bugs and security vulnerabilities [11, 29]. Recently, researchers have tried to apply the concepts behind software fuzzing to the hardware verification domain.

Figure 2 shows the basic workflow of hardware fuzzing. The fuzzing process starts with a selection from the seed pool. The seed undergoes mutation to generate new inputs, which are then translated into DUT-compatible stimuli by a preprocessor. The HDL simulator runs the simulation and collects coverage data as feedback. Inputs revealing new hardware states (*i.e.,* new coverpoints) are retained as seeds for future fuzzing iterations. This process repeats until no new states emerge for a set number of iterations. Throughout, assertions or golden-model cross-checkings ensure the DUT's correct behaviors. In the following section, we categorize existing hardware fuzzing studies and evaluate their compatibilities with the industrial hardware verification environment.

3 INDUSTRY-COMPATIBLE HARDWARE FUZZING

The workflow of hardware fuzzing consists of three main parts: a HDL simulation platform for correctness checking, a coverage mechanism for feedback collecting, and a mutation engine for input space exploring, as shown in Figure 2. Based on the different implementation strategies of each stage, we categorize existing fuzzing research from three dimensions, as depicted in Figure 3. Each type of methods is evaluated for its advantages and potential limitations in the context of industrial verification.

3.1 Simulation Platform

In hardware fuzzing, a simulation platform is essential for observing hardware behaviors and collecting coverage metrics. Traditional hardware simulators, such as Synopsys VCS [25] and Siemens Questa [23], often require significant time to run simulations, particularly when dealing with complex hardware designs.

To accelerate the fuzzing process, some studies use FPGA platforms for hardware simulation [2, 10, 14]. This approach necessitates substantial modifications to normal fuzzing operations. One challenge is that FPGA implementations can only collect coverage data at the netlist level, which is often incompatible with existing industrial test plans. Besides, utilizing FPGAs efficiently poses technical difficulties for industrial verification engineers, particularly in terms of the communication between the CPU-based fuzzer and the FPGA-based simulator. Similar challenges also exist when using

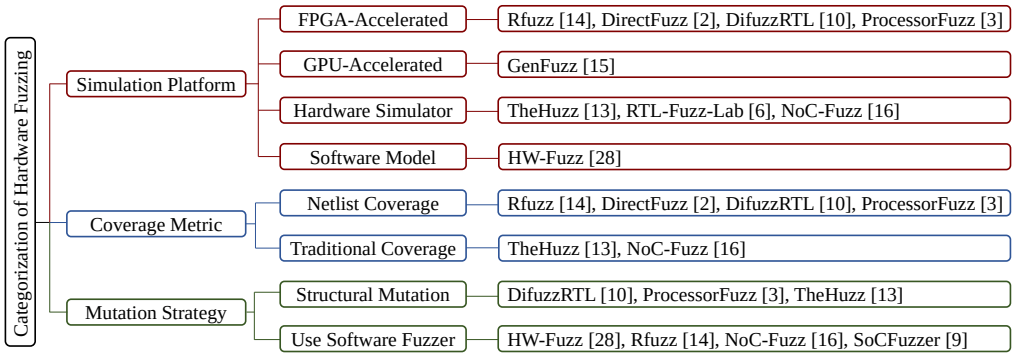

Figure 3: A taxonomy of hardware fuzzing studies with representative examples.

GPUs for simulation acceleration [15], making these acceleration platforms less suitable for industrial hardware verification.

An alternative approach involves fuzzing hardware in the software domain by translating hardware RTL designs into software models using open-source tools like Verilator [1]. It allows for the application of software fuzzers [28] to directly verify the RTL software model. However, this method faces challenges in ensuring the software model's equivalence to the original hardware design, which is unacceptable in industrial verification [8].

Therefore, in industrial hardware verification, using a standard hardware simulator as the simulation platform is necessary, even though it may not provide the fastest execution speed. Such simulators offer mature support for coverage collection, making hardware fuzzing compatible with existing industrial workflows.

3.2 Coverage Metric

Hardware fuzzing is primarily guided by coverage. Initially, netlist-level coverage metrics are the focus of hardware fuzzing studies [2, 3, 10, 14]. For example, RFuzz [14] introduces muxiplexer coverage, treating the select signal of each 2:1 multiplexer as a distinct coverpoint. DifuzzRTL [10] employs register coverage by using the states of CPU control registers as coverpoints.

However, industry engineers often show limited interest in these novel coverage types, as they are not included in the test plans [8, 13]. They prefer verification methods and coverage metrics for RTL, the same level at which hardware designs are written. Moreover, translating these new coverage types into traditional metrics is challenging. It is hard to determine whether a line of code has been thoroughly tested based on multiplexer or register coverage alone.

To bridge this gap, some studies have developed hardware fuzzing frameworks that align with traditional coverage metrics [13, 17]. This makes hardware fuzzing more compatible with industrial verification plans and also promotes the use of industrial hardware simulators in hardware fuzzing workflows.

3.3 Mutation Strategy

Mutation strategy is crucial for efficient space exploration in hardware fuzzing. By mutating high-quality seed inputs, newly generated test cases progressively uncover hardware behaviors. Typically, hardware designs necessitate valid stimuli that adhere to specific protocols or standards, indicating that mutation engines must produce diverse inputs while ensuring their compliance with the design's interface and functionality.

Some studies implement structured and grammar-aware mutation strategies tailored to the input format [3, 10, 13]. For instance, when fuzzing a CPU, the mutation process is customized for the specific instruction set architecture, generating valid CPU instruction sequences. However, this approach is overcomplicated for verification engineers, which requires considerations of the instruction formats as well as the evolutionary algorithms. Given that developing such a complex mutation engine is not an ordinary skill for industrial engineers, and considering the need of unique mutation strategies for each design, it is impractical to apply structural mutation-based fuzzing to industrial verification workflows.

A more accessible and general solution is to use mature software fuzzing tools [11, 29] as mutation engines [4, 6, 9, 14, 28]. This approach requires an additional translation step to convert the generated inputs from software fuzzers into valid hardware stimuli. For instance, AFL [29] generates input files in a binary format. To produce valid stimuli, a hardware fuzzing grammar is designed to map the binary streams to specific hardware transactions. This allows fuzzer to use the same mutation flow across different hardware designs. Consequently, verification engineers only need to focus on the mapping strategies, while mature software fuzzers provide powerful heuristic exploration capabilities, rendering hardware fuzzing a more user-friendly method for industrial verification.

In summary, we categorize existing hardware fuzzing studies from three dimensions and analyze their suitabilities for integration into industrial hardware verification workflows. We outline three criteria to help identify hardware fuzzing methods that are compatible with existing industrial verification:

- Firstly, it should use industrial hardware simulators rather than acceleration platforms with FPGAs or GPUs.
- Secondly, it should aim at achieving traditional coverage metrics which are included in industrial test plans.
- Thirdly, it should use a grammar-agnostic mutation strategy to avoid the need of design-specific structural mutations.

4 ANALYSIS OF INDUSTRIAL GAPS

In this section, we investigate whether existing industrial hardware verification tools efficiently support hardware fuzzing. If not, what additional support is needed? We use RtlFuzzLab [6] as a case study. This open-source framework meets our defined industrial criteria well and implements two of the most influential hardware fuzzing algorithms [14, 28] in recent years, offering representative

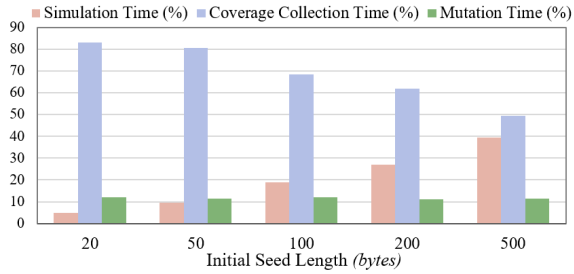

(a) I2C

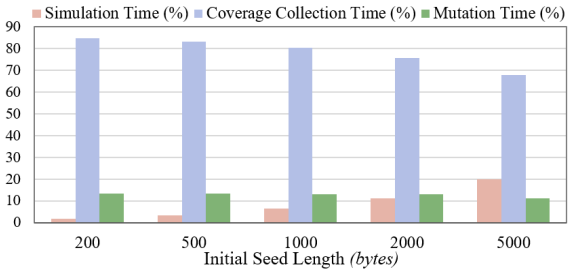

(b) Rocket Core

Figure 4: Time distribution across hardware fuzzing stages.

and referential significance. Our analysis below evaluates its performance from a time perspective, and we aim to identify the reasons for any inefficiencies observed.

4.1 Time Distribution of Hardware Fuzzing

We evaluate the performance of the hardware fuzzing framework on two designs of different scales: a TLI2C IP comprising several thousand lines of code [24] and the Rocket Core with over 40,000 lines of code [22]. In hardware fuzzing, the initial seed is crucial, as it affects the generation of new inputs, the simulation time for each fuzzing iteration, and the overall frequency of fuzzing. Therefore, we conduct our experiments with initial seeds of varying lengths.

As illustrated in Section 3, hardware fuzzing mainly consists of three stages: coverage collection, mutation, and simulation, with their time distribution shown in Figure 4. Counter-intuitively, the coverage collection and mutation stages occupy a significant portion of the total time, exceeding 90% when the initial seed length is small. As the seed length increases, the simulation time grows, while the proportion of time spent on coverage collection and mutation remains over 50%. The results indicate that existing industrial verification tools are not adequately equipped to support hardware fuzzing, resulting in significant time inefficiencies.

4.2 Industrial Gaps for Hardware Fuzzing

Coverage Collection Inefficiency. Existing fuzzing works use Synopsys VCS [25] or Verilator [1] to collect coverage information at simulation runtime. VCS maintains a structured and hierarchical coverage database for engineers to extract coverage reports [26], while Verilator outputs coverage report in a text format, including each coverpoint's definition and hit count [1]. Both VCS and Verilator only support a detailed and structured coverage collection mechanism, which is untailored and redundant for hardware fuzzing. Specifically, this mechanism is designed for long-term simulation, providing engineers with comprehensive reports for analysis. In contrast, hardware fuzzing requires frequent coverage

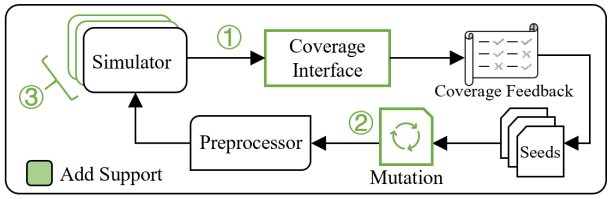

Figure 5: Overview of our fuzzing environment prototype.

updates in each iteration, which necessitates a more sketched and efficient coverage interface provided by industrial simulators.

Mutation Inefficiency. Many hardware fuzzing studies use mature software fuzzers like AFL [29] as their mutation engine [6, 14, 17, 28]. However, this can slow down the fuzzing process for two reasons. Firstly, AFL requires inter-process communications with the hardware simulation environment, which is time-consuming. Secondly, AFL utilizes complex mutation operations (*e.g.,* bitflip, havoc, trim, splice *etc.*) to mutate structured software inputs in a grammar-agnostic manner. Given the generally simpler nature of hardware interface protocols and the potential of designing hardware fuzzing grammars [28] to generate valid stimuli, these mutation operations can be simplified. Considering these factors, there is a clear need of an efficient mutation engine for hardware fuzzing, which is seamlessly integrated within the industrial simulation environment.

Serial Execution Inefficiency. As illustrated in 4.1, the simulation stage only occupies a relatively small fraction of time in the fuzzing process. Therefore, most previous methods are restricted to a single-input fuzzing mechanism that runs one simulation thread at a time, overlooking the potential advantages of batch simulation with multiple threads. Hardware fuzzing is not strictly serial and can benefit greatly from parallel execution. After the coverage collection and mutation stages are supported and optimized, we re-examine the entire hardware fuzzing process, trying to enable the simulation environment to efficiently support hardware fuzzing parallelism and fully exploit the capabilities of multi-core processors.

5 HARDWARE FUZZING ENVIRONMENT PROTOTYPE

The analysis in Section 4 shows that current industrial environment cannot adequately support hardware fuzzing, leading to significant time inefficiencies. Therefore, focused on the three inefficiency reasons above, we develop an industrial environment prototype, HwFuzzEnv, which enhances support for hardware fuzzing. Figure 5 illustrates the overview of our prototype. It introduces three simple but effective enhancements. Firstly, a sketched coverage interface is proposed for the hardware simulator to improve coverage collection efficiency (① in Figure 5). Secondly, a simplified mutation engine is developed within the hardware simulation environment to support efficient mutation (② in Figure 5). Lastly, a multi-thread pipelined mechanism is proposed to support efficient parallelism and speedup the whole fuzzing workflow (③ in Figure 5).

5.1 Sketched Coverage Interface

Previous analysis reveals that the inefficiency of coverage collection arises from industrial simulators' lack of support for an efficient coverage interface, which is essential for hardware fuzzing. Generally, hardware fuzzing engines require only a coverage vector (*i.e.,* a vector of coverpoints' hit counts) as feedback for heuristic

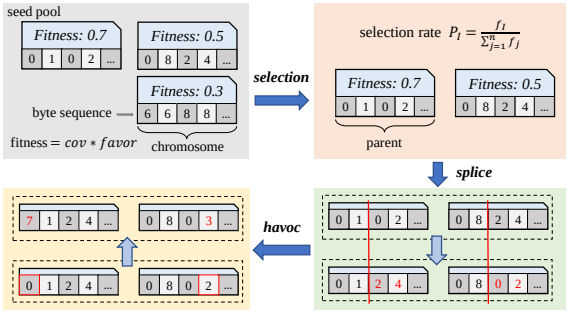

Figure 6: Simplified mutation engine in HwFuzzEnv.

mutation [2, 10, 13, 14, 17, 28]. The specific locations and order of coverpoints in the source code do not influence the performance of fuzzing. To support efficient industrial hardware fuzzing, simulators need an interface that can directly acquire such a coverage vector instead of parsing comprehensive coverage reports during fuzzing. Since we do not have access to the source code of commercial simulators, we investigate the open-source Verilator [1] to implement the sketched coverage interface in HwFuzzEnv.

Verilator instruments coverage monitors in the abstract syntax tree (AST) nodes generated from the hardware design. These monitors trigger coverage counters when specific simulation conditions are met. Each counter is then mapped to an element within a coverage vector, incrementing the corresponding value continuously throughout the runtime. We modify a few lines of source code in Verilator to directly access this coverage vector. We have guaranteed that the coverage vector derived from our sketched coverage interface is equivalent to the one obtained from the original coverage interface. The only difference lies in the order of the coverpoints within the vector, which is unnecessary for hardware fuzzing.

5.2 Simplified Mutation Engine

According to analysis in Section 4, it is necessary to integrate a lightweight mutation engine into the hardware simulation environment. This mutation engine should perform effective and efficient mutation operations. Figure 6 illustrates the simplified mutation engine for hardware fuzzing in HwFuzzEnv.

In the mutation flow, each chromosome is represented as a sequence of bytes, which allows us to define a parametric generator similar to the one proposed by Zest [20]. The parametric generator is a function that takes a sequence of untyped parameters and produces a structured input according to the hardware fuzzing grammar. A byte serves as the basic unit of the parametric generator, enhancing the granularity of mutation operations while reducing the time required.

Selection: The selection process prioritizes individuals with higher fitness from the input pool for future mutation. The fitness function takes into account two primary factors. Firstly, it considers the coverage rate achieved by the input, favoring those that achieve more coverage. Secondly, it assesses whether the input is uniquely responsible for reaching certain coverpoints. If a coverpoint can only be covered by a specific input seed, the fitness of the seed is multiplied by a $favor$ factor, increasing its chance for selection.

Mutation: By carefully examining previous hardware fuzzing works that use software fuzzer AFL [14, 16, 28] as mutation engine, we find that all the interesting inputs are generated from its havoc and splice operation. Therefore, we only implement these two operations in our simplified mutation engine. The havoc operation randomly modifies byte values in the chromosome, and the splice operation involves a crossover between two selected chromosomes.

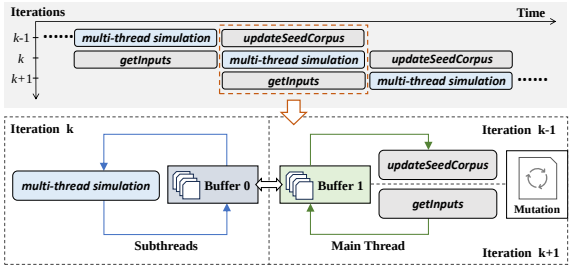

Figure 7: Multi-thread pipelined mechanism in HwFuzzEnv.

5.3 Simulation Parallelization

The simulation workflow is an iterative process that includes the following steps: i) getting input from mutation engine, ii) running HDL simulation, and iii) updating seed corpus using new coverage. After the industrial verification environment has been enhanced to support efficient coverage collection and mutation, we can focus on accelerating the simulation step within HwFuzzEnv. To speed up the simulation process, we implement batch simulations by running multiple simulations concurrently across multiple threads. The framework gets multiple inputs per iteration and dispatches them to individual threads. Each thread operates a hardware simulator in parallel, and the simulation time can be amortized thereby. Once all threads complete the simulation tasks, their coverage information is collected to update the seed corpus in the mutation engine.

The time costs associated with steps i) and iii) are relatively minimal for a single thread. However, as the number of threads increases, these costs escalate due to the multiplication of input acquisitions and coverage updates. Specifically, the time consumed by these two steps begins to approximate that of the simulation stage. Given that these steps do not strictly demand serial processing, we utilize a pipelined approach in HwFuzzEnv to improve the throughput of simulation and hardware fuzzing.

Figure 7 illustrates our multi-thread pipelined fuzzing mechanism, in which different parts of three iterations run concurrently at every timestamp. While each child thread is running simulation for iteration k, the main thread simultaneously `updateSeedCorpus` using the coverage of iteration $k-1$ and `getInputs` for iteration $k+1$. Notably, the seed corpus for the `getInputs` function during iteration $k+1$ does not include the updated information from iteration k (as it is still in simulation). To transfer data as a pipeline, we utilize a ping-pong buffer to store the inputs and coverage data. In this manner, we optimize the utilization of multi-threading and speed up hardware fuzzing further.

6 EXPERIMENTS AND RESULTS

We evaluate HwFuzzEnv based on the open-source hardware fuzzing framework RtlFuzzLab [6], which employs AFL [29] for mutation and Verilator [1] for simulation. We collect the branch coverage as fuzzing guidance in our experiments and the workflows are the same for other code coverage or functional coverage. We comprehensively evaluate on a range of open-source RTL designs, including TileLink Peripheral IPs [24], RISC-V Sodor Cores [21], and the

Table 1: Benchmarks and characteristics.

Source	Name	Input Width	FIRRTL Lines	Cover Points
[24]	I2C	165	2373	245
	PWM	163	2452	181
	UART	164	2416	136
[21]	Sodor1Stage	35	3617	303
	Sodor3Stage	35	4020	314
	Sodor5Stage	35	4088	341
[22]	Rocket Core	239	43856	2378

RISC-V Rocket Core [22]. A detailed list of the benchmarks and their characteristics is available in Table 1.

We conduct our experiments on a Ubuntu 20.04 system with 384GB RAM and two 3GHz Intel Xeon Gold-6248R CPUs, offering 48 physical cores (96 threads) in total.

6.1 Coverage Interface Comparison

In Section 5.1, we implement the sketched coverage interface in Verilator to support efficient coverage collection for hardware fuzzing. A comparative experiment is conducted to evaluate the time efficiency of our sketched coverage interface against the original coverage interface. The fuzzing process is iterated 10,000 times, and the total time spent in coverage collection is shown in Table 2.

The results indicate that, by using the sketched coverage interface of HwFuzzEnv, the coverage collection process achieves a several-hundred-fold speedup. According to the analysis in Section 4.1, coverage collection is the most time-consuming stage in hardware fuzzing. By implementing the sketched coverage interface, we reduce its time consumption to a negligible amount.

6.2 Mutation Engine Comparison

In Section 5.2, we propose a simplified mutation engine in HwFuzzEnv. We conduct a comparative experiment to evaluate its time efficiency against the AFL mutation engine. The fuzzing process is iterated 10,000 times, and the total time spent in mutation is presented in Table 2. The results indicate that our simplified mutation engine achieves a speed dozens of times faster than AFL.

However, it is also crucial to investigate whether this reduction in mutation operations adversely affects the quality of the generated inputs. We compare the number of fuzzing iterations required to achieve K% branch coverage using our mutation engine and AFL, starting with the same initial seeds. For different designs, the value of K is different because the maximum coverage that can be achieved is different. As demonstrated in Table 3, our simplified mutation engine attains equivalent coverage with much fewer fuzzing iterations. The results suggest that our mutation engine improves not only its speed but also the quality of the generated inputs.

6.3 Fuzzing Speed Comparison

After evaluating the efficiency of the sketched coverage interface and simplified mutation engine, we integrate these components and evaluate their combined performance improvement for RtlFuzzLab. In this subsection, we adopt a single-thread fuzzing workflow. The experimental results are presented in Figure 8.

The results indicate that, with the support of HwFuzzEnv, by optimizing the coverage collection and mutation stage, RtlFuzzLab improves the fuzzing speed by 6.2×-21.5× (on average 12.7×). Remarkably, this substantial improvement is achieved with only a single thread. In the following subsection, we extend our investigation to include multi-thread fuzzing experiments.

Table 2: Comparison on coverage collection and mutation efficiency of RtlFuzzLab with and without HwFuzzEnv's support.

Benchmark	Time on Coverage (s)		Time on Mutation (s)	
	Original	Sketched	Original	Simplified
I2C	41.61	0.09 (**462**×)	17.20	0.37 (**46**×)
PWM	34.85	0.08 (**435**×)	16.73	0.41 (**40**×)
UART	31.73	0.07 (**453**×)	16.88	0.35 (**48**×)
Sodor1Stage	51.53	0.05 (**1030**×)	17.86	0.30 (**59**×)
Sodor3Stage	53.27	0.05 (**1024**×)	18.16	0.30 (**60**×)
Sodor5Stage	55.82	0.05 (**1053**×)	18.06	0.31 (**58**×)
Rocket Core	262.73	0.26 (**1010**×)	21.43	0.81 (**26**×)

Table 3: Comparison on mutation quality of AFL's mutation and HwFuzzEnv's simplified mutation.

Benchmark	Iterations for Achieving K% Coverage	
	AFL Mutation	Simplified Mutation
I2C (K=95)	13268	6663 (**2.0**×)
PWM (K=95)	69311	36757 (**1.9**×)
UART (K=80)	81291	50908 (**1.6**×)
Sodor1Stage (K=90)	104384	48301 (**2.2**×)
Sodor3Stage (K=90)	124100	48895 (**2.5**×)
Sodor5Stage (K=90)	127446	76110 (**1.7**×)

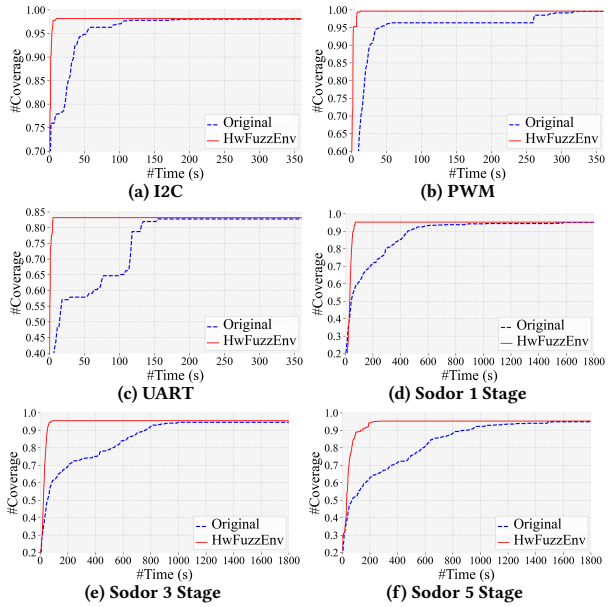

(a) I2C (b) PWM (c) UART (d) Sodor 1 Stage (e) Sodor 3 Stage (f) Sodor 5 Stage

Figure 8: Comparison on single-thread performance of RtlFuzzLab with and without HwFuzzEnv's support.

6.4 Multi-Thread Comparison

We develop a multi-thread pipelined mechanism in HwFuzzEnv to speedup hardware fuzzing further in an industrial environment. Since multi-thread techniques do not modify the original single-thread hardware fuzzing flow, we confine our comparison to the speed of hardware fuzzing (measured in the number of tested input clock cycles per second) in the experiments. Table 4 illustrates the speedups achieved by different multi-thread mechanisms (with or without pipeline) in comparison to single-thread.

When the thread count is small, both the simple multi-thread framework and the pipelined framework achieve a speedup roughly equal to the number of threads utilized. However, as the thread count increases, the speedup of the simple multi-thread mechanism does not exhibit linear growth. In this case, the multi-thread

Table 4: Speedups of different parallel mechanisms of HwFuzzEnv.

DUT	I2C	Sodor1Stage	Rocket Core
Original	0.05×	0.06×	0.16×
1 Thread	1.00×	1.00×	1.00×
4 Thread	3.26×	3.19×	3.57×
Pipelined 4 Thread	3.44×	3.34×	3.83×
16 Thread	9.90×	8.85×	11.24×
Pipelined 16 Thread	12.99×	12.45×	14.60×
64 Thread	15.45×	8.94×	21.11×
Pipelined 64 Thread	31.05×	25.79×	28.26×

pipelined framework demonstrates a noticeable advantage, especially for small designs like I2C or Sodor1Stage. As shown in Table 4, the multi-thread pipelined fuzzing mechanism with 64 threads achieves a speedup of up to 31.0× (on average 28.2×) compared with the single-thread. When compared to RtlFuzzLab without HwFuzzEnv (*i.e.,* Original), it achieves a fuzzing speed of up to 621× (on average 361×) faster. These advancements significantly improve the efficiency of hardware fuzzing for industrial verification.

7 DISCUSSION

As hardware designs scale up and human costs rise, hardware fuzzing emerges as a promising technique for hardware verification automation. This paper aims to advance the integration of hardware fuzzing into industrial verification. Through our analysis and prototype development, we gain a clearer understanding of the current gap between hardware fuzzing and its practical application.

Currently, EDA tools lack sufficient support for hardware fuzzing operations, which substantially limits the efficiency of hardware fuzzing and obstructs its industrial application. For the coverage collection stage, hardware simulators should incorporate a faster mechanism akin to our sketched coverage interface. For the mutation stage, a well-designed mutation engine integrated into the testbench development environment is essential. Additionally, the batch simulation of multiple stimuli is a highly desirable feature for hardware simulators intending to support hardware fuzzing.

For hardware fuzzing to be truly applicable in future hardware verification, the issues mentioned above must be overcome first. Our research could serve as a case study, providing industrial EDA companies with valuable insights to enhance their toolkits.

8 CONCLUSION

In this study, we explore the industrial applications of hardware fuzzing. We begin by summarizing the standards of compatible hardware fuzzing techniques for industrial verification. Then, we analyze the current industrial environment to assess its ability to support these fuzzing methods. Finally, we design a prototype environment HwFuzzEnv to facilitate hardware fuzzing. With this environment, the previous hardware fuzzing method achieves a substantial speedup of several hundred times. Our work provides beneficial perspectives for EDA companies. By adding only a tiny extra effort, EDA tools can efficiently support hardware fuzzing, encouraging them to consider incorporating these features into their upcoming verification toolkits.

ACKNOWLEDGMENTS

This work was partly supported by the National Natural Science Foundation of China (Grant No. 62090021) and the National Key R&D Program of China (Grant No. 2022YFB4500500).

REFERENCES

[1] CHIPS Alliance. 2024. Verilator. *https://verilator.org/guide/latest/*.
[2] Sadullah Canakci, Leila Delshadtehrani, Furkan Eris, Michael Bedford Taylor, Manuel Egele, and Ajay Joshi. 2021. DirectFuzz: Automated test generation for RTL designs using directed graybox fuzzing. In *58th ACM/IEEE Design Automation Conference*. 529–534.
[3] Sadullah Canakci, Chathura Rajapaksha, Leila Delshadtehrani, Anoop Nataraja, Michael Bedford Taylor, Manuel Egele, and Ajay Joshi. 2023. ProcessorFuzz: Processor Fuzzing with Control and Status Registers Guidance. In *IEEE International Symposium on Hardware Oriented Security and Trust*. 1–12.
[4] Chen Chen, Rahul Kande, Nathan Nguyen, Flemming Andersen, Aakash Tyagi, Ahmad-Reza Sadeghi, and Jeyavijayan Rajendran. 2023. HyPFuzz: Formal-Assisted Processor Fuzzing. In *32nd USENIX Security Symposium*. 1361–1378.
[5] Marek Cieplucha. 2019. Metric-driven verification methodology with regression management. *Journal of Electronic Testing* 35, 1, 101–110.
[6] Brandon Fajardo, Kevin Laeufer, Jonathan Bachrach, and Koushik Sen. 2021. RtlFuzzLab: Building A Modular Open-Source Hardware Fuzzing Framework. In *Workshop on Open-Source EDA Technology*.
[7] Harry Foster. 2022. Wilson Research Group Functional Verification Study. *Siemens*.
[8] Weimin Fu, Orlando Arias, Yier Jin, and Xiaolong Guo. 2021. Fuzzing Hardware: Faith or Reality? Invited Paper. In *2021 IEEE/ACM International Symposium on Nanoscale Architectures*. doi:10.1109/nanoarch53687.2021.9642252
[9] Muhammad Monir Hossain, Arash Vafaei, Kimia Zamiri Azar, Fahim Rahman, Farimah Farahmandi, and Mark Tehranipoor. 2023. SoCFuzzer: SoC Vulnerability Detection using Cost Function enabled Fuzz Testing. In *Design, Automation & Test in Europe Conference & Exhibition*. doi:10.23919/date56975.2023.10137024
[10] Jaewon Hur, Suhwan Song, Dongup Kwon, Eunjin Baek, Jangwoo Kim, and Byoungyoung Lee. 2021. DifuzzRTL: Differential fuzz testing to find CPU bugs. In *IEEE Symposium on Security and Privacy*. 1286–1303.
[11] The LLVM Compiler Infrastructure. 2018. LibFuzzer: a library for coverage-guided fuzz testing. *https://llvm.org/docs/LibFuzzer.html*.
[12] A Jayasena and P Mishra. 2024. Directed test generation for hardware validation: A survey. *Comput. Surveys* 56, 5, 1–36.
[13] Rahul Kande, Addison Crump, Garrett Persyn, Patrick Jauernig, Ahmad-Reza Sadeghi, Aakash Tyagi, and Jeyavijayan Rajendran. 2022. TheHuzz: Instruction Fuzzing of Processors Using Golden-Reference Models for Finding Software-Exploitable Vulnerabilities. In *31st USENIX Security Symposium*. 3219–3236.
[14] Kevin Laeufer, Jack Koenig, Donggyu Kim, Jonathan Bachrach, and Koushik Sen. 2018. RFUZZ: Coverage-directed fuzz testing of RTL on FPGAs. In *IEEE/ACM International Conference on Computer-Aided Design*. 1–8.
[15] Dian-Lun Lin, Yanqing Zhang, Haoxing Ren, Brucek Khailany, Shih-Hsin Wang, and Tsung-Wei Huang. 2023. GenFuzz: GPU-accelerated Hardware Fuzzing using Genetic Algorithm with Multiple Inputs. In *60th ACM/IEEE Design Automation Conference*. 1–6.
[16] Ruiyang Ma, Jiayi Huang, Shijian Zhang, Yuan Xie, and Guojie Luo. 2024. NoC-Fuzzer: Automating NoC Verification in UVM. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*.
[17] Ruiyang Ma, Huatao Zhao, Jiayi Huang, Shijian Zhang, and Guojie Luo. 2024. An Endeavor to Industrialize Hardware Fuzzing: Automating NoC Verification in UVM. In *Design, Automation & Test in Europe Conference & Exhibition*.
[18] Valentin J.M. Manes, HyungSeok Han, Choongwoo Han, Sang Kil Cha, Manuel Egele, Edward J. Schwartz, and Maverick Woo. 2021. The Art, Science, and Engineering of Fuzzing: A Survey. *IEEE Transactions on Software Engineering*, 2312–2331. doi:10.1109/tse.2019.2946563
[19] AnaI. Molina and Oswaldo Cadenas. 2007. Functional Verification: Approaches and Challenges. *Latin American Applied Research*.
[20] Rohan Padhye, Caroline Lemieux, Koushik Sen, Mike Papadakis, and Yves Le Traon. 2019. Semantic fuzzing with Zest. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis*. doi:10.1145/3293882.3330576
[21] Berkeley Architecture Research. 2014. RISC-V Sodor. *https://github.com/ucb-bar/riscv-sodor*.
[22] Berkeley Architecture Research. 2016. Rocket Chip. *https://github.com/chipsalliance/rocket-chip*.
[23] Siemens. 2023. Siemens Questa. *https://eda.sw.siemens.com/*.
[24] SiFive. 2016. SiFive Blocks. *https://github.com/sifive/sifive-blocks*.
[25] Synopsys. 2023. Synopsys VCS. *https://www.synopsys.com/verification/simulation/vcs.html*.
[26] Synopsys. 2023. VCS Coverage Technology User Guide.
[27] Fabian Thomas, Lorenz Hetterich, Ruiyi Zhang, Daniel Weber, Lukas Gerlach, and Michael Schwarz. 2024. RISCVuzz: Discovering Architectural CPU Vulnerabilities via Differential Hardware Fuzzing. https://ghostwriteattack.com/.
[28] Timothy Trippel, Kang G Shin, Alex Chernyakhovsky, Garret Kelly, Dominic Rizzo, and Matthew Hicks. 2022. Fuzzing Hardware like Software. In *31st USENIX Security Symposium*. 3237–3254.
[29] Michał Zalewski. 2013. American fuzzy lop. *http://lcamtuf.coredump.cx/afl*.