

# Enhanced Statistical Debugging for Adaptive Monitoring

---

Mohammed Adib Khan  
ak19qp@brocku.ca

Dr. Morteza Noferesti  
mnoferesti@brocku.ca

Dr. Naser Ezzati-Jivan  
nezzati@brocku.ca

Department Of Computer Science  
Brock University, Canada

# Introduction


---

- ❑ Monitoring and debugging functions is essential, particularly when kernel-level operations experience long wait times.
- ❑ Some examples of kernel-level waiting time metrics could derive from:
  - Syscalls
  - Block\_rqs
  - Sched\_Switch
  - Irq\_handler
- ❑ Prolonged delays in kernel-level operations can severely impact overall system performance and could potentially lead to system failures.
- ❑ Statistical debugging techniques offer powerful solutions.

# Problem Statement & Objectives

---

- ❑ The primary goal is to identify candidate functions and function paths for adaptive or selective tracing, crucial for diagnosing performance issues.
- ❑ To achieve this, we have two key objectives:
  - Identifying problematic functions
  - Determining potential problematic function paths



# Objective 1 – Identifying Problematic Functions

---

- ❑ The first step involves pinpointing the functions responsible for kernel-level activity delays by examining correlation.
- ❑ Our focus lies on detecting application-related bugs within specific functions.
- ❑ If a function and a performance bug shows a high correlation, our method will identify that function.

## Objective 2 – Determining Function Paths

---

- ❑ Simply identifying the functions does not suffice for our ultimate objective.
- ❑ The next step is to ascertain the function path leading to these problematic or intriguing functions.
- ❑ This path is indispensable for adaptive tracing, as merely tracing the function itself would lack sufficient context for effective analysis.

$$Path^i = \langle F^{main}, F^1, F^2, \dots, F^i \rangle$$

# Background – Statistical Debugging

- ❑ Statistical debugging is a method that identifies defects by examining the correlation between program states and failures.
- ❑ It collects program execution data, isolating elements most related to failures by comparing 'successful' and 'failing' runs.
- ❑ Typically, it focuses on line-code level and conditions within the code.
- ❑ It is effective for consistently reproducible issues, aiding in bug localization in complex and large codebases.
- ❑ It provides vital clues to developers to understand and fix the defects.

	<b>■: covered statements</b>	<b>x</b>	3	1	3	5	5	2	
1	<code>int middle(x, y, z) {</code>	<b>y</b>	3	2	2	5	3	1	
2	<code>int x, y, z;</code>	<b>z</b>	5	3	1	5	4	3	
3	<code>int m = z;</code>		■	■	■	■	■	■	3
4	<code>if (y &lt; z) {</code>		■	■	■	■	■	■	4
5	<code>if (x &lt; y)</code>		■	■	□	□	■	■	5
6	<code>m = y;</code>		□	■	□	□	□	□	6
7	<code>else if (x &lt; z)</code>		■	□	□	□	■	■	7
8	<code>m = y;</code>		■	□	□	□	□	■	8
9	<code>} else {</code>		□	□	■	■	□	□	9
10	<code>if (x &gt; y)</code>		□	□	■	■	□	□	10
11	<code>m = y;</code>		□	□	■	□	□	□	11
12	<code>else if (x &gt; z)</code>		□	□	□	■	□	□	12
13	<code>m = x;</code>		□	□	□	□	□	□	13
14	<code>}</code>		□	□	□	□	□	□	14
15	<code>return m;</code>		■	■	■	■	■	■	15
16	<code>}</code>		✓	✓	✓	✓	✓	✗	

Image credits: Jones and Harrold 2005

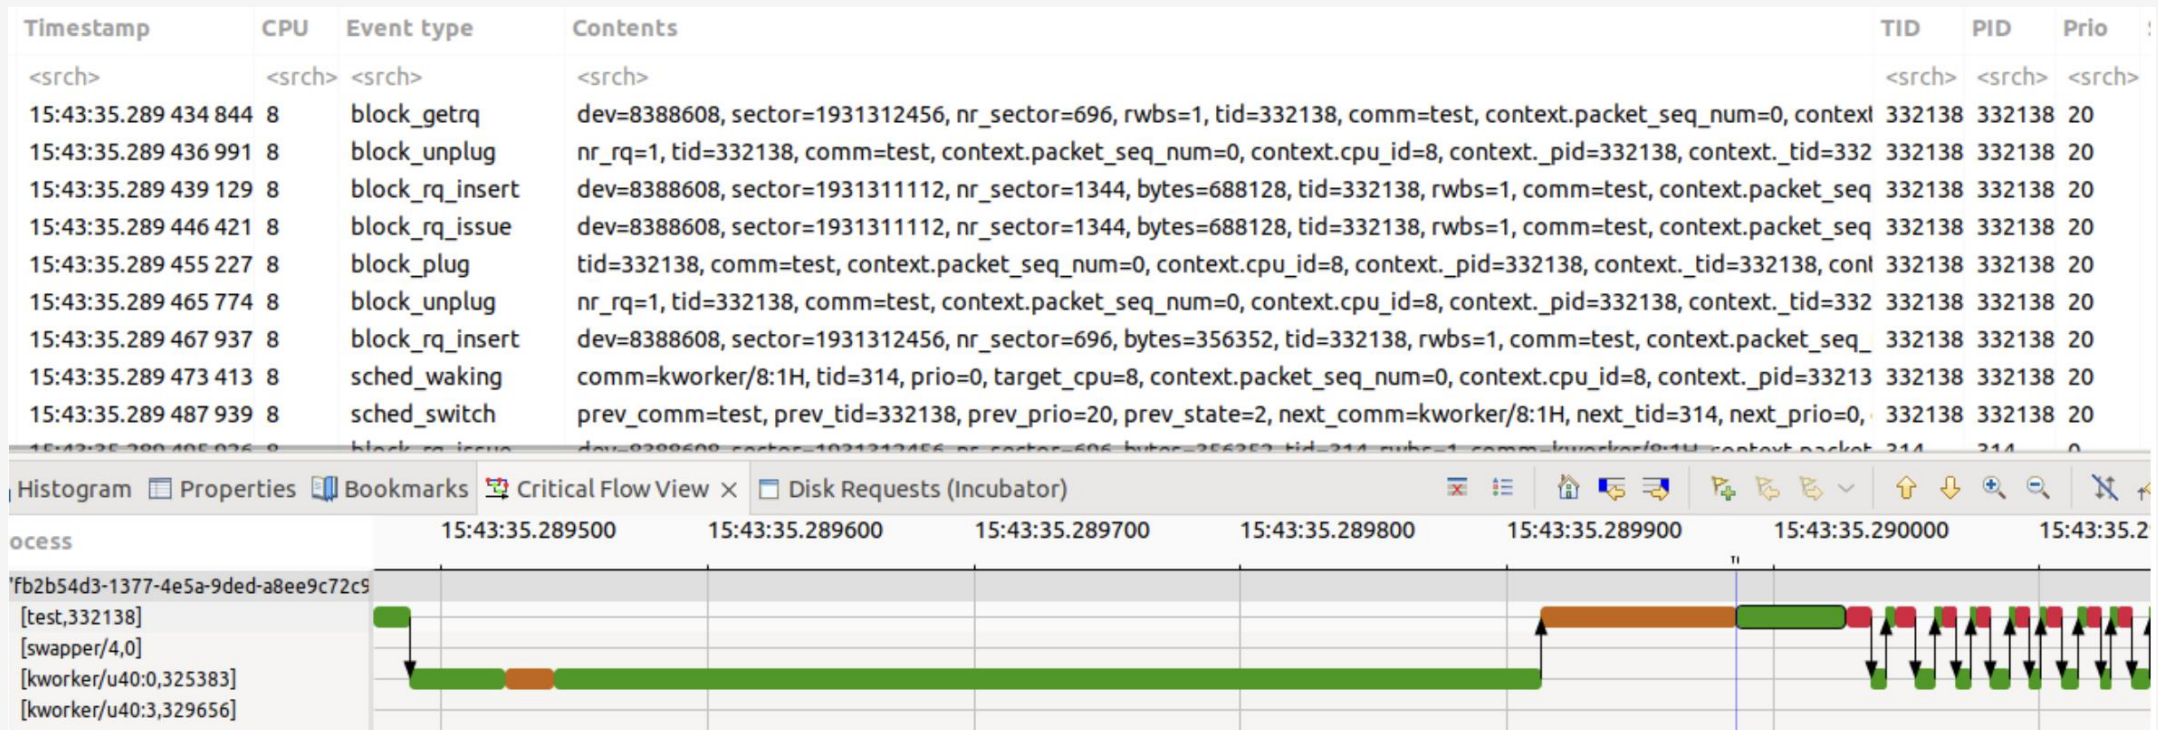
# Background – Enhanced Statistical Debugging

---

- ❑ Enhanced Statistical Debugging is our refined version of traditional statistical debugging.
- ❑ Traditional statistical debugging analyses bug correlations at a line-code level, examining conditional constructs within the program code.
- ❑ In contrast, our Enhanced Statistical Debugging shifts focus to a function level, examining the correlation between specific function execution and performance issues.
- ❑ This function-level analysis aids in identifying potential bottlenecks, offering a more precise tool for diagnosing and resolving performance concerns.

# Methodology – Monitoring Kernel Trace

- ❑ Our approach includes monitoring prolonged activities in the kernel trace of our application, such as waiting times (for net, disk, CPU, etc.), syscalls, interrupts, etc.
- ❑ For each protracted activity, we ensure to gather enough stack trace data to perform meaningful analysis on them.





# Methodology – Enhanced Statistical Debugging

---

- ❑ These stack trace samples undergo analysis using statistical debugging.
- ❑ We accumulate ample data for both faulty and successful activities.
- ❑ Then, via statistical debugging, we strive to identify any correlation between the application's functions and these lengthy or faulty activities.
- ❑ We define fail runs as wait times exceeding mean + std, and the opposite for success. Other thresholds could also be used depending on the activity.
- ❑ Success/fail predicates correspond to the last function in the call stack.
- ❑ Success(observed) and failed(observed) corresponds to the functions which are observable at any place in the call stack.

$$Failure(P) = F(P)/(S(P) + F(P))$$

$$Observation(P) = \frac{F(P_{observed})}{S(P_{observed}) + F(P_{observed})}$$

$$Increase(P) = Failure(P) - Observation(P)$$

# Methodology – Function Path Sequence Mining

- ❑ Upon identification of the functions, we conduct further analysis (path sequence mining) to locate the most common paths leading to these kernel-level metrics delays.
- ❑ These functions, together with their associated paths, become candidates for tracing/logging.
- ❑ This targeted approach to tracing allows us to effectively concentrate our resources, negating the need for comprehensive tracing.




Image credits: Hassani, Marwan, et al. 2019

# Methodology – Adaptive Monitoring

---

- ❑ Adaptive monitoring uses a dynamic, targeted approach to system tracing and debugging.
- ❑ It builds on insights from statistical debugging and path mining to focus on problematic functions and paths.
- ❑ Unlike exhaustive logging, this method saves resources by tracing only areas of interest.
- ❑ By reducing data volume and overhead, it prevents crucial information from being overlooked.

# Case Study – Firefox

- ❑ We applied our methodology to a case study of a performance bug in Firefox.
- ❑ Our method successfully identified the correct function paths causing the bug.

**Open** Bug 1565019 Opened 4 years ago Updated 3 months ago

## High CPU usage in parent process main thread when loading tripadvisor.ca, for over 700ms

▸ **Categories** (Core :: Networking, defect, P2)


▸ **Tracking** (NEW bug found in Firefox 70 which should be worked on in the next release/iteration)

▸ **People** (Reporter: mstange, Unassigned)

▸ **References** (Depends on 1 open bug, Blocks 1 open bug)

▸ **Details** (Whiteboard: [necko-triaged])

Bottom ↓ Tags ▼ Timeline ▼

 **Markus Stange [mstange]** Reporter  
Description • 4 years ago

Here's a profile of loading tripadvisor.ca on my machine: <https://perfht.ml/2JvpP9n>

A lot of work seems to be happening in the parent process main thread, most of it triggered by initiating network requests. It would be good to reduce the amount of work and/or move it to other threads.

(I'm not expecting any immediate remedies; I'm filing this bug mostly to have an example that I can point others at, and so that we have a profile that we can compare ourselves to, in the future as this improves.)

# Case Study – Firefox

```
Downloads — -zsh — 208x24
Last login: Tue May 30 15:29:49 on ttys000
adib@Adibs-MacBook-Pro ~ % cd downloads
adib@Adibs-MacBook-Pro downloads % logmine paths_names.csv -m0.2 -sdesc
2058 -0.002808069, "getifaddrs_internal>>event_base_loop>>base::MessagePumpLibevent::Run>>MessageLoop::Run>>base::Thread::ThreadMain>>ThreadFunc>>set_alt_signal_stack_and_start>>start_thread"
1525 -0.001091107, "putspent>>rayon_core::registry::ThreadBuilder::run>>std::sys_common::backtrace::_rust_begin_short_backtrace>><core::panic::unwind_safe::AssertUnwindSafe<F> as core::ops::function::FnOnce<>>::call_once>>std::panicking::try>>std::panic::catch_unwind>>core::ops::function::FnOnce::call_once{{vtable-shim}}>>std::sys::unix::thread::Thread::new::thread_start>>start_thread"
1360 0.010767933, "getifaddrs_internal"
1355 0, "getifaddrs_internal>>event_base_loop>>base::MessagePumpLibevent::Run>>MessageLoop::Run>>base::Thread::ThreadMain>>ThreadFunc>>set_alt_signal_stack_and_start>>start_thread"
1029 0.010767933, "getifaddrs_internal>>event_base_loop>>base::MessagePumpLibevent::Run>>MessageLoop::Run>>base::Thread::ThreadMain>>ThreadFunc>>set_alt_signal_stack_and_start>>start_thread"
1029 -0.001543896, "getifaddrs_internal>>event_base_loop>>base::MessagePumpLibevent::Run>>MessageLoop::Run>>base::Thread::ThreadMain>>ThreadFunc>>set_alt_signal_stack_and_start>>start_thread"
1029 -0.001601887, "getifaddrs_internal>>event_base_loop>>base::MessagePumpLibevent::Run>>MessageLoop::Run>>base::Thread::ThreadMain>>ThreadFunc>>set_alt_signal_stack_and_start>>start_thread"
812 0, "__lll_lock_wait>>IPC::Channel::Send>>mozilla::ipc::NodeChannel::SendMessage>>mozilla::ipc::NodeController::ForwardEvent>>mojo::core::ports::Node::SendUserMessageInternal>>mojo::core::ports::Node::SendUserMessage>>mozilla::ipc::NodeController::SendUserMessage>>mozilla::ipc::PortLink::SendMessage>>mozilla::ipc::MessageChannel::Send>>mozilla::ipc::IPProtocol::ChannelSend>>mozilla::dom::PVsyncParent::SendNotify>>mozilla::dom::VsyncParent::DispatchVsyncEvent>>mozilla::detail::RunnableMethodImpl<mozilla::dom::VsyncParent*, void (mozilla::dom::VsyncParent::*)(mozilla::VsyncEvent const&), true, (mozilla::RunnableKind)0, mozilla::VsyncEvent>>Run>>nsThread::ProcessNextEvent>>NS_ProcessNextEvent>>mozilla::ipc::MessagePumpForNonMainThreads::Run>>MessageLoop::Run>>nsThread::ThreadFunc>>_pt_root>>set_alt_signal_stack_and_start>>start_thread"
752 -0.002455997, "putspent>>crossbeam_channel::channel::Receiver<T>::recv>>std::sys_common::backtrace::_rust_begin_short_backtrace>><core::panic::unwind_safe::AssertUnwindSafe<F> as core::ops::function::FnOnce<>>::call_once>>std::panicking::try>>std::panic::catch_unwind>>core::ops::function::FnOnce::call_once{{vtable-shim}}>>std::sys::unix::thread::Thread::new::thread_start>>start_thread"
573 0.010767933, "getifaddrs_internal>>[unknown] ([unknown])"
460 -0.002275313, "putspent>>crossbeam_channel::channel::Receiver<T>::recv>>webrender::render_backend::RenderBackend::run>>std::sys_common::backtrace::_rust_begin_short_backtrace>><core::panic::unwind_safe::AssertUnwindSafe<F> as core::ops::function::FnOnce<>>::call_once>>std::panicking::try>>std::panic::catch_unwind>>core::ops::function::FnOnce::call_once{{vtable-shim}}>>std::sys::unix::thread::Thread::new::thread_start>>start_thread"
376 -0.0024667, "putspent>>crossbeam_channel::channel::Receiver<T>::recv>>std::sys_common::backtrace::_rust_begin_short_backtrace>><core::panic::unwind_safe::AssertUnwindSafe<F> as core::ops::function::FnOnce<>>::call_once>>std::panicking::try>>std::panic::catch_unwind>>core::ops::function::FnOnce::call_once{{vtable-shim}}>>std::sys::unix::thread::Thread::new::thread_start>>start_thread"
```

# Conclusion

---

- ❑ Our study demonstrated the effectiveness of statistical debugging techniques in identifying performance issues.
- ❑ The results have significant implications for improving system call performance and can be applied to other systems.

# Future Work & Questions

---

- ❑ Future work includes refining these techniques, advancing sequence pattern mining and exploring broader applications.
- ❑ We are looking forward to discussing and receiving use-cases from our industrial partners.
- ❑ Special thanks to CINEA for their sponsoring of this project.
- ❑ Any questions or feedback?

Mohammed Adib Khan  
ak19qp@brocku.ca

Dr. Morteza Noferesti  
mnoferesti@brocku.ca

Dr. Naser Ezzati-Jivan  
nezzati@brocku.ca