

Titre: Title:	Hardware-assisted instruction profiling and latency detection
Auteurs: Authors:	Suchakrapani Datt Sharma et Michel R. Dagenais
Date:	2016
Type:	Article de revue / Journal article
Référence: Citation:	Sharma, S. D. & Dagenais, M. R. (2016). Hardware-assisted instruction profiling and latency detection. <i>The Journal of Engineering</i> , 2016(10), p. 367-376. doi: 10.1049/joe.2016.0127



Document en libre accès dans PolyPublie

Open Access document in PolyPublie

URL de PolyPublie: PolyPublie URL:	https://publications.polymtl.ca/3068/
Version:	Version officielle de l'éditeur / Published version Révisé par les pairs / Refereed
Conditions d'utilisation: Terms of Use:	CC BY



Document publié chez l'éditeur officiel

Document issued by the official publisher

Titre de la revue: Journal Title:	The Journal of Engineering
Maison d'édition: Publisher:	IET
URL officiel: Official URL:	https://doi.org/10.1049/joe.2016.0127
Mention légale: Legal notice:	

**Ce fichier a été téléchargé à partir de PolyPublie,
le dépôt institutionnel de Polytechnique Montréal**

This file has been downloaded from PolyPublie, the
institutional repository of Polytechnique Montréal

<http://publications.polymtl.ca>

Hardware-assisted instruction profiling and latency detection

Suchakrapani Datt Sharma, Michel Dagenais

Department of Computer and Software Engineering, École Polytechnique de Montréal, Montréal, Québec, Canada
E-mail: suchakrapani.sharma@polymtl.ca

Published in *The Journal of Engineering*; Received on 27th April 2016; Accepted on 24th August 2016

Abstract: Debugging and profiling tools can alter the execution flow or timing, can induce heisenbugs and are thus marginally useful for debugging time critical systems. Software tracing, however advanced it may be, depends on consuming precious computing resources. In this study, the authors analyse state-of-the-art hardware-tracing support, as provided in modern Intel processors and propose a new technique which uses the processor hardware for tracing without any code instrumentation or tracepoints. They demonstrate the utility of their approach with contributions in three areas - syscall latency profiling, instruction profiling and software-tracer impact detection. They present improvements in performance and the granularity of data gathered with hardware-assisted approach, as compared with traditional software only tracing and profiling. The performance impact on the target system – measured as time overhead – is on average 2–3%, with the worst case being 22%. They also define a way to measure and quantify the time resolution provided by hardware tracers for trace events, and observe the effect of fine-tuning hardware tracing for optimum utilisation. As compared with other in-kernel tracers, they observed that hardware-based tracing has a much reduced overhead, while achieving greater precision. Moreover, the other tracing techniques are ineffective in certain tracing scenarios.

1 Introduction

Modern systems are becoming increasingly complex to debug and diagnose. One of the main factors is the increasing complexity and real-time constraints which limit the use of traditional debugging approaches in such scenarios. Shorter task deadlines mean that the faithful reproduction of code execution can be very challenging. It has been estimated that developers spend around 50–75% of their time debugging applications at a considerable monetary cost [1]. In many scenarios, *heisenbugs* [2] become nearly impossible to detect. Long-running systems can have bugs that display actual consequences much later than expected, either due to tasks being scheduled out or hardware interrupts causing delays. Important parameters that need to be analysed while doing a root cause analysis for a problem include the identification of costly instructions during execution, the detection of failures in embedded communication protocols and the analysis of instruction profiles that give an accurate representation of which instructions consume the most central processing unit (CPU) time. Such latent issues can only be recorded faithfully using tracing techniques. Along with accurate profiling, tracing provides a much needed respite to developers for performance analysis in such scenarios.

We focus in this paper on two important common issues in current systems: the efficient detection/tracking of hardware latency and the accurate profiling of syscalls and instructions, with an ability to detect programme control flow more accurately than with current software approaches. We discuss our new analysis approach, which utilises conditional hardware tracing in conjunction with traditional software tracing to accurately profile latency causes. The trace can be decoded offline to retrieve the accurate control flow, data even at instruction granularity, without any external influence on the control flow. As software tracing can induce changes in the control flow, with our system we can further detect the cause of latency induced by the software tracers themselves, on the software under observation, to nanosecond range accuracy.

Pure software profiling and tracing tools consume the already constrained and much needed resources on production systems. Over the years, hardware tracing has emerged as a powerful technique for tracing, as it gives a detailed and accurate view of the system with almost zero overhead. The IEEE Nexus 5001 standard [3] defines four classes of tracing and debugging approaches for embedded systems. Class 1 deals with basic debugging operations

such as setting breakpoints, stepping instructions and analysing registers – often directly on target devices connected to hosts through a joint test action group (JTAG) port. In addition to this, Class 2 supports capturing and transporting programme control-flow traces externally to host devices. Class 3 adds data-flow traces support, in addition to control-flow tracing and Class 4 allows emulated memory and I/O access through external ports. Hardware-tracing modules for recent microprocessors (Class 2–Class 4) can either utilise (i) on-chip buffers for tracing, recording trace data from individual CPUs on the system-on-chip, and send it for internal processing or storage or (ii) off-chip trace buffers that allow trace data to flow from on-chip internal buffers to external devices, with specialised industry standard JTAG ports and to development host machines, through high-performance hardware-trace probes [4, 5]. These hardware-trace probes contain dedicated trace buffers (as large as 4 GB) and can handle high-speed trace data. As we observed in our performance tests (Section 4), the former approach can incur overhead in the range of 0.83–22.9%, mainly due to strain on memory accesses. We noted that trace streams can generate data in the range of hundreds to thousands of MB/s (depending on trace filters, trace packets and packet generation frequency). Thus, there is a trade-off in choosing either an external analysis device or on-chip buffer recording. The former gives a better control (less dependency on external hardware which is crucial for on-site debugging), but incurs a small overhead for the memory subsystems on the target device. The latter provides a very low overhead system, but requires external devices and special software (often proprietary) on development hosts. The generated trace data is compressed for later decoding with specialised debug/trace tools [6, 7] that run on host machines, as illustrated in Fig. 1.

Device memory is limited, thus there are multiple ways to save tracing data using either of the two approaches discussed above. Therefore, to achieve maximum performance, recent research deals with compressing the trace output during the decoding phase to save transfer bandwidth [1]. Earlier, part of the focus was on the unification of the traces, which is beneficial for Class 3 devices [8]. This provides a very detailed picture of the execution at almost no overhead on the target system.

In this paper, we mainly focus on on-chip local recorded traces, pertaining to Class 2 devices, owing to their low external hardware dependency and high availability in commonly used architectures




Fig. 1 Hardware-tracing overview

such as Intel x86-64. With our proposed approach, using hardware-trace assistance, we were able to trace and profile short sections of code. This would ensure that precious input/output (I/O) bandwidth is saved while maintaining sufficient granularity. We used Intel's new Processor Trace (PT) features and were able to gather accurate instruction profiling data such as syscall latency and instruction counts for interesting events such as abnormal latency. In profile mode, we can also selectively isolate sections of code in the operating system, for instance idling the CPU, taking spinlocks or executing floating point unit (FPU) bound instructions, to further fine-tune the systems under study.

The remainder of this paper is organised as follows. Section 2 gives a general overview of programme-flow tracing, its requirements and limitations of the current sampling systems to handle these. It also introduces the concept of hardware tracing to overcome such limitations. We discuss state-of-the-art techniques used in software tracing and finally concentrate on our research scope. In Section 3, we introduce our hardware and hardware-assisted software-tracing-based architecture. We then elaborate our three contributions that introduce an instruction and time-delta profiling technique to identify interrupt latencies, profiling syscall latency in short sections of code and identify causes of latency in software tracers. These contributions utilise our hardware-assisted approach. In Section 4, as our final contribution, we start with a detailed experiment, measuring overhead and trace size, for Intel's PT hardware-tracing infrastructure as it forms the core of our hardware-based approach. We have also proposed a new metric to define the granularity of the temporal and spatial resolution in a trace. We then show how the hardware-based trace profiler can help visualise anomalies through histograms.

2 Background

In an ideal environment, developers would want to know as much as possible about the flow of their programmes. There are three important artefacts that can be gathered during a programme execution – flow of instructions during programme, their classification and deduction of the programme-flow with timing information. Static analysis of binaries, to understand how the programme runs, allows the developers to visually analyse how the compiler generates instructions, estimate how the instructions may execute, and can be used further for code coverage [9, 10]. Such information is also vital for debuggers to generate and aid in the breakpoint debugging approach. Recently, the focus on pure static code analysis tools has been mostly in the security domain, for analysing malicious or injected code in binaries [11, 12] or optimising compilers based on the analysis of generated code [13]. However, the actual execution profiles can differ from what static tools can anticipate, due to the complexities of newer computer architectures in terms of pipelines, instruction prefetching, branch prediction and unforeseen run-time behaviour such as hardware interrupts. Therefore, to understand the effect of individual instructions or function blocks, the instructions executed can be profiled at run-time. The use of counting instructions for blocks of code, at programme execution time, has been proposed and explored in-depth before [14]. Therefore, the instruction sequence and profile can be recorded and then replayed later on. However, some of these earlier approaches dealt with inserting instrumentation code, to profile instructions and other interesting events. Sampling-based techniques, developed earlier such as digital continuous profiling infrastructure [15, 16], have also been discussed extensively before, where authors demonstrated the use of hardware counters provided by the processor for profiling instructions. Merten *et al.* [17] have earlier proposed the use of a branch trace buffer and their hardware table extension for profiling branches. Custom hardware-based path profiling has been discussed by Vaswani *et al.* [18], where they observe that low overhead with hardware-assisted path profiling can be achieved. Recent advances, especially in the Linux kernel, discuss how profiling tools such as Perf can be used to generate execution profiles, based on data collected from special hardware counters, hardware blocks that record branches or pure software controlled sampling [19, 20].

2.1 Programme-flow tracing

Recording instruction flow or branches in a programme can provide information about how a programme actually executes in comparison with how the expected execution. The comparison of an anomalous programme-flow trace with that of a previous one can let the developer know what was the effect of changes on the system. It can also be used to track regressions during new feature additions. At lower levels such as instructions flow, bugs that occur in long-running real-time systems can now be detected with more confidence, as the complete execution details are available. With recent hardware support from modern processors, this has become easier than ever. We discuss details about such hardware support further in Section 2.2. Larus *et al.* discussed quite early about using code instrumentation to inject tracing code in function blocks or control-flow edges to track instructions or deduce the frequency of their execution [14, 21]. They observed overhead of 0.2–5%, without taking into consideration the effect of the extra overhead of disk writes (which they observed as 24–57% in those days). Other more powerful tools, which effectively perform execution profiling or control-flow tracing, can be built using similar binary modifying frameworks such as Valgrind [22]. Even though this framework is more data-flow tracing oriented [23], some very insightful control-flow tools have been developed such as Callgrind and Kcachegrind [24]. Programme-flow tracing can either encompass a very low-level all-instruction trace generation

scheme or a more lightweight branch-only control-flow trace scheme.

Instruction tracing: Tracing each and every instruction to deduce the programme-flow can be quite expensive if instrumentation is required. Hence, architectures such as advanced RISC machines (ARM) and PowerPC provide hardware support for such mechanisms in the form of NSTrace (PowerPC), EmbeddedICE, embedded trace macrocell (ETM), programme trace macrocell (now part of ARM CoreSight) and microprocessor without interlocked pipeline stages PDTrace [25, 26]. The basic idea is to snoop the bus activity at a very low-level, record such data and then reconstruct the flow offline from the bus data. External hardware is usually connected as bus data *sink* and special software can then use architecture level simulators to decode the data. The benefit of a complete instruction flow trace is that there is highly detailed information about each and every instruction for accurate profiles, in-depth view of memory access patterns and, with the support of time-stamped data, a very accurate overall tracer as well. However, the amount of data generated is too high if external devices are not used to sink the data. Indeed, memory buses are usually kept busy with their normal load, and an attempt to store the tracing data locally incurs bus saturation and additional overhead. An approach to reduce such bandwidth and yet keep at least the programme-flow information correct is to use branch-only traces.

Branch tracing: The issue of memory-related overhead for hardware programme/data-flow traces has been observed earlier as well [14, 21]. Even though hardware can generate per-instruction trace data at zero execution overhead, such an additional data-flow may impact the memory subsystem. Hence, just choosing the instructions that cause a programme to change its flow greatly reduces the impact. Such control-flow instructions (such as direct/indirect jumps, calls, exceptions etc.) can indeed be enough to reconstruct the programme-flow. Dedicated hardware blocks in the Intel architecture such as last branch record (LBR), branch trace store (BTS) [27] and more recently Intel PT chooses to only record branches in the currently executing code on the CPU cores. By following the branches, it is quite easy to generate the instruction flow with the help of additional offline binary disassembly. For example, for each branch instruction encountered in the flow, a record for the branch taken/not-taken and its target can be recorded externally. This is then matched with the debug information from the binary to reconstruct how the programme was flowing. We detail and discuss the branch-tracing approach, as well as instruction tracing, in Section 2.2, where we show a state-of-the-art branch tracing approach using Intel PT as an example.

2.2 Hardware tracing

As discussed previously, the complete instruction and branch tracing is supported by dedicated hardware in modern multi-core processors. They provide external hardware recorders and tracing devices access to the processor data and address buses. ARM's early implementation of EmbeddedICE (in-circuit emulator) was an example of this approach. Eventually, processor chip vendors formally introduced dedicated and more advanced hardware-tracing modules such as CoreSight, Intel BTS and Intel PT. In a typical setup such as shown in Fig. 1, trace data generated from the trace hardware on the chip can be funnelled to either the internal buffer for storage or observed externally through an external hardware buffer/interface to the host development environment, for more visibility. In both cases, the underlying techniques are the same, but performance varies according to the need of the user and the hardware implementation itself.

2.2.1 Tracing primitives: Since an important part of our research deals with programme-flow tracing, we discuss how hardware-tracing blocks can be used to implement it. The basic idea is to record the control-flow instructions along with some timing

information (if needed) during the execution of the programme. Different architectures have different approaches for deciding on the optimum buffer size, trace compression techniques and additional meta-data such as timing information, target and source instruction pointers (IPs) etc. We explain such techniques along with an overview of the tracing process in this section. A programme-flow trace can be broadly broken down into following elements.

Trace configuration: Most of the hardware-trace modules can be fine-tuned by writing data to certain control registers such as model specific registers (MSRs) in Intel or the CoreSight ETM/ETB configuration registers for ARM. For example, writing specific bits in MSRs can control how big a trace buffer will be or how fine-grained or accurate the timing data will be generated. An optimum configuration leads to better trace output – the effect of which is discussed later in this paper.

Trace packets: A hardware-trace-enabled execution generates all the hardware-trace data in a compressed form for eventual decoding. This can consist of different distinguishable elements called trace packets. For example, in the context of Intel PT, these hardware-trace packets can contain information such as paging (changed CR3 value), time-stamps, core-to-bus clock ratio and taken-not-taken (tracking conditional branch directions), record target IP of branch, exceptions, interrupts, source IP for asynchronous events (exceptions, interrupts). The amount or type of packets enabled and their frequency of occurrence directly affect the trace size. In the control-flow trace context, the most important packets that are typically common to different architectural specifications of trace hardware are:

Taken-not-taken: For each conditional direct branch instruction encountered (such as jump on zero and jump on equal), the trace hardware can decode if that specific branch was taken or not. This is illustrated with Intel PT's trace output as an example in Fig. 2. We can observe that Intel PT efficiently utilises 1 bit per branch instruction to encode it as a taken or not-taken branch.

The earlier implementations such as Intel BTS used 24 bits per branch, which caused an overhead between 20 and 100% as the CPU enters the special debug mode, causing a 20–30 times slowdown [28, 29]:

Target IP: Indirect unconditional branches (such as register indirect jumps) depend on register or memory contents; they require more bits to encode the destination IP of the indirect branch. This can also be the case when the processor encounters interrupts, exceptions or far branches. Some implementations such as Intel PT provide other packets for updating control flow such as flow update packet, which provide source IP for asynchronous events such as interrupts and exceptions. In other scenarios, the binary analysis can usually be used to deduce the source IP.

Timing: Apart from deducing the programme-flow, packets can be timed as well. However, time-stamping each and every instruction can be expensive in terms of trace size as well as extra overhead incurred. ARM CoreSight provides support for accurate time-stamp per-instruction. However, the use-case is mainly aimed at usage of an external high-speed buffer and interface hardware through a JTAG port. For on-device tracing such as Intel PT, the packet size can be kept small by controlling the frequency of time-stamps being generated and the type of time-stamps. For example, a timing




Fig. 2 Odd-even test generates corresponding taken-not-taken packets

packet can either be the lower 7 bytes of the time-stamp counter (TSC) value as an infrequently recorded TSC packet or can be just a more frequent 8 bit mini-time-stamp counter (MTC) packet occurring in between two TSC packets. MTC packets record incremental updates of CoreCrystalClockValue and can be used to increase the timing precision with fewer bits utilised. Trace timing in some implementations can further be improved by a cycle-accurate mode, in which the hardware keeps a record of cycle counts between normal packets.

In the next section, we discuss how we can leverage hardware-tracing techniques and utilise it for efficient and more accurate profiling and tracing.

3 Trace methodology

To get useful instruction profiling and tracing data for use-cases such as accurate detection of interrupt latency, we propose a framework that utilises an hardware-assisted software-tracing approach. The major focus of our work is on post-mortem analysis of production systems. Hence, the underlying technologies used aim at recording raw trace or programme-flow data at run-time, and eventually perform an offline merge and analysis, to get in-depth information about abnormal latency causes or generate instruction execution profiles. The data generated in hardware tracing can reach a range of hundreds of megabytes per second. Various approaches have been taken to reduce this overhead. Apart from careful configuration of the trace hardware, various methods such as varying the length of taken-not-taken (TNT) packets (short/long), IP compression and indirect transfer return compression [30] are employed to control precisely the trace size, with the aim of reducing memory bus bandwidth usage. Previous work often focused on trace compression and even better development of tracing blocks itself [1]. In contrast, we chose to leverage the latest state-of-the-art hardware such as Intel PT and carefully isolate interesting sections of the executed code to generate short hardware traces. These short traces can be eventually tied to the corresponding software-trace data to generate a more in-depth view of the system at low cost. This can also be used to generate instruction execution profiles in those code sections for detecting and pinpointing anomalous sections of the programme, right down to the executed instruction. This gives a unique and better approach as compared with other techniques of sample-based profiling (such as Perf) or simulation/translation-based profiling (such as Valgrind) mainly due to the fact that there is no information loss, as the inferences are based on the real instruction flow in a programme, and the overhead of simulated programme execution is completely removed. Choosing only specific sections of code, to trace and profile with hardware, also means that we do not require external trace hardware and can rely on internal trace buffers for post-mortem analysis. In this section, we first show the design of our framework itself and demonstrate how we can use Intel PT hardware-assisted software tracing. We also explain our three main contributions, detailing how we could profile interrupts/syscall latency and evaluate the impact of software tracers themselves. We start with some background on Intel PT, and then explain the architecture of our technique.

3.1 Intel PT

Intel’s MSR-based LBR and the BTS approach for branch tracing have been widely explored before [28, 29]. Eventually, the benefits of the hardware-tracing approach advanced the branch-tracing framework further, in the form of Intel PT. Branch trace data with PT can now be efficiently encoded and eventually decoded offline. The basic idea, as shown in Fig. 2, is to save the branching information during programme execution, encode and save it. Later on, the trace data along with run-time information such as process maps, debug-info and binary disassembly, we can fill in the gaps

between the branches and form a complete execution flow of the application. During the decoding of the compressed recorded branch data, whenever a conditional or indirect branch is encountered, the recorded trace is browsed through to find the branch target. This can be merged with debug symbols and static analysis of the binary code to get the intermediary instructions executed between the branches. Therefore, with Intel PT’s approach, we do not need to exclusively store each and every instruction executed, but just track branches – allowing on-device debugging and less complex implementation of hardware and debugging software. Apart from that, with the simple MSR-based configuration of the hardware, we have the ability to set hardware trace start and stop filters based on an IP range to allow a more concise and efficient record of trace at run-time. The decoder has been open sourced by Intel for a rapid adoption in other tools such as Perf [31]. We incorporated PT in our framework, owing to its low overhead and versatility as presented later in Section 4.2 where we discuss its performance and overhead.

3.2 Architecture

We developed a framework based on the PT library, for decoding the hardware trace, and a reference PT driver implementation provided by Intel to enable, disable and fine-tune trace generation [32, 33]. An overview of our hardware-assisted trace/profile system is shown in Fig. 3. The control block is a collection of scripts to control the PT hardware in CPUs through the simple-PT module. The control scripts can be used for configuring the PT hardware for filtering, time-stamp resolution and frequency control. It can enable/disable traces manually for a given time period. This is achieved by setting the **TRACE_EN** bit in the **MSR_IA32_RTIT_CTL** control register that activates or

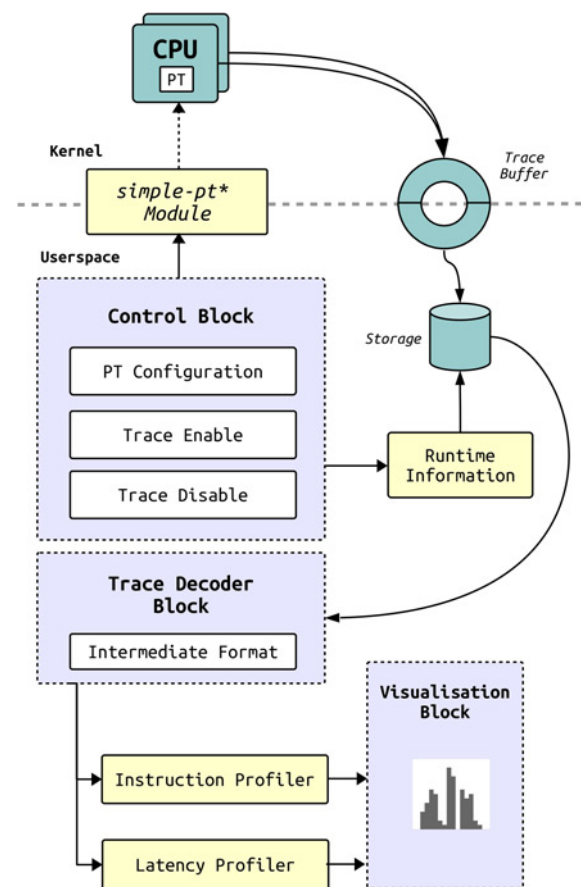


Fig. 3 Architecture of our proposed hardware-assisted trace/profile framework. Simple PT [32] is used for trace hardware control

deactivates the trace generation. We use the control scripts for generating raw instruction and latency profiling data.

Any PT data generated is sent to a ring-buffer, the content of which is dumped to the disk for offline analysis. Along with the trace data, run-time information such as process maps, CPU information and features are saved as *sideband* data. This is essential for trace reconstruction. The trace decoder reconstructs the control flow based on run-time information, debug information from binary and per-CPU PT data. The decoding process involves reading the PT binary byte by byte. Much such as assembly opcodes, packets are identified by an arrangement of bytes. We go through individual incoming bytes from the trace buffer to identify the packet and its contents. This is merged with the disassembled binary information to give the packets an execution context. This data is converted to our intermediate format that is consumed by the instruction profiler and latency profiler modules which can generate visualisations. Our intermediate format consists of a stream of instructions with markers identified by function names. This is converted to visualisation data by transforming the markers to a callstack with instruction or time-delta. More about this is explained in subsequent sections. We now elaborate on our three contributions that cover instruction/syscall latency profiling and the performance impact of software-tracing systems on modern production systems.

3.3 Delta profiling

As seen in Fig. 3, the raw PT data from the processor can be reconstructed to generate a programme control flow. It is, however, important to analyse this huge information in a meaningful manner. Therefore, we present an algorithm to sieve through the data and generate instruction execution profiles based on patterns of occurrence of instructions in the control flow.

These profiles can be used to represent the histograms based on a time-delta or an instruction count delta. This can work for single instructions to observe histograms during a simple execution, as well by just counting the occurrence of a set of instructions. This approach is significantly different from sample-based profiling, as it is based on true instruction flow and can pinpoint errors at finer granularity in short executions. As interrupts are quite significant in embedded production systems, we choose profile instructions that are responsible for disabling and enabling interrupts in the Linux kernel. Thus, we generate two histograms that represent time-delta and instruction count delta of intervals between interrupt enabling and disabling instructions. We observed that interrupt disabling and enabling in Linux on an x86 machine is not just dependent on two instructions, **sti** and **cli**, but also on a pattern of instructions that use **pushf** and **popf** to push and pop the entire **EFLAGS** register, thus clearing or setting the interrupt flag in the process. Thus, to effectively profile the interrupt cycle, we identified the instruction patterns during decoding and grouped and identified them as **superSTI** and **superCLI** instructions. For incoming coded hardware-trace streams, we devised an algorithm as shown in listing Algorithm 1 (see Fig. 4) that is able to generate these profiles. For example, when we apply this during the trace decode time, we can obtain the time taken between two consecutive interrupts enable and disable in the execution or the number of instructions executed between them. These target super-instruction pairs (S_i), which are actually a pseudo-marker in the instruction stream based on pattern matching, can be given as input to the profiler. On the basis of the mode, it can either begin instruction counting or time-stamp generation and stores it in a database. We then iterate over the database and generate the required visualisations.

We can extend this technique of identifying patterns in the code to record more interesting scenarios. For example, in the Linux kernel, CPU idling through the **cpu_relax()** function generates a series of repeating **nop** instructions. Similarly, the *crypto* subsystem in the Linux kernel aggressively uses the less-recommended

Algorithm 1: Hardware Delta Profiling

Data: $\{\mathbf{D} : t_p \dots t_q\}$, where \mathbf{D} is the coded stream for time T_{q-p} , Target Instruction Set \mathbf{I} (either individual instruction pair I_t or super-instruction pair S_i) and **Mode** ($T\Delta$ or $IC\Delta$)

Result: Time- Δ histogram, Instruction Count- Δ histogram, Instruction Count histogram of \mathbf{I}

```

begin
   $iC \leftarrow 0$ 
  resetFlags()
  for  $i \in D$  do
     $x = \text{decodeInstruction}(i)$ 
    incrementCount()
    if  $x = (I_t \text{ or } S_i)$  and Mode =  $T\Delta$  then
      if set( $F$ ) then
        unset( $F$ )
         $ts_{n+m} \leftarrow t_x$ 
         $\Delta t = ts_{n+m} - ts_n$ 
        addToDatabase( $DB, \Delta t$ )
      else
         $ts_n \leftarrow t_x$ 
        set( $F$ )
    if  $x = (I_t \text{ or } S_i)$  and Mode =  $IC\Delta$  then
      if set( $F$ ) then
        unset( $F$ )
         $IC\Delta = \text{getDelta}()$ 
        resetCounter( $iC$ )
        addToDatabase( $DB, IC\Delta$ )
      else
        startCounter( $iC$ )
        set( $F$ )
    count( $iC$ )
  generateHistogram( $DB$ )

```

Fig. 4 Hardware-delta profiling algorithm generates time-delta and instruction count-delta histogram from encoded trace stream

FPU. We were able to successfully identify such patterns in the system based purely on PT, without any active software tracer.

We present an implementation of the algorithm in Section 4.3 where we elaborate more on our delta profiling experiment. Decoded traces from Intel PT were chosen to demonstrate utility of this approach.

3.4 Syscall latency profiling

Syscalls affect the time accuracy of systems, especially in critical sections, as they form a major chunk of code executed from user-space. For example, filesystem syscalls such as **read()**, **open()**, **close()** etc. constitute 28.75% of code in critical sections of Firefox [34]. Profiling syscall counts for a given execution is easy and can be performed with simple profilers such as Perf or even through static or dynamic code analysis techniques based on **ptrace()**. However, to understand, the extra time incurred in the syscalls, we can get help from software tracers. As the software tracers themselves affect the execution of syscalls, an accurate understanding can only be achieved by an *external observer* which does not affect the execution flow. In such scenarios, hardware tracing is a perfect candidate for such an observer. We used hardware traces and devised a way to visualise syscall stacks in our proposed technique, after decoding, to compare them between multiple executions. This gave us a deep and accurate understanding of any extra time incurred in syscalls, right down to individual instructions. We devised a way such that, post-decoding, the trace data was converted to our visualisation path format that prepares a raw callstack. See Fig. 5.

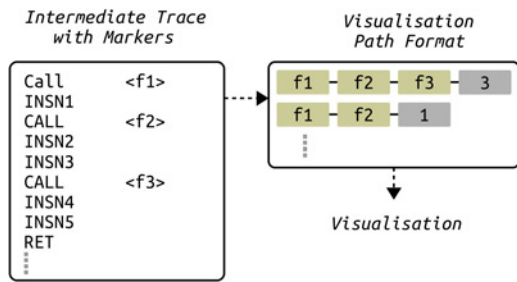


Fig. 5 Sample trace sequence converted to a section of visualisation path format

The function markers in the decoded data are converted to individual execution paths in the hierarchy. Each node in the path represents a function. To each path, we append the instructions that executed for traversing the path up to its tail. Each line in the new data contains a single complete path in the callstack. As an example, Fig. 6a illustrates the effect of an external tracer (LTTng) on the `mmap()` syscall with the help of a callstack. The callstack shown is for a short section of code from the library `mmap` call to the start of the syscall in the kernel. We see in this figure that the highlighted call-path reached the `lttng_event_write()` function from the second ring of the `entry_SYSCALL_64()` function in the kernel. The layers represent calls in a callstack, with the call depth going from the innermost to the outermost layers. Here, the path to the `lttng_event_write()` function took 9.3% of all instructions in the recorded callstack. As the code sections are short, and the data represented is hierarchical in nature, it is easy to visualise them on sunburst call-graphs [35, 36] for a clear visual comparison.

We can observe such a short callstack from another execution, in Fig. 6b, where LTTng tracing is disabled, and we note the absence of extra calls which were added as layers and peaks in Fig. 6a. The metrics in the sunburst graphs are calculated based on the number of instructions executed along a particular call-path based on the visualisation path format. The visualisation is interactive and its implementation is based on the D3 javascript library.

3.5 Software-tracer impact

With our hardware-assisted profile, we can observe how much extra time and instructions any external software tracer added to the normal execution of the syscall. This can also be used as a basis for analysing the overhead of known tracers on the test system itself. For example, we observed that the extra time taken in the syscall is due to the different paths the syscall has when tracing is enabled, as compared with when tracing is disabled. A lot of code in the kernel is untraceable such as C macros and blacklisted functions built with the `__attribute__((no_instrument_function))` attribute that do not allow tracers to trace them. This is usually a mandatory precaution taken in the kernel to avoid tracers going in sections of code that would cause deadlocks. However, our PT-based approach allowed us to get much finer details than other pure software tracers, as it acts as an *external observer* and can even record calls to those functions. We monitored how LTTng changes the flow of seven consecutive syscalls in a short section of traced code. As an example, for `mmap()` calls, we observed that with software tracing and recording enabled, a total of 917 additional instructions were added to the normal flow of the syscall, which took an extra 173 ns. For short tracing sections, an overhead of 579 ns was observed on average for `open()` syscalls, with 1366 extra instructions. We observed that the overhead also varies according to the trace payload for syscalls, as LTTng specific functions in the kernel modules copy and commit the data to trace events. Therefore, with the hardware-trace




Fig. 6 Effect of an external software tracer on the `mmap()` syscall, obtained from a near zero overhead hardware trace, is visible in a and b. The rings represent the callstack and are drawn based on instruction count per call a Extra layers in callstack as it took a longer syscall path b Shallower callstack as it took a shorter syscall path

assisted tool, we can get instruction and time accurate overhead of the tracer's impact on the system itself. Such detailed information about a software tracer's impact is not possible to obtain by conventional software tracers themselves. Hardware-assisted profiles allow to study the flow through those unreachable sections of code (assembly, non-traceable functions in the kernel) along with a higher granularity. We tested this with our PT-based approach. In Section 4.2, we further compare the overhead of Linux kernel's Ftrace tracer with Intel PT to see the impact of hardware traces as compared with current function tracing facility.

4 Experimentation and results

4.1 Test setup

The test machine has an Intel Skylake i5-6600K processor which supports Intel PT and runs a patched Linux kernel version 4.4-rc4 on a Fedora 23 operating system. To get minimum jitter in our tests, we disabled CPU auto-scaling and fixed the operating frequency to 3.9 GHz. This was done to just ensure that synthetic benchmarks were accurate and reproducible. This would not affect the result as the ratio of core crystal clock value to TSC frequency is considered during decoding time to make sure the effect of auto-scaling is accounted for in the time calculations. The system has 16 GB main memory and a 500 GB solid-state drive.

4.2 PT performance analysis

The most important requirement for a performance analysis framework is that it should have minimum impact on the test system itself. Therefore, before deciding on the trace hardware for our framework, we characterised PT's performance. The impact of a PT-based hardware-assisted tracer on the system has not been thoroughly characterised before. As it formed the basis of our other contributions; therefore, as a major part of our work, we developed a series of benchmarks to test how much overhead the tracing activity itself causes. We measured four aspects – the execution overhead in terms of extra time, the trace bandwidth, and the trace size and temporal resolution with varying time accuracy. We also compared the overhead of our PT-based approach with that of current default software tracers in Linux kernel.

4.2.1 Execution overhead: Similar to such synthetic tests done for measuring the Julia Language performance [37] against C and Javascript, we added more indirect branch intensive tests such as TailFact which causes multiple tail-calls for factorial computation and Fibonacci to illustrate conditional branches. We also tested an un-optimised and optimised Canny edge detector to check the effect of jump optimisations in image processing tasks. As conditional branches constitute most of the branch instructions, to get a precise measurement for a conditional branch, we tested a million runs of an empty loop (Epsilon) against a loop containing an un-optimised conditional branch (Omega) Therefore, the TailFact test gives us the upper limit of overhead for indirect branch instructions while the Omega test gives us the upper limit for conditional branches.

Observations: Our test results have been summarised in Table 1. We can see that excessive TIP packets generated due to tail-calls from the recursive factorial algorithm cause the maximum overhead of 22.9%, whereas the optimised random matrix multiplication (RandMatMul) overhead is 0.83%. The optimisation in the C version of RandMatMul is evident as it aggressively used vector instructions (Intel AVX and SSE) from the BLAS library [38] during a DGEMM, thus generating very few TIP or TNT packets, as compared with the un-optimised loop-based multiplication in Javascript, which through the V8's JIT got translated to conditional branches. This explains the difference, as seen in the table, where the overhead for V8 is 11.08%. Same is the case with RandMatStat, which also generated more TIP packets thus pushing the overhead to 20%. To observe the TNT packet overhead, the Omega test generated pure TNT packets. As this includes one million extra conditional jump overhead from the test loop for both Epsilon and Omega, we can normalise the overhead and observe it to be 8.68%.

Table 1 Execution overhead and trace bandwidth of Intel PT under various workloads. The TailFact and Omega tests define the two upper limits

Benchmark	Bandwidth, MBps	Time overhead	
		C, %	V8, %
TailFact	2200	22.91	–
ParseInt	1420	9.65	10.36
Fib	1315	5.86	5.80
RandMatStat	340	2.58	20.00
CannyNoOptimise	303	2.55	–
PiSum	339	2.47	6.20
CannyOptimise	294	2.34	–
Sort	497	1.05	6.06
RandMatMul	186	0.83	11.08
Omega	205	11.78(8.68)	–
Epsilon	217	3.10(0.0)	–

4.2.2 Trace bandwidth: The direct correlation between the trace size, packet frequency and hence the trace bandwidth is quite evident. To quantify it, we record the trace data generated per time unit and quantify the trace bandwidth for our micro-benchmarks in Table 1. To calculate the bandwidth, we record the size of raw trace data generated and the time taken for execution of the individual benchmarks.

Observations: We see that the trace bandwidth is quite high for workloads with high-frequency TIP packets such as TailFact. Larger TIP packets increase the bandwidth and cause a considerable load on the memory bus. Overall, for moderate workloads, the median bandwidth lies between 200 and 400 MBps.

4.2.3 Trace size: Apart from the inherent character of the applications (more or less branches) that affect the trace size and timing overhead, Intel PT provides various other mechanisms to fine-tune the trace size. The basic premise is that the generation and frequency of other packets such as timing and cycle count information can be configured before tracing begins. To test the effect of varying such helper packets, we ran the PiSum micro-benchmark from our overhead experiments, which mimics a more common workload with userspace-only hardware-trace mode. We first started with varying the generation of synchronisation packets called PSB, whereas the cycle-accurate mode (CYC packets) and the MTC packets were disabled. The PSB frequency can be controlled by varying how many bytes are to be generated between subsequent packets. Thus, for a higher number of bytes, those packets will be less frequent. As PSB packets are accompanied with a TSC packet, this also means that the granularity of timing varies. The same was repeated with MTC packets while the PSB packet generation was kept constant and the CYC mode was disabled. Similar tests were done with CYC packets, where PSB packet generation was kept constant and MTC packet generation was disabled. Figs. 7–9 show the effect of varying the frequency of packets on the generated trace data size.

Observations: We observe in Figs. 7–9 that, as expected, when the time period (indicated by number of cycles or number of bytes in-between) for CYC, MTC or PSB packets is increased, the trace size decreases. However, it moves toward saturation, as opposed to a linear decrease. The reason we observed is that, for a given trace duration, there is a fixed number of packets that are always generated from the branches in the test application. This sets a minimum threshold. Furthermore, in Fig. 7, the trace size did not increase further for time periods $<2^6$ cycles, because the maximum number of CYC packets that can be generated for our synthetic workload was reached at 2^6 cycles. The trace data size




Fig. 7 Trace size and resolution while varying valid CPU cycles between two subsequent CYC packets. Lower TRF value is better




Fig. 8 Trace size and resolution while varying valid CPU cycles between two subsequent MTC packets. Lower TRF value is better

can however further increase for other workloads with higher frequency of CYC packets, when kernel tracing is enabled. In our tests, we found that the lower and upper bounds of trace data size, based on lowest and highest possible frequencies of all packets combined, are, respectively, 819 and 3829 KB.

4.2.4 Temporal resolution: In addition to the effect of different packet frequencies on the trace data size, it is also important to observe how much timing granularity we lose or gain. This can help the user decide what would be the trade-off between size, timing granularity of the trace and the tracing overhead, to better judge the tracer’s impact on the target software. We therefore define a temporal resolution factor (TRF). For a given known section of code, with equidistant branches

$$TRF = \frac{N_f}{\max(P) - \min(P)} \times [p - \min(P)]$$

where

$$p = \left(\frac{\Delta I_c \text{Sort}[n-1]}{2} \right) \times \left(\frac{\Delta T \text{Sort}[n-1]}{2} \right)$$




Fig. 9 Trace size and resolution while varying valid bytes of data between two subsequent PSB packets

and P is the set of all p that represents a median factor in the observed sets. Here, ΔT and ΔI_c are the time and instruction deltas between subsequent decoded branches and n is the length of the total decoded branches. The formula calculates the median of the sorted datasets of ΔT and ΔI_c and normalises them with a factor of N_f for an accurate representation. The value of n , however, varies according to the frequency of packets. Hence, with the packets we estimate the averages based on the maximum we can obtain. This experiment was coupled with the trace size experiment above so that, for the same executions, we could observe our temporal resolution as a function of the trace size as well. The choice of equidistant branches is intentional for a more controlled and repeatable micro-benchmark. The results are presented in Figs. 7–9 with TRF on the second Y-axis. TRF varies between 0 and 100. Lower TRF values represent better resolutions.

Observations: It is interesting to see a clear trend that the data size is inversely proportional to TRF. Hence, the larger the trace size, the better is the temporal resolution in the trace. Another important observation is the sudden increase in resolution when CYC packets are introduced. We observed that the highest resolution we obtained for our tests was 14 ns and 78 instructions between two consecutive events in the trace ($TRF = 4.0 \times 10^{-9}$). This compares with the lowest resolution of 910.5 μ s and 2.5 million instructions ($TRF = 100$) between two events with no CYC and far apart PSB packets. The reason for such a huge variation is that the introduction of the cycle-accurate mode generates CYC packets before all CYC-eligible packets (such as TNT and TIP). The CYC packets contain the number of core clock cycles since the last CYC packet received, which is further used to improve the time resolution. With a large number of CYC-eligible packets being generated in quick succession, the trace size as well as the resolution increases drastically, as compared with MTC and PSB packets. For CYC observations in Fig. 7, the resolution for $<2^6$ cycles is not shown, as it covers the whole execution in our workload for which we are interested. Thus, we always get a constant maximal number of CYC packets and the TRF value saturates. Therefore, we only included valid cycles $>2^6$. This can however vary for real life use-cases such as kernel tracing where the branches are not equally spaced and the code section is not linear. However, the TRF values obtained can be a sufficient indicator of upper and lower bounds of resolution.

4.2.5 Ftrace and PT analysis: We also compared the hardware control-flow tracing with the closest current software solutions in the Linux kernel. An obvious contender in-kernel control-flow tracing for our Intel PT-based framework is Ftrace [39]. Both can be used to get the execution flow of the kernel for a given workload – with our approach providing a much more detailed view than Ftrace. We therefore used the Sysbench synthetic benchmarks to gauge the overhead of both approaches for disk I/O and memory and CPU intensive tests. We configured Ftrace in a mode as close as possible to Intel PT by outputting raw binary information to a trace buffer. We also set the per-CPU trace buffer size to 4 MB. Our findings are presented in Table 2. We can see that, as compared with PT, Ftrace was 63% slower for a Sysbench File I/O workload with random reads and writes. This generates numerous kernel events for which Ftrace had to obtain time-stamps at each function

Table 2 Comparison of Intel PT and Ftrace overheads for synthetic loads

Benchmark	Baseline	With PT	With Ftrace	Overhead	
				PT, %	Ftrace, %
File I/O, MBps	32.32	31.99	19.76	1.00	63.56
Memory, MBps	5358.25	5355.59	4698.52	0.00	14.04
CPU, s	19.001	19.007	19.121	0.00	0.6

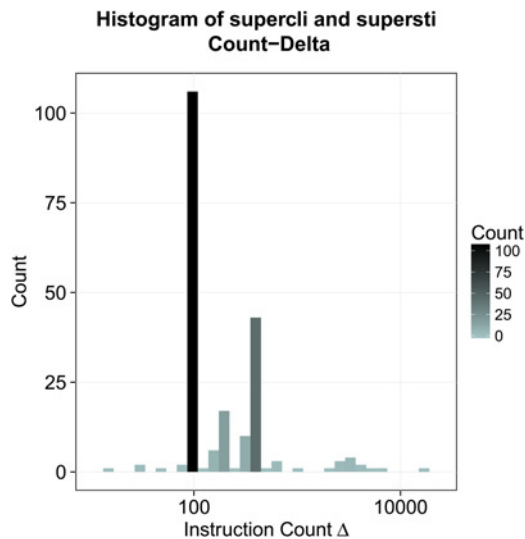


Fig. 10 Histogram of instruction count delta for *superSTI* and *superCLI* instructions generated using delta profiling algorithm

entry. In the PT case, the time-stamps are an inherent part of the trace data generated in parallel through trace hardware. This explains the huge difference in overhead. A similar difference is observed in the memory benchmark as well. In the case of the CPU benchmark, the work was userspace bound and hence the trace generated was smaller – thus a non-statistically significant overhead in PT and 0.6% overhead in Ftrace.

4.3 Delta profiling instructions

In these sets of experiments, we show how our Algorithm 1 (Fig. 4) is able to generate histograms for instruction and time-delta for *superSTI* and *superCLI* instructions groups. To quantify how much time is spent in a short section where interrupts are disabled, we created a synthetic test module in the kernel that disables and enables interrupts as our input. We control the module using `ioctl()` to cycle the interrupts. We pin the userspace counterpart of our test module on CPU0 and analyse the hardware trace for CPU0. Our algorithm is implemented during the trace decoding phase to generate the instruction delta and time-delta in an intermediate format which we then plot as histograms. This helps to

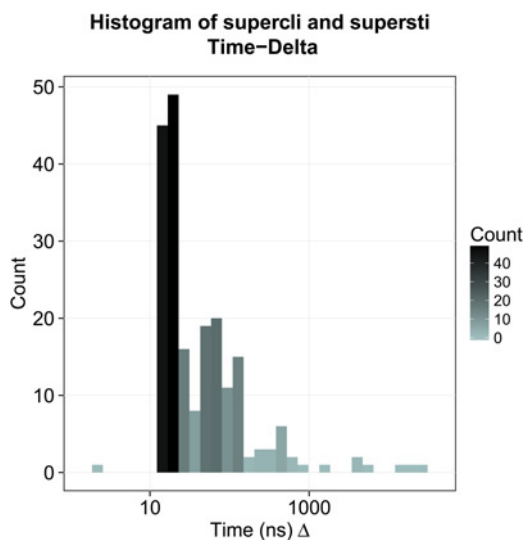


Fig. 11 Histogram of time-delta for *superSTI* and *SuperCLI* instructions generated using delta profiling algorithm

pinpoint how many instructions were executed in the interval between two consecutive interrupts disable and enable. Looking at Fig. 10, we can see that most of the interrupts disabled intervals executed around 90–100 instructions. For the same execution, we observe in Fig. 11 that most of the interrupts disabled intervals have a duration in the range 40–80 ns. We can then look for the interrupts disabled intervals of abnormally high duration which are at the far right in the histogram. Delta profiling of actual instruction flows therefore allows for an overview of the interrupt cycling in the kernel for a particular short running task. As discussed in Section 5, to get more in-depth analysis, we can take snapshots when abnormal latencies are encountered, to get a more in-depth view on identifying them from the histogram profiles.

5 Conclusion and future work

New techniques used in hardware tracing are now empowering developers in the domain of software performance debugging and analysis. We observe that hardware-assisted tracing and profiling provides a very fine granularity and accuracy in terms of control flow and time and presents a trace framework that utilises hardware-assistance from Intel PT. Our PT-based approach, with a minimal overhead in the range of 2–5%, was able to provide a highly detailed view of control flow. We also present a detailed analysis of trace size and temporal resolution provided by PT, while fine tuning its configuration. With the help of our framework, we were able to generate detailed targeted callstacks for syscalls and observe differences between multiple executions. We also demonstrated a way to trace the software tracers themselves and show how similar kernel control-flow tracers such as Ftrace cause overheads as high as 63%, while PT was able to generate similar yet more detailed results with 1% overhead. Hardware tracing also allowed us to gather traces from parts of the kernel that are invisible to traditional software tracers. PT-assisted cycle-accurate profiling was able to provide resolution as high as 14 ns, with timed events only 78 instructions apart. However, our analysis of Intel PT and the information released by Intel suggests that many additional features such as using PT in virtual machine tracing are left yet to be explored.

5.1 Latency snapshots

We hinted how, in addition to syscall latency profiles, more in-depth analysis on other non-deterministic latencies can be done. An interesting observation relevant for real-time systems is also an in-depth analysis of IRQ latency. Newer tracepoints in the kernel introduced by tracers such as LTTng [40] allow recording software-trace events when IRQ a latency beyond a certain threshold is reached. We can further refine our idea by recording hardware-trace snapshots in such conditions, thus obtaining more detailed control-flow and timing information.

5.2 Virtual machine trace and analysis

Our Intel PT-based hardware-assisted technique can also be used to detect virtual machine (VM) state transitions, in host or guest only hardware tracing, for more in-depth analysis of VMs without any software-tracing support. We observed that PT can also generate VMCS packets and set the non-root bit in PIP packets when hardware tracing is enabled in VM context [30]. The extra information in hardware traces allows the decoder to identify VM entry and exit events and load specific binaries for rebuilding control-flow across VMs and host. Thus, with no software intrusion and a low overhead, we can get accurate VM traces, compare and quantify their performance.

6 Acknowledgments

We are grateful to Ericsson, EfficIOS, NSERC and Prompt for funding and Francis Giraldeau for his technical advice and comments.

7 References

- [1] Tewar A.K., Myers A.R., Milenković A.: 'mcfTRaptor: toward unobtrusive on-the fly control-flow tracing in multicores', *J. Syst. Archit.*, 2015, **61**, (10), pp. 601–614, doi: 10.1016/j.sysarc.2015.07.005
- [2] Ball T., Burckhardt S., Halleux J., *ET AL.*: 'Deconstructing concurrency heisenbugs'. 31st Int. Conf. on Proc. Software Engineering – Companion Volume, 2009. ICSE-Companion 2009, 2009, pp. 403–404, doi: 10.1109/ICSE-COMPANION.2009.5071033
- [3] IEEE Nexus 5001. Available at <http://www.nexus5001.org/>, accessed March 2016
- [4] <http://www.ds.arm.com/ds-5/debug/dstream/>, accessed March 2016
- [5] <http://www.ghs.com/products/supertraceprobe.html>, accessed March 2016
- [6] <http://www.ds.arm.com/ds-5/>, accessed March 2016
- [7] <http://www.ghs.com/products/timemachine.html>, accessed March 2016
- [8] Tallam S., Gupta R.: 'Unified control flow and data dependence traces', *ACM Trans. Archit. Code Optim.*, 2007, **4**, (3), doi: 10.1145/1275937.1275943
- [9] Boogerd C., Moonen L.: 'On the use of data flow analysis in static profiling'. 2008 Eighth IEEE Int. Working Conf. on Proc. Source Code Analysis and Manipulation, 2008, pp. 79–88, doi: 10.1109/SCAM.2008.18
- [10] Wichmann B.A., Canning A.A., Clutterbuck D.L., *ET AL.*: 'Industrial perspective on static analysis', *Softw. Eng. J.*, 1995, **10**, (2), pp. 69–75
- [11] Livshits B.: 'Improving software security with precise static and runtime analysis'. PhD thesis, Stanford University, Stanford, CA, USA, 2006
- [12] Goseva-Popstojanova K., Perhinschi A.: 'On the capability of static code analysis to detect security vulnerabilities', *Inf. Softw. Technol.*, 2015, **68**, (C), pp. 18–33, doi: 10.1016/j.infsof.2015.08.002
- [13] Lee J., Shrivastava A.: 'A compiler optimization to reduce soft errors in register files', *SIGPLAN Not.*, 2009, **44**, (7), pp. 41–49, doi: 10.1145/1543136.1542459
- [14] Ball T., Larus J.R.: 'Optimally profiling and tracing programs', *ACM Trans. Program. Lang. Syst.*, 1994, **16**, (4), pp. 1319–1360, doi: 10.1145/183432.183527
- [15] Anderson J.M., Berc L.M., Dean J., *ET AL.*: 'Continuous profiling: where have all the cycles gone?', *ACM Trans. Comput. Syst.*, 1997, **15**, (4), pp. 357–390, doi: 10.1145/265924.265925
- [16] Dean J., Hicks J.E., Waldspurger C.A., *ET AL.*: 'ProfileMe: hardware support for instruction-level profiling on out-of-order processors'. Thirtieth Annual IEEE/ACM Int. Symp. on Proc. Microarchitecture, 1997. Proc., 1997, pp. 292–302, doi: 10.1109/MICRO.1997.645821
- [17] Merten M.C., Trick A.R., Nystrom E.M., *ET AL.*: 'A hardware mechanism for dynamic extraction and relay of program hot spots', *SIGARCH Comput. Archit. News*, 2000, **28**, (2), pp. 59–70, doi: 10.1145/342001.339655
- [18] Vaswani K., Thazhuthaveetil M.J., Srikant Y.N.: 'A programmable hardware path profiler'. Proc. Int. Symp. on Code Generation and Optimization, CGO '05, Washington, DC, USA, 2005, pp. 217–228, doi: 10.1109/CGO.2005.3
- [19] Nowak A., Yasin A., Mendelson A., *ET AL.*: 'Establishing a base of trust with performance counters for enterprise workloads'. Proc. 2015 USENIX Annual Technical Conf. (USENIX ATC 15), Santa Clara, CA, July 2015, pp. 541–548
- [20] Bitzes G., Nowak A.: 'The overhead of profiling using PMU hardware counters'. Technical Report, CERN, openlab, 2014
- [21] Larus J.R.: 'Efficient program tracing', *Computer*, 1993, **26**, (5), pp. 52–61, doi: 10.1109/2.211900
- [22] Nethercote N., Seward J.: 'Valgrind: a framework for heavyweight dynamic binary instrumentation', *SIGPLAN Not.*, 2007, **42**, (6), pp. 89–100, doi: 10.1145/1273442.1250746
- [23] Nethercote N., Mycroft A.: 'Redux: a dynamic dataflow tracer', *Electron. Notes Theor. Comput. Sci.*, 2003, **89**, (2), pp. 149–170, doi: [http://www.dx.doi.org/10.1016/S1571-0661\(04\)81047-8](http://www.dx.doi.org/10.1016/S1571-0661(04)81047-8)
- [24] Weidendorfer J.: 'Sequential performance analysis with callgrind and kcachegrind', in Resch, M., Keller, R., Himmler, V., Krammer, B., Schulz, A. (Eds.) 'Tools for high performance computing' (Springer, Berlin Heidelberg, 2008, 1st edn.), pp. 93–113
- [25] Sandon P.A., Liao Y.-C., Cook T.E., *ET AL.*: 'NTrace: a bus-driven instruction trace tool for PowerPC microprocessors', *IBM J. Res. Dev.*, 1997, **41**, (3), pp. 331–344, doi: 10.1147/rd.413.0331
- [26] Vermeulen B.: 'Functional debug techniques for embedded systems', *IEEE Des. Test Comput.*, 2008, **25**, (3), pp. 208–215, doi: 10.1109/MDT.2008.66
- [27] Vasudevan A., Qu N., Perrig A.: 'XTRec: secure real-time execution trace recording on commodity platforms'. 2011 44th Hawaii Int. Conf. on Proc. System Sciences (HICSS), 2011, pp. 1–10, doi: 10.1109/HICSS.2011.500
- [28] Pedersen C., Acampora J.: 'Intel code execution trace resources', *Intel Technol. J.*, 2012, **16**, (1), pp. 130–136
- [29] Soffa M.L., Walcott K.R., Mars J.: 'Exploiting hardware advances for software testing and debugging (NIER track)'. Proc. 33rd Int. Conf. on Software Engineering, ICSE '11, Honolulu, USA, 2011, pp. 888–891, doi: 10.1145/1985793.1985935
- [30] Intel: 'Intel processor trace' (Intel Press, 2015, 1st edn.), pp. 3578–3644, accessed March 2016
- [31] Kleen A.: 'Adding processor trace support to linux'. Available at <http://www.lwn.net/Articles/648154/>, accessed March 2016
- [32] <https://www.github.com/andikleen/simple-pt>, accessed March 2016
- [33] <https://www.github.com/01org/processor-trace>, accessed March 2016, 2015
- [34] Baugh L., Zilles C.: 'An analysis of I/O and syscalls in critical sections and their implications for transactional memory'. IEEE Int. Symp. on Proc. Performance Analysis of Systems and Software, 2008. ISPASS 2008, ICSE '11, Honolulu, USA, 2011, pp. 54–62, doi: 10.1109/ISPASS.2008.4510738
- [35] Moret P., Binder W., Villazón A., *ET AL.*: 'Visualizing and exploring profiles with calling context ring charts', *Softw. Pract. Exper.*, 2010, **40**, (9), pp. 825–847, doi: 10.1002/spe.v40:9
- [36] Adamoli A., Hauswirth M.: 'Trevi: a context tree visualization & analysis framework and its use for classifying performance failure reports'. Proc. Fifth Int. Symp. on Software Visualization, SOFTVIS '10, Salt Lake City, UT, USA, 2010, pp. 73–82, doi: 10.1145/1879211.1879224
- [37] <http://www.julialang.org/benchmarks/>, accessed March 2016
- [38] <http://www.openblas.net>, accessed March 2016
- [39] <https://www.kernel.org/doc/Documentation/trace/trace.txt>, accessed March 2016
- [40] <http://lists.lttng.org/pipermail/lttng-dev/2015-October/025151.html>, accessed March 2016