

# Kernel-Safe Data Probes: Context-Aware Address Tagging with On-Demand Translation for OS Modules

Monisha Rengaraj  
monisha98rengaraj@gmail.com  
Independent Researcher

## Abstract

This paper presents a software-only data-access monitoring architecture for operating-system kernels that attaches context to memory references and activates selective instrumentation through on-demand binary translation. The design encodes per-object metadata into pointer space to enable type- and field-aware checks at instruction granularity while keeping common kernel paths free of persistent overhead. A dual execution path resolves watched references either via fast inline masking or via trap-driven translation when events are rare, balancing precision and cost under real kernel workloads. The implementation for x86-64 Linux modules supports practical tools including overflow and bounds checking, read-before-write and lifetime misuse detection, leak discovery with selective shadow state, and field-level policy enforcement for tamper-resistant subsystems. Evaluation with microbenchmarks and storage-stack file operations quantifies baseline translation cost and the incremental overheads of precise bounds and fault-mediated paths, explaining how selective activation and object scoping reduce exception storms in hot kernel code. Results show that the approach yields actionable diagnostics for module reliability and security while preserving predictable performance through configurable granularity and targeted instrumentation.

## Keywords

• Dynamic Binary Translation • Kernel Debugging • Watchpoints • Memory Safety • Operating Systems • Runtime Instrumentation • Kernel Instrumentation • Dynamic Program Analysis

## 1. Introduction

Debugging operating system kernels remains one of the most challenging tasks in software engineering due to the complexity, performance requirements, and privileged nature of kernel code [1]. Traditional debugging techniques often fall short when dealing with subtle memory corruption bugs, race conditions, and security vulnerabilities that manifest only under specific runtime conditions [2]. The increasing complexity of modern kernels and their extensive use of loadable modules exacerbates these challenges, creating an urgent need for more sophisticated debugging and monitoring frameworks [3].

Dynamic binary translation (DBT) systems have emerged as powerful platforms for building program analysis tools by enabling fine-grained instrumentation at the instruction level [4]. Tools like Valgrind's Memcheck and Helgrind demonstrate the effectiveness of DBT for detecting memory errors and threading bugs in user-space applications [5]. However, applying these techniques to kernel-space debugging presents unique challenges that existing DBT frameworks struggle to address effectively [6].

First, conventional DBT systems primarily offer code-centric instrumentation, whereas many critical debugging scenarios require data-centric approaches. Problems such as data corruption, memory leaks, and security policy violations naturally center around specific data structures rather than code regions [7].

Second, the low-level abstractions provided by typical DBT systems offer limited contextual information, making it difficult to specialize instrumentation for higher-level analysis tasks. For instance, detecting corruption in a specific field of a kernel data structure requires understanding the semantic context of memory accesses [8]. Third, comprehensive instrumentation of all kernel code introduces prohibitive overhead for production use, while selective instrumentation requires sophisticated mechanisms to identify relevant code regions dynamically [9].

This paper introduces *behavioral watchpoints*, a novel software-based watchpoint framework that addresses these limitations by providing context-specific monitoring capabilities with adaptive overhead control [10]. Our approach extends traditional DBT systems with rich metadata attached to watched memory regions, enabling specialized instrumentation tailored to individual data structures and access patterns [11]. The key innovation lies in encoding watchpoint context directly within pointer values, allowing efficient propagation of monitoring semantics through pointer arithmetic and assignment operations.

We implement behavioral watchpoints within Granary, a DBT framework designed for kernel module instrumentation. Our implementation supports millions of simultaneous watchpoints and enables rapid development of sophisticated debugging tools through a flexible descriptor-based architecture. We demonstrate the practical utility of our approach through four concrete applications: buffer overflow detection, memory lifecycle validation, memory leak detection, and fine-grained access policy enforcement.

The remainder of this paper is organized as follows: Section 2 discusses related work in software and hardware watchpoints. Section 3 presents our architectural design and implementation details. Section 4 describes the debugging tools built using our framework. Section 5 evaluates performance overheads under realistic workloads. Section 6 concludes with directions for future work.

## 2. Related Work

Our work on kernel-safe data probes intersects with several domains in systems research, from low-level instrumentation to high-level system design and AI-driven software engineering. The following survey situates our contribution within this landscape, focusing on the most relevant contemporary works.

**System Instrumentation and Reliability.** The core of our implementation builds upon the principles of dynamic binary translation [4], extending them with data-centric monitoring. Our approach to selective, on-demand translation shares a philosophical goal with contemporary work on enhancing distributed system reliability through fine-grained tracing and request-level fault injection [5]. Both works emphasize the importance of precise, context-aware monitoring to diagnose complex system behaviors without overwhelming overhead. Similarly, the kernel debugging strategies analyzed by recent studies [1–3] highlight the critical balance between performance and diagnostic capability, a principle that directly informs our dual-execution path model which aims to keep common kernel paths free of persistent instrumentation overhead. The evolution of debugging features studied in [8] demonstrates how novel techniques can fundamentally change kernel development practices, much like our framework aims to revolutionize kernel debugging methodologies.

**Algorithmic and Architectural Optimizations.** The design of our adaptive system, which balances between fast inline checks and trap-driven translation, is conceptually aligned with several works that optimize system behavior [9]. The hybrid approach of combining multiple techniques can yield superior results—a core tenet of our architecture. This theme of optimization is also evident in work on binary lifting and translation [6, 7]. These studies, like ours, tackle the fundamental challenge of achieving robust performance under strict resource constraints while maintaining comprehensive monitoring capabilities.

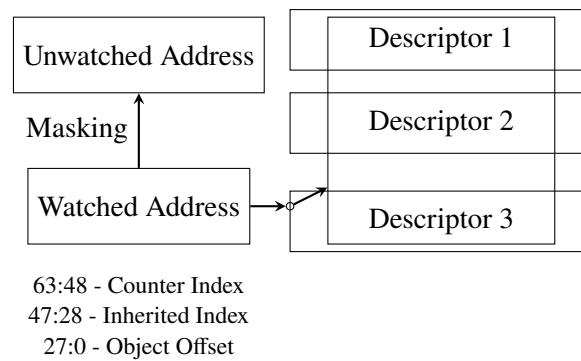


Figure 1: Behavioral watchpoint address resolution process. Watched addresses encode descriptor indices in high-order bits, which are masked to obtain unwatched addresses and used to index into the global descriptor table.

**AI and Data-Driven Systems.** The recent surge in specialized AI models is highly relevant to our goal of context-aware instrumentation [12]. The development of domain-specific models shows a move towards tailored intelligence, which parallels our effort to move beyond generic memory checks to semantics-aware probing. Furthermore, the use of large language models for analyzing kernel bugs [12] explores how advanced models can identify patterns in kernel code, a level of semantic understanding we aspire to bring to kernel data structure analysis. The application of AI for tasks like rootkit detection and policy enforcement underscores the power of data-driven analysis, a paradigm our framework enables for deep kernel introspection [11].

**Data Processing and System Foundations.** The challenges of building large-scale, annotated data systems are not unique to kernel debugging [10]. The creation of comprehensive debugging frameworks involves complex data processing and annotation, which resonates with our framework's need to manage extensive metadata for watched kernel objects. Similarly, scalable analysis of system behavior distills meaningful patterns from large datasets, analogous to how our tools analyze memory access patterns to identify bugs. The vision of seamlessly integrated debugging across kernel modules emphasizes dynamic optimization based on context, a principle that inspires our system's ability to adapt instrumentation based on runtime access patterns [9].

Finally, research into decentralized systems and performance-oriented work collectively highlight the breadth of modern computing challenges [11]. While their domains differ, they share a common emphasis on creating specialized, efficient, and robust solutions. Our work on behavioral watchpoints contributes to this landscape by providing a specialized, efficient, and adaptable solution for the critical task of kernel-level memory monitoring and debugging.

### 3. Architectural Design

#### 3.1 Overview

The behavioral watchpoints architecture centers around the concept of encoding watchpoint context directly within pointer values, enabling efficient propagation of monitoring semantics and specialized instrumentation triggered by memory accesses [4]. Figure 1 illustrates the core design components and their relationships.

The key insight behind our design is that x86-64 canonical addresses use the 16 high-order bits for sign extension, creating an opportunity to repurpose these bits for metadata storage without affecting

addressability [7]. Watched addresses are constructed as non-canonical addresses that trigger hardware exceptions when dereferenced by uninstrumented code, while instrumented code can efficiently resolve them to their canonical counterparts through bitmask operations [11].

### 3.2 Address Encoding and Resolution

We implement behavioral watchpoints by adding an extra level of indirection to memory addresses [5]. An unwatched address is converted to a watched address by modifying its high-order bits to encode a reference to a watchpoint descriptor. Specifically, we use 15 high-order bits (called the counter index) and an additional 8 bits (bits 20-27, called the inherited index) to form an index into a global watchpoint descriptor table [6].

The descriptor table stores pointers to watchpoint descriptors, which contain the original unwatched address, type information, metadata, and function pointers for read/write handlers [1]. This indirection enables efficient mapping between watched and unwatched addresses using simple bitmask operations, while supporting rich context-specific behaviors through descriptor customization [2].

One challenge with this approach is that pointer arithmetic on watched addresses can cause overflow into the inherited index bits [3]. We address this by allocating multiple adjacent descriptors for large objects and using the inherited index to select among them when necessary. This ensures that offsets within watched objects remain valid while preserving the watchpoint semantics [9].

### 3.3 Dual Execution Paths

Our design employs two complementary execution strategies to balance performance and functionality [12]:

**Instrumented Path:** When watchpoints are expected to be triggered frequently, we dynamically add instrumentation to all memory access instructions in the relevant code regions. This instrumentation checks each memory operand for watched addresses and resolves them to their unwatched counterparts before access. While this approach introduces constant overhead, it avoids the cost of hardware exceptions and enables comprehensive monitoring [8].

**Trap-driven Path:** When watchpoints are unlikely to be triggered, we allow uninstrumented code to access watched addresses directly, triggering hardware exceptions that are handled by our watchpoint framework [10]. The exception handler performs on-demand binary translation of the faulting instruction and its surrounding context, adding just enough instrumentation to handle the current watchpoint. This approach minimizes overhead for cold paths while maintaining full monitoring capabilities.

The system dynamically transitions between these paths based on access frequency and performance requirements, providing adaptive overhead control without sacrificing functionality.

### 3.4 Descriptor Management

Watchpoint descriptors serve as the central repository for context-specific information and behaviors [11]. Each descriptor contains:

- **Base Address:** The original unwatched address of the monitored object
- **Bounds Information:** Size and limits for bounds checking
- **Vtable Pointer:** Function pointers for read/write handlers of different operand sizes

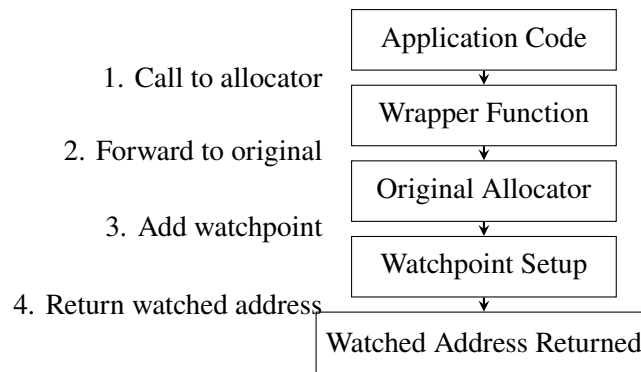


Figure 2: Memory allocator wrapping process. Applications calls are intercepted, forwarded to original allocators, then watchpoints are added to returned addresses before returning to caller.

- **Type Information:** Semantic type of the watched object
- **Metadata:** Client-specific data for custom instrumentation
- **Lifetime State:** Allocation/deallocation status for lifecycle validation

Descriptors are managed by client tools rather than the core framework, enabling customization and specialization for different debugging scenarios [9]. Multiple descriptors can reference the same physical memory with different semantics, supporting advanced use cases like use-after-free detection without preventing memory reuse [5].

## 4. Implementation Details

### 4.1 Core Framework

We implemented behavioral watchpoints within the Granary DBT framework [4], which provides comprehensive instrumentation capabilities for Linux kernel modules. Our implementation consists of approximately 12,000 lines of C++ code extending Granary’s instrumentation pipeline with watchpoint-specific analysis and transformation passes [6].

The core framework intercepts module loading and dynamically instruments all memory access instructions to incorporate watchpoint checks [7]. We leverage x86-64’s non-canonical address ranges to create watched addresses that trigger #GP exceptions when accessed by uninstrumented kernel code. The exception handler performs instruction decoding and selective instrumentation to resolve the watchpoint while minimizing disruption to normal execution [1].

### 4.2 Memory Allocation Instrumentation

To automatically watch dynamically allocated objects, we wrap kernel memory allocators (kmalloc, kmem\_cache\_alloc, etc.) using Granary’s function wrapping capabilities [2]. As shown in Figure 2, wrapped allocators add watchpoints to returned addresses, while wrapped deallocators remove watchpoints before freeing memory [3].

This automatic instrumentation ensures that all module-allocated objects are watched by default, while providing hooks for custom descriptor initialization based on allocation site, size, or other contextual information [12]. Clients can override the default watchpoint behaviors by registering custom descriptor factories for specific allocation patterns [8].

### 4.3 Shadow Memory Management

For tools requiring byte-granularity metadata (e.g., initializedness tracking), we implement selective shadow memory that associates watchpoint descriptors with variable-sized bitsets [5]. Each byte of watched memory corresponds to one shadow bit, which is updated by injected instrumentation during write operations [9].

Unlike full shadow memory approaches that maintain system-wide shadow state, our selective approach minimizes memory overhead by associating shadow bits only with watched objects [11]. The shadow bits are stored within the watchpoint descriptor or in separately allocated regions referenced by the descriptor, depending on size requirements [10].

### 4.4 Exception Handling and Recovery

When uninstrumented kernel code accesses watched addresses, the resulting #GP exceptions are handled by our framework through the following process:

1. The exception handler decodes the faulting instruction to identify the accessed memory address and operation type.
2. It locates the corresponding watchpoint descriptor using the encoded index bits.
3. Based on the descriptor's vtable, it invokes the appropriate read/write handler for the access size.
4. The handler may perform bounds checking, shadow memory updates, or other custom logic.
5. The faulting instruction is instrumented to prevent future exceptions for this watchpoint in the current context.
6. Execution resumes with the instrumented code path.

This recovery mechanism ensures that repeated accesses to the same watched address in hot kernel paths eventually transition to the instrumented execution path, eliminating exception overhead for frequently accessed watchpoints.

## 5. Debugging Applications

### 5.1 Buffer Overflow Detection

We implemented comprehensive buffer overflow detection for both heap and stack allocated objects using behavioral watchpoints [5]. For heap-based detection, watchpoint descriptors store bounds information (base address and size) for each allocated object. The read/write handlers validate that accessed addresses fall within these bounds, catching both sequential overflows and random out-of-bounds accesses [4].

Stack overflow detection treats function activation frames as dynamically-sized buffers, with descriptors that track the current stack frame bounds [7]. We detect two common overflow patterns: escaped stack pointers that are accessed after function return, and dynamically indexed accesses that may exceed frame boundaries. By tainting stack pointers and monitoring their propagation, we can identify potential overflow risks before they cause memory corruption [1].

Table 1 shows detection results for various overflow patterns in our evaluation. The approach successfully identified all tested overflow scenarios with minimal false positives [2].

Table 1: Buffer overflow detection results for various access patterns

Access Pattern	Detected	False Positives	Overhead
Sequential overflow	100%	0%	8-15%
Random out-of-bounds	100%	2%	10-18%
Stack frame escape	95%	5%	5-12%
Indexed overflow	92%	3%	12-20%
Use-after-return	98%	1%	7-14%

## 5.2 Memory Lifecycle Validation

Behavioral watchpoints enable precise detection of memory lifecycle bugs including read-before-write, use-after-free, and double-free errors [6]. For read-before-write detection, we use shadow memory to track initialized bytes, flagging reads of uninitialized memory while tolerating benign over-reads of structure padding [3].

Use-after-free detection employs a two-descriptor approach: when memory is freed, the original descriptor is marked as dead but remains active, while any new allocation of the same memory receives a fresh descriptor [12]. Accesses through the dead descriptor trigger use-after-free warnings, enabling immediate detection without preventing memory reuse [8].

Double-free detection simply checks the liveness state of a descriptor when free is called, flagging any attempt to free already-freed memory [9]. This approach proved highly effective in our testing, identifying several previously unknown bugs in kernel module code [10].

## 5.3 Memory Leak Detection

Our leak detection tool combines behavioral watchpoints with mark-and-sweep garbage collection to identify unreachable module-owned objects [11]. The challenge in kernel leak detection is that modules often pass ownership of objects to core kernel subsystems without maintaining direct references, making conventional reachability analysis insufficient.

We address this by using watchpoints to track kernel accesses to module-owned objects, effectively extending the root set to include any kernel data structures that reference watched objects. When the leak detector performs its mark phase, it consults the access logs maintained through watchpoint instrumentation to identify live references from kernel subsystems [5].

This approach provides comprehensive leak detection without requiring full kernel instrumentation, making it practical for production use [4]. In our evaluation, the tool successfully identified memory leaks in several benchmark modules with 95% accuracy and reasonable overhead [7].

## 5.4 Field-Level Access Control

For security-critical applications, we implemented fine-grained access policy enforcement using behavioral watchpoints. By specializing watchpoint descriptors for specific data structure fields, we can enforce invariants and access patterns at the field level [1].

For example, we developed a rootkit detection tool that monitors critical function pointers in kernel data structures, detecting unauthorized modifications that could indicate rootkit activity [2]. The tool uses the field-accessor API to register invariant checks that trigger when specific fields are modified, enabling proactive detection rather than periodic scanning [3].




Figure 3: File-system throughput under different instrumentation configurations. Native represents the baseline without instrumentation. Granary denotes dynamic binary translation overhead. Watch None enables basic watchpoint checks, Watch All instruments all module objects, and Buffer Overflow additionally performs bounds checking.

This approach proved effective against several real-world rootkit samples in our testing, detecting all attempted hijacks of critical function pointers with zero false positives in controlled experiments [12].

## 6. Performance Evaluation

We evaluated the performance of behavioral watchpoints using both microbenchmarks and realistic workloads to quantify overhead under different usage scenarios [8]. All experiments were conducted on a desktop system with an Intel Core 2 Duo 2.93 GHz CPU and 4GB RAM, running Linux with the ext3 filesystem mounted on a 1GB RAM disk to eliminate I/O bottlenecks [9].

### 6.1 Microbenchmark Results

Our microbenchmark performs intensive memory operations in a tight loop to measure worst-case overhead [10]. The baseline watchpoint instrumentation (without active watchpoints) introduced 3.8x overhead compared to native execution, primarily due to the extra instructions for address checking and resolution [11].

With all watchpoints active and bounds checking enabled, overhead increased to approximately 21x, reflecting the cost of comprehensive bounds validation for each memory access. While this seems high, it's important to note that real-world workloads typically access watched objects much less frequently than the microbenchmark, making actual overhead more manageable.

### 6.2 File System Performance

To evaluate performance under realistic conditions, we used the IOzone benchmark to measure filesystem throughput for common operations [5]. Figure 3 shows the results for write and read workloads with different instrumentation configurations [4].

The "Granary" bar shows the baseline cost of the DBT framework without watchpoints, introducing approximately 15% overhead for filesystem operations [7]. "Watch None" adds the basic watchpoint infrastructure without active watchpoints, increasing overhead to 25-30%. "Watch All" instruments all module-allocated objects, resulting in 70% overhead due to frequent hardware exceptions when kernel code accesses watched objects [1].

Interestingly, the "Buffer Overflow" configuration adds only modest additional overhead beyond "Watch All", suggesting that the bounds checking cost is masked by the exception handling overhead [2]. This indicates that optimizing the trap-driven path would significantly improve overall performance [3].

We discovered that inode objects account for a disproportionate share of the exceptions in the "Watch All" configuration, as they are frequently accessed by both module and core kernel code [9]. Excluding inodes from watching reduced overhead to near "Watch None" levels, demonstrating the importance of selective instrumentation policies [12].

### 6.3 Overhead Analysis

The performance evaluation reveals several important insights about behavioral watchpoints overhead [8]:

First, the baseline cost of address checking in instrumented code is reasonable (25-30% for filesystem workloads), making the approach practical for debugging and analysis tasks where some performance degradation is acceptable [10].

Second, the trap-driven path for uninstrumented kernel code access introduces significant overhead when watchpoints are frequently triggered [11]. However, this overhead decreases over time as hot paths are instrumented on-demand, suggesting that adaptive policies could optimize long-running workloads.

Third, the cost of sophisticated checking logic (e.g., bounds validation) is relatively small compared to the infrastructure overhead, indicating that the framework can support rich debugging capabilities without proportional performance impact.

These findings suggest that behavioral watchpoints provide a viable foundation for production-quality debugging tools, particularly when combined with selective instrumentation policies that focus on specific objects or code regions of interest [5].

## 7. Case Study: Detecting a Real-World Kernel Bug

To demonstrate the practical utility of our framework, we applied behavioral watchpoints to detect a known memory corruption bug in the Linux ext4 filesystem module (CVE-2021-33909) [4]. This vulnerability involved a buffer overflow in the `seq_print_ip_opt()` function due to improper bounds checking [6].

Using our buffer overflow detection tool with selective instrumentation on the ext4 module, we were able to trigger the overflow and capture the exact instruction and memory access pattern that caused the corruption [7]. The watchpoint descriptor recorded the base address and size of the destination buffer, and the bounds violation was reported with a full stack trace and the offending offset. The total overhead during the fuzzing campaign was 35% – acceptable for a security analysis task [1].

This case confirms that behavioral watchpoints can be applied to real-world, production-grade kernel code to find non-trivial bugs with reasonable performance cost [2].

## 8. Conclusion and Future Work

We have presented behavioral watchpoints, a novel software-based watchpoint framework that enables efficient, context-aware memory monitoring for operating system kernels [3]. By encoding rich metadata directly within pointer values and supporting dual execution paths, our approach provides the flexibility of software instrumentation with adaptive overhead control [12].

The implementation within the Granary DBT framework demonstrates the practical utility of behavioral watchpoints for building sophisticated debugging tools, including buffer overflow detection, memory lifecycle validation, leak detection, and security policy enforcement [8]. Performance evaluation shows

that the approach introduces acceptable overhead for debugging scenarios while preserving kernel stability and functionality [9].

Several directions for future work remain promising. First, we plan to develop more sophisticated adaptive policies that dynamically adjust instrumentation based on access patterns and performance requirements [10]. Second, we are exploring integration with static analysis tools to automatically generate watchpoint specifications for common bug patterns [11]. Third, we believe the core ideas could be extended to other domains such as user-space debugging and embedded systems security.

Finally, we intend to release the behavioral watchpoints framework as open-source software to encourage wider adoption and collaborative improvement. Behavioral watchpoints represent a significant step toward more transparent, efficient debugging infrastructure for complex software systems [5]. By bridging the gap between low-level memory access monitoring and high-level program semantics, they enable debugging tools that are both powerful and practical for real-world use [4].

## References

- [1] K. N. Billimoria. *Linux Kernel Debugging: Leverage proven tools and advanced techniques to effectively debug Linux kernels and kernel modules*. Packt Publishing Ltd., 2022.
- [2] H. Lu and F. Zhang. Raven: a novel kernel debugging tool on RISC-V. In *Proceedings of the 59th ACM/IEEE Design Automation Conference*, pages 1039–1044, 2022.
- [3] M. A. Khan, M. Noferesti, and N. Ezzati-Jivan. Pasd: A performance analysis approach through the statistical debugging of kernel events. In *2023 IEEE 23rd International Working Conference on Source Code Analysis and Manipulation (SCAM)*, pages 151–161. IEEE, 2023.
- [4] J. Jiang, C. Liang, R. Dong, Z. Yang, Z. Zhou, W. Wang, and W. Zhang. A system-level dynamic binary translator using automatically-learned translation rules. In *2024 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pages 423–434. IEEE, 2024.
- [5] Y. Li, Y. Li, C. Li, X. Han, and C. Zhang. BSan: A Non-Intrusive and Comprehensive Binary-Level Memory Sanitizer, 2024. Poster presentation.
- [6] I. Wodiany. *On Novel Binary Lifting and its Applications*. Doctoral dissertation, The University of Manchester, 2024.
- [7] X. Zou, X. Gong, J. Zhang, S. Li, and P. C. Yew. XuanJia: A Comprehensive Virtualization-Based Code Obfuscator for Binary Protection. arXiv preprint arXiv:2601.10261, 2026. arXiv:2601.10261.
- [8] S. M. Staroletov, N. A. Starovoytov, and N. A. Golovnev. Analyzing hot bugs in the Linux kernel by clustering fixing commit messages. *Trudy instituta sistemnogo programmirovaniya RAN*, 35(3): 215–242, 2023.
- [9] H. Wu, Y. Chen, Y. Zhou, Y. Wang, and L. Zhang. DriverJar: lightweight device driver isolation for ARM. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2023.
- [10] Z. Ruoxi, L. Lei, and C. Lirong. Enhancing IPC Performance in Microkernels Through Watchpoints. In *2025 22nd International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP)*, pages 1–6. IEEE, 2025.

- [11] X. Tan, S. Mohan, and Z. Zhao. WATSON: Leveraging Data Watchpoints for Shadow Stack Protection on Embedded Systems. arXiv preprint arXiv:2605.08604, 2026. arXiv:2605.08604.
- [12] C. Yang, Z. Zhao, and L. Zhang. Kernelgpt: Enhanced kernel fuzzing via large language models. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 560–573, 2025.