

# KPAC: Efficient Emulation of the ARM Pointer Authentication Instructions

Illia Ostapishyn<sup>1b</sup>, Gabriele Serra<sup>1b</sup>, *Member, IEEE*, Tim-Marek Thomas<sup>1b</sup>, and Daniel Lohmann<sup>1b</sup>

**Abstract**—ARMv8.3-A has introduced the pointer authentication (PA) feature, a new set of measures and instructions to sign and validate pointers. PA is already used and supported by the major compilers to protect the return addresses on the stack as a measure against memory corruption attacks. As more and more SoCs implement ARMv8.3-A and code compiled with PA is even fully backwards compatible on CPUs without (where the new instructions are just ignored), we can expect PA-enabled binaries to become standard in the near future. This gives rise to the question, if and how also systems without the native PA could benefit from the extra security provided by the return address protection. In this article, we explore KPAC, a set of efficient software-based approaches to bring the PA-based return-address protection onto the platforms without the hardware support in an easily adoptable (binary-compatible) and scalable manner. Technically, KPAC achieves this by either a synchronous trap-based emulation inside the kernel or an asynchronous novel memory-based invocation of a dedicated CPU core. Our experiments with the CortexSuite benchmarks, Chromium, and Memcached on a variety of platforms running Linux ranging from a Xilinx ZCU102 board over a Raspberry Pi 4 up to an 80-core Ampere Altra demonstrate the broad applicability and scalability of our approach. Furthermore, we discuss how the principles of KPAC can be generalized to the other suited problem areas.

**Index Terms**—Computer security - application security, modeling - emulation, software - embedded software, software - system software - operating systems.

## I. INTRODUCTION

**H**ARDWARE-BASED implementations for control-flow integrity (CFI) are becoming increasingly popular with Intel's control-flow enforcement technology (CET) [1], [2], [3] and ARM's pointer authentication (PA) [4] features being the most prominent candidates. Both provide measures to ensure the integrity of the programmer-intended control-flow by protecting the return addresses on the stack, a frequent target for the buffer-overflow attacks in combination with techniques like return- or jump-oriented programming

(ROP/JOP) [5], [6], [7]. The hardware-based implementations overcome the most significant acceptance limitations of software-based CFI techniques: poor performance [8] and issues regarding the protection of the protection measure itself [9], [10], [11], [12]. While the ARMv8.3-A PA feature is long supported by standard compilers [13], [14] and the Linux kernel (in contrast to the Intel's CET, which only very recently made it into Linux [15]), for the last five years only Apple's A12/M1 actually implemented it. However, this is currently changing with Qualcomm's Snapdragon 8cx Gen 3 [16], which includes the PA support. As the PA-enabled binaries are fully backwards compatible (the special new instructions inserted by the compiler to encode/decode return addresses resolve to NOPs on CPUs without), we can expect to see a much broader adoption in the near future. Therefore, we consider it worthwhile to explore how and at what costs it would be possible to emulate the PA feature for the return address protection on the platforms without native PA.

### A. About This Article

In this article, we provide, discuss, and evaluate four different approaches to emulate the PA-based return-address protection on the ARM processors without the PA support. We compare our results to the only attempt in this direction we are aware of, which is PAC-PL of Serra and colleagues [17], who employed an FPGA for the hardware-based encryption/decryption of return addresses. While PAC-PL provides an acceptable performance impact (negligible in many cases, up to  $3\times$  in some cases), it also comes with a number of drawbacks. First, PAC-PL is not binary compatible, as the code has to be compiled with a custom GCC extension. Furthermore, it requires the availability of an FPGA, which alone makes it unsuitable for many application scenarios. Consequently, their work triggered our attempt to look for more efficient software-based and, if possible, also binary-compatible approaches.

In a nutshell, we present a  $2 \times 2$  matrix of software-based approaches that *either* require recompilation (like PAC-PL) *or* are binary compatible (via code patching) and *either* execute synchronously (by trapping) *or* asynchronously (by employing a dedicated CPU core) and compare them to *kpacpl*, a reimplement of PAC-PL by its author. Our results show that with the extra core (which in real-world settings is arguably more available/affordable than the on-board-FPGA), we outperform the programmable logic (PL)-based approach in all the cases. Without the extra core, the synchronous and binary-compatible variant comes at a *worst-case* overhead of

Manuscript received 12 August 2024; accepted 12 August 2024. This work was supported in part by the German Research Foundation (DFG) under Grant LO 1719/4-1 (391395160). This article was presented at the International Conference on Embedded Software (EMSOFT) 2024 and appeared as part of the ESWEEK-TCAD special issue. This article was recommended by Associate Editor S. Dailey. (*Corresponding author: Illia Ostapishyn.*)

Illia Ostapishyn, Tim-Marek Thomas, and Daniel Lohmann are with the Institute of Systems Engineering (ISE-SRA), Leibniz Universität Hannover, 30167 Hannover, Germany (e-mail: ostapishyn@sra.uni-hannover.de; thomas@sra.uni-hannover.de; lohmann@sra.uni-hannover.de).

Gabriele Serra is with the TECIP Institute, Scuola Superiore Sant'Anna, 56127 Pisa, Italy (e-mail: gabriele.serra@santannapisa.it).

Digital Object Identifier 10.1109/TCAD.2024.3443773

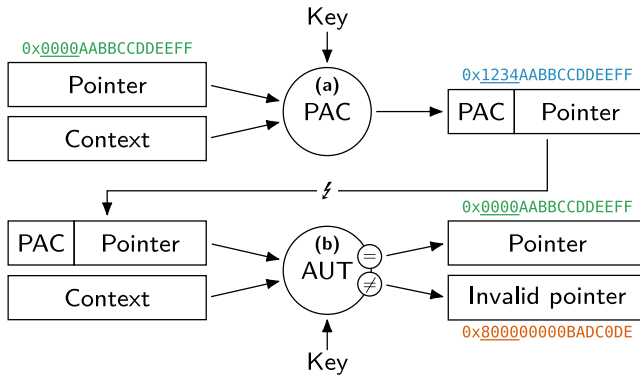


Fig. 1. PA mechanism.

85 17.37 $\times$ , which is orders of magnitude below the costs for  
 86 the software emulation reported so far [17]. For instance,  
 87 Chromium on a Raspberry Pi 4 receives an actual slowdown by  
 88 7.13 $\times$  in the JetStream benchmark, which could be considered  
 89 as acceptable for the security-critical Web applications.

90 In particular, we claim the following contributions.

- 91 1) We describe KPAC, an approach for efficient software-
- 92 based and optionally ARMv8.3-ABI-compatible PA as
- 93 an extension to the Linux kernel.
- 94 2) We provide the remote-core system call (RCSC), a novel
- 95 mechanism for efficient and safe interaction between the
- 96 user-mode threads and dedicated kernel cores.
- 97 3) We explore and evaluate the design space for the KPAC
- 98 on a variety of benchmarks and platforms.

99 The remainder of this article is organized as follows.

100 Section II presents the ARM PA mechanism and Serra  
 101 et al.'s [17] implementation based on the programmable logic.  
 102 Section III describes the assumed threat model, Section IV  
 103 our approach, and Section V the concrete implementation.  
 104 In Section VI we evaluate the implementation variants and  
 105 discuss our findings in Section VII. Finally, we review the  
 106 further relevant literature in Section VIII and conclude this  
 107 article in Section IX.

## 108 II. BACKGROUND

### 109 A. ARMv8.3-A Pointer Authentication

110 pointer authentication (PA) is an approach to protect  
 111 the code and data pointers with negligible footprint in  
 112 performance, memory, and hardware. The key idea is to utilize  
 113 the free bits in the unused upper part of pointers to store a  
 114 cryptographic hash of the pointer value as its signature, so  
 115 that unintended modifications can easily be detected. Typical  
 116 memory configurations on AArch64 require only 48-bit virtual  
 117 addresses, which leaves 16 bits for the signature, called the  
 118 pointer authentication code (PAC) [4]. The signature algo-  
 119 rithm is left to the implementation; ARMv8.3-A suggests the  
 120 QARMA block cipher [18], which can efficiently be realized  
 121 in hardware.

122 The mechanism features instructions for the creation and  
 123 validation of these signatures. Fig. 1(a) visualizes the signing  
 124 instructions using the mnemonic PAC. These instructions take  
 125 three values, the 64-bit pointer itself, a 128-bit key (implicitly),  
 126 and a 64-bit context information to produce a pointer with

a PAC in its upper bits. The ARM implementation features  
 127 five keys: two for instructions and data pointers each, and one  
 128 general-purpose key. They are stored in the system control  
 129 registers that are accessible only by higher privilege levels (i.e.,  
 130 the operating system kernel) and, thus, kept secret from the  
 131 user applications requesting authentication. Linux, Windows,  
 132 and XNU [16], [19], [20] manage these keys on a perprocess  
 133 basis for the systems with the PA extension; Linux 5.7+, and  
 134 XNU even support PA inside the kernel itself [21].  
 135

136 After the pointer has been signed using PAC instructions,  
 137 their counterparts based on the mnemonic AUT are responsible  
 138 for verification of signature before usage [Fig. 1(b)]: The AUT  
 139 instructions take the authenticated pointer, recompute the PAC,  
 140 and compare the result with the code stored in the signed  
 141 pointer. If the signature matches, the PAC is stripped from  
 142 the pointer. Otherwise a trap will occur, either immediately  
 143 (ARMv8.6-A) or upon dereferencing of the pointer.

144 GCC and LLVM compilers already employ PA [13], [14] to  
 145 protect the function return addresses (the *backward edges* of  
 146 the control flow) that might be stored on the stack, where they  
 147 would become vulnerable for the buffer overflow attacks. This  
 148 is done by inserting PACIASP and AUTIASP instructions,  
 149 operating on the link register (LR/X30) with the stack pointer  
 150 (SP) as the context value in the function prologues and  
 151 epilogues, respectively. In the ARM ISA, these instructions  
 152 are located in the NOP instruction space, which ensures the  
 153 backward compatibility of newly compiled programs with  
 154 CPUs lacking the PA extension. As leaf functions never push  
 155 their return address onto the stack, the standard setting is to  
 156 omit the PA instructions in them.

157 However, half a decade following the introduction of the PA  
 158 mechanism in the ARM specification and despite ubiquitous  
 159 compiler and OS support, only few systems are readily avail-  
 160 able that implement it in hardware. Most notably, the A12 chip  
 161 presented by Apple in 2018 and all its successors come with  
 162 PA [22] mechanism. This has only recently be complemented  
 163 by Qualcomm's Snapdragon 8cx Gen 3 SoC [16], which  
 164 brings the PA also to the Windows and Android domains.  
 165 Nevertheless, we face a plethora of systems with no support  
 166 for the PA and adoption will continue to be slow, especially  
 167 in the embedded domain.

### 168 B. PA Using FPGA: The PAC-PL Approach

169 As a solution for this, Serra et al. [17] implemented the PA  
 170 mechanism on an SoC featuring an field programmable gate  
 171 array (FPGA). Since, we base our work on theirs and use a  
 172 reimplementaion as a comparison point, we briefly present  
 173 and discuss it here.<sup>1</sup>

174 The main idea of PAC-PL is to perform the signing and  
 175 authentication of pointers using PL on the SoC. Its architecture  
 176 consists of two components: 1) a QARMA block cipher [18]  
 177 crypto engine and 2) an AXI subordinate device, which  
 178 handles interaction between the crypto engine and the host  
 179 over the AXI bus via the memory-mapped registers, which  
 180 are mapped into the kernel- and user-level address spaces,

<sup>1</sup>Unfortunately, the original PAC-PL code underlies the IP restrictions, but its author provided us with a personal reimplementaion of its core features.

181 respectively. Instead of using the ARMv8.3-A PAC/AUT  
 182 instructions, the signing/authentication of pointers is triggered  
 183 by writing into the corresponding registers, which lets the PAC-  
 184 PL accelerator generate, remove, and check the PAC. Hence,  
 185 the approach is not binary-compatible: the software has to  
 186 be compiled with a custom GCC plugin that generates the  
 187 necessary instructions.

188 Since, QARMA is designed to be particularly fast in  
 189 hardware, the overheads are dominated by the communication  
 190 latency, which is costly due to the mismatch in the clock  
 191 frequencies between the FPGA and the host CPU. While  
 192 calculating the cipher itself only requires ten host cycles, a  
 193 complete PAC/AUT operation takes at least 426 cycles. In our  
 194 measurements on a Xilinx ZCU102 at 1.2 GHz this approach  
 195 leads (with the most strict PA application mode all explained  
 196 later) to an average overhead of 34% for the CortexSuite [23]  
 197 benchmarks. The operation time is bounded, making it suitable  
 198 for the real-time systems.

199 In the paper [17], the utilization of an FPGA is partly  
 200 justified by comparing it to the performance results from  
 201 a software-based emulation of their approach, which bears  
 202 much higher (up to five orders of magnitude!) overheads.  
 203 However, this extremely high overhead is likely caused by  
 204 the employed user-kernel interface, which induces two page  
 205 faults per PAC/AUT transaction (hence, the four page faults per  
 206 protected function) to emulate a PAC-PL device. Furthermore,  
 207 while QARMA is optimized for the hardware implementa-  
 208 tions, another cipher might be more suitable for a CPU-based  
 209 software implementation. Last but not least, the work does not  
 210 evaluate nor mention support of the multithreaded applications.  
 211 In the remainder of this article, we explore the options for  
 212 more efficient and optionally binary-compatible PA emulation  
 213 that scales well in concurrent environments.

### 214 III. SECURITY OBJECTIVES AND THREAT MODEL

215 ARMv8.3 pointer authentication was developed to accom-  
 216 plish the pointer integrity. Intuitively, pointer integrity seeks to  
 217 prevent the alterations to pointers while residing in memory,  
 218 ensuring that the value of a pointer at the time of its use  
 219 (i.e., dereferencing) remains consistent with the value intended  
 220 during its creation or storage. Control-flow attacks and numer-  
 221 ous other data-oriented attacks hinge on manipulating the  
 222 susceptible pointers. Consequently, the enforcement of pointer  
 223 integrity defends against these attacks. The security objective  
 224 of ARMv8.3 PA, therefore, consists of preventing the attacker  
 225 from forging pointers used by a vulnerable program.

226 Likewise, KPAC pursues the same security guarantees. Our  
 227 approach shall satisfy the following functional requirements.

- 228 1) *Pointer Integrity*: Prevent and detect the use of the  
 229 corrupted code or data pointers.
- 230 2) *Attack Resistance*: Resist attempts to forge the valid  
 231 pointers and resist pointer reuse attacks.

232 Further, we identify nonfunctional requirements, which  
 233 allow wider compatibility as follows.

- 234 1) *Compatibility*: Enabling pointer integrity protection of  
 235 existing programs without interfering with their opera-  
 236 tion even without dedicated hardware support.

- 2) *Performance*: Minimize run-time overhead by providing  
 237 configurable protection scopes as a tradeoff between the  
 238 hardening and performance. 239

The following assumptions define the attacker’s capability,  
 consistent with the prior works in this area ([24], [25]). Our  
 adversary model reckons with an attacker as follows. 242

- 1) With unrestricted user-space memory read and write  
 243 capabilities, constrained exclusively by the data exe-  
 244 cution prevention (DEP) mechanism, therefore with  
 245 the ability to read any program memory and write  
 246 to the nonexecutable segments exploiting the input-  
 247 controlled memory corruption errors in the victim  
 248 process (e.g., controlling return addresses, function  
 249 pointers, or VTable pointers). 250
- 2) Disposes of a full knowledge of the process memory  
 251 layout and has successfully bypassed the address space  
 252 layout randomization (ASLR), if present. 253
- 3) With no control over privilege levels higher than the user  
 254 level, meaning without the ability to access the kernel  
 255 space or higher privilege levels. 256

Note that, assumptions 1 and 2 rule out the feasibility of  
 randomization-based defenses susceptible to the information  
 disclosure, such as stack canaries, ASLR, or software shadow-  
 stacks. KPAC was designed to maintain its effectiveness even  
 when the complete memory layout of the victim process is  
 disclosed as long as the assumption 3 holds. Therefore, the  
 attacker cannot deduce the keys, which are located in memory  
 not directly readable from the user space. 264

According to the presented threat model, KPAC is as secure  
 as ARMv8.3 PA. The PAC-PL Serra et al. [17] have obsoleted  
 the assumption 3 by employing ARM TrustZone, which cre-  
 ates isolated secure environments to protect the sensitive data  
 for the key management. This extra protection is applicable to  
 KPAC as well, but not further explored in this article. 270

### 271 IV. KPAC APPROACH

A key point of the ARMv8.3-A PA (also mimicked by PAC-  
 PL) is that it delegates key management to the OS running  
 on the EL1 privilege level (the supervisor mode) ensuring  
 higher protection. In order to stay true to this property, KPAC  
 delegates key management and exception handling to the OS  
 kernel. As a corollary, the partial interpretation of the PAC  
 and AUT operations by software has to take place inside the  
 kernel, which generally induces the significant overhead, as  
 every operation thereby comes with a minimum of two user-  
 kernel context switches. Mitigating this overhead as far as  
 possible is one key to an efficient software implementation.  
 The other key is the overhead of the signing algorithm itself. 283

The central component of KPAC is a Linux kernel extension,  
 which implements the PA backend. Since, QARMA is not  
 suitable for the fast software implementation, *SipHash* [26] has  
 been selected as the cryptographic hashing algorithm instead.  
 It is designed to be efficient and secure with short inputs  
 to compute a 64-bit message authentication code, which is  
 truncated to the unused bits of the pointer. 290

The kernel extension exposes two interfaces for the user-  
 space applications to request the PA (Fig. 2) as follows. 292

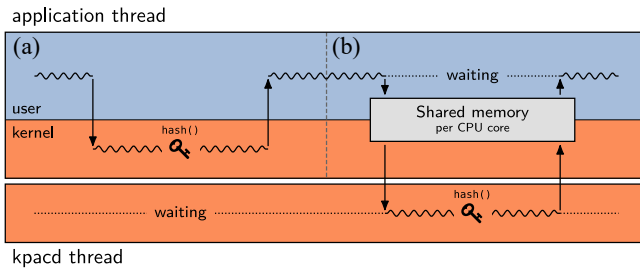


Fig. 2. Application requesting PA by making (a) an *SVC* request followed by (b) a *kpacd* request.

TABLE I  
PROPERTIES OF PRESENTED EMULATION APPROACHES

Approach	Hardware requirements	Multi-threaded	Binary compat.	Bounded WCET	Average overhead
<i>native</i>	ARMv8.3-A	✓	✓	✓	0 %
(C) <i>kpacpl-static</i>	FPGA			✓	34 %
(AC) <i>svc-static</i>		✓		✓	88 %
(BC) <i>kpacd-static</i>	extra core	✓		✓	17 %
(D) <i>kpacpl-libkpac</i>	FPGA		✓	✓	44 %
(AD) <i>svc-libkpac</i>		✓	✓	✓	87 %
(BD) <i>kpacd-libkpac</i>	extra core	✓	✓	✓	31 %

- 1) *Synchronous system calls (SVC requests)*, which execute the PA within the invoking thread. This is the canonical way to implement an user kernel interaction.
- 2) *RCSC*, a novel asynchronous communication protocol based on the per-CPU shared memory, which executes the PA on a dedicated kernel core running the *kpacd* daemon. This (kind of) mimics the idea of PAC-PL to use extra hardware (here a CPU core instead of an FPGA) for the PA.

To instrument applications with either invocation scheme, two methods have been investigated as follows.

- 1) *Static instrumentation* by a compiler plugin (as in PAC-PL). This is, assumingly, the most run-time efficient way, as it facilitates static optimization of the code and also provides configurable protection scopes for the overhead mitigation.
- 2) *Load-time instrumentation* by a dynamic library (*libkpac*) that is applied by the `LD_PRELOAD` feature of the system's dynamic loader and patches at load time all PACIASP/AUTIASP instructions in the code to invoke KPAC instead. This provides full binary compatibility for ARMv8.3-A binaries that were compiled with PA support.

Table I briefly summarizes the resulting four KPAC variants, together with PAC-PL and a native ARMv8.3-A processor. The given overhead numbers should be considered as a ballpark figure only. They describe the geometric mean over all the CortexSuite benchmarks on a Xilinx ZCU102 at 1.2 GHz. Our experiments with Apple's M1 Ultra did not yield any measurable overhead for the native ARMv8.3-A PA.

## V. IMPLEMENTATION

We integrated KPAC into the Linux kernel version 6.1. The compiler support for the static instrumentation is provided as a GCC 12.2 plugin.

### A. *SVC: The Synchronous System Call Interface*

The AArch64 instruction set defines the *SVC* (supervisor call) instruction, which transfers the control flow to the EL1 privilege level running the OS kernel. This instruction is used across the operating systems to implement the system calls and has a 16-bit immediate argument. On Linux, the system call number is passed in the `W8` register and the immediate argument of the *SVC* instruction is ignored.

We extend the Linux system call interface to emulate the ARMv8.3-A PACIASP and AUTIASP instructions by reserving two values of the *SVC* instruction's immediate argument. These new emulation calls thereby require a single instruction in the code that only alters the LR, making them semantically equivalent to the PACIASP/AUTIASP instructions emitted by the standard compilers. As the execution time of the *SVC* instruction and the SipHash algorithm takes bounded time [26], the emulation is also suitable for the hard real-time settings that demand bounded WCETs.

### B. *kpacd: The Remote-Core System Call Interface*

On many platforms, context switches into the OS kernel induce a high overhead for changing the privilege level and the address space. Furthermore, the executed kernel code may put extra pressure to the CPU-local caches and the TLB, significantly impairing the performance [27]. An alternative approach is to run the kernel services asynchronously on a dedicated core [28], [29] that always stays in the kernel mode. The services are invoked by a shared-memory interface between both the cores, omitting the above overhead altogether. Our *rsc* implements this idea for the Linux, while additionally providing for the lock-free per-core separation.

With RCSC, one or several CPU cores are reserved for the KPAC and execute the *kpacd* (*Kernel PAC Daemon*) in the kernel mode, which polls a shared memory page for the PA requests. This is comparable to the PAC-PL, where the service core acts as the accelerator instead of an FPGA.

Sacrificing a full core just for the PA purposes might appear as an odd design decision, given that such core induces a *much* higher hardware overhead than a small FPGA. However, in practice, an unused core is way more often available and actually cheaper for many embedded systems than an FPGA.

*Invocation of kpacd:* Listing 1 demonstrates the assembly code corresponding to an authentication RCSC to the *2kpacd* thread. After storing the pointer and the context value at the respective offsets in the RCSC page (`L3`), the application hands off the request by writing the operation code into the first *status word* of the page (`L6`), which wakes the remote *kpacd* to perform the requested operation. The status word is then checked in a loop (`L9–11`) by loading the first word of the page with the exclusive load (`LDXR`) instruction, branching to `WFE` if the value is not zero (zero signals completion). The `WFE` instruction hints the CPU core to enter a low-power state, until a wake-up event occurs [4]. A remote store (by the *kpacd* core) to this location, which was recently read using an exclusive load (`LDXR`), generates such an event. Hence, on both the sides the polling does not come with an extra energy/heat overhead. The combination of `LDXR` and `WFE` is

```

1  mov    x9, #KPAC_BASE
2  mov    x10, sp
3  stp    lr, x10, [x9, #REG_PLAIN] // store pointer and context
4
5  mov    x10, #OP_PAC
6  stlr   x10, [x9]                // request operation
7
8  sevl
9 1:    wfe                          // sleep for event <--+
10     ldxr  x10, [x9]              //
11     cbnz  x10, 1b                // until completion ----+
12     ldr   lr, [x9, #REG_CIPHER] // obtain result

```

Listing 1. Assembly code of a function prologue requesting a signed pointer from *kpacd*.

also used in the AArch64 `__CMPWAIT_CASE` macro of the Linux kernel.

**Multithreading Support:** On multiprocessor systems, multiple threads from within the same or different processes might invoke *kpacd* simultaneously. These concurrent requests need to be isolated and coordinated. RCSC solves this by providing an individual RCSC page for each core, which is (implicitly) used by the thread currently executing on this core. Hence, no synchronization is required when accessing the RCSC page, enabling scalability. As each core executes exactly one thread at a time, the per-core pages also ensure isolation. Upon a switch to another thread, the scheduler completes any pending RCSC requests and saves the relevant content (24B) of the shared page in the thread control block.

Technically, the provision of per core pages (which we consider a general mechanism) has to be integrated with the virtual memory subsystem. For this, the data structure representing the address space and containing the pointer to the top-level page directory (*page global directory* and *PGD*), is extended to support a different PGD per core. As illustrated in Fig. 3, all the entries in these PGDs are kept synchronized except for one. The entry leading to the core-local RCSC page. Thereby, all cores use the same virtual address to access their core-specific memory. This comes at the cost of duplicated PGDs for processes using the *kpacd* service. Moreover, when a PGD entry is modified, the changes have to be mirrored into the PGDs of other CPU cores. The performance overhead of this is negligible, since the top-level page-directory entries are only populated at the process start and rarely modified during the execution. As the underlying page tables are shared, all the changes in them (e.g., induced by an `mmap()`) are immediately seen by other CPU cores and require no further mirroring nor synchronization.

For load balancing in larger multicore systems, an arbitrary number of cores could be assigned to *kpacd*. Each *kpacd* core services a fixed set of application cores in a round-robin manner. Hence, the worst-case service time of *kpacd* is also bounded in multicore settings.

### C. Static Instrumentation via Compiler Plugin

Applying either of the KPAC invocation approaches for the return-address protection requires adding the signing and authentication code in the prologues and epilogues of functions. One way to achieve this, also taken by Serra and associates [17], is to employ a compiler plugin and add an additional pass working on the register-transfer language representation of the program.

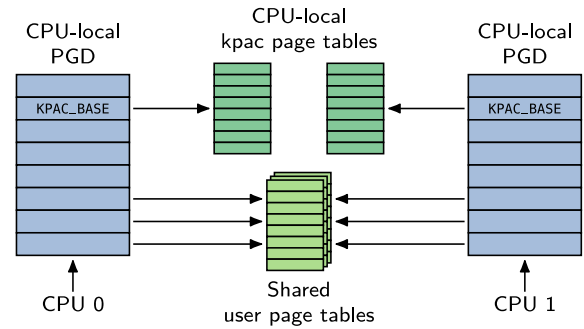


Fig. 3. Page table arrangement introduced by CPU-local top-level page directories (PGDs).

**Protection Scopes:** If compiling code for ARMv8.3-A with `-mbranch-protection`, GCC would apply the PA-based return address protection on the prologues and epilogues of all the nonleaf functions (which push the return address to the memory). As the PA-based return address protection basically just adds two instructions to a protected function, this does not induce any measurable overhead. In contrast, a PA emulation induces a much higher overhead, so it might be worthwhile to explore different protection scopes and let the compiler plugin only instrument the most vulnerable functions.

Our compiler plugin therefore resembles the protection levels of GCC's and Clang's `-fstack-protector` feature [13], [14], a purely compiler-based CFI measure that comes with the common limitations regarding performance and actual protection (cf. Section I). However, its defined protection levels are established among the developers who have to trade between the hardening and performance of their software. The three different protection scopes are referred to in the following as *char*, *strong*, and *all*.

**char** protects the nonleaf functions that place char arrays of at least size 8 (*ssp-buffer-size*) on the stack and nonleaf functions, that perform dynamic stack allocation with `alloca()`. This protection scope bears the lowest performance impact, while already providing some protection for simple but common buffer-overflow attacks.

**strong** extends the scope of protected functions to the nonleaf functions that accommodate any arrays or variables that have their address taken on the stack. This further mitigates the range of some advanced attack techniques based on the ROP at a moderate performance impact.

**all** extends the scope even further to all the (nonleaf) functions, which provides the highest protection level, but also induces significant performance costs.

Our plugin supports all the optimization levels, but automatically disables the *ipa-ra* and *shrink-wrap* optimizations, as we depend on the caller-saved registers to be actually saved and the function prologue at the beginning of a function.

### D. Load-Time Instrumentation via *libkpacd*

While recompiling the existing applications might be feasible in some settings (e.g., embedded applications), this is often

<pre> 1 paciasp // PAC lr 2 sub sp, sp, #0x60 3 stp x20, x19, [sp, #32] 4 stp fp, lr, [sp, #16] 5 6 // function body omitted 7 8 ldp fp, lr, [sp, #16] 9 ldp x20, x19, [sp, #32] 10 add sp, sp, #0x60 11 autiasp // AUT lr 12 ret </pre>	<pre> 1 sub sp, sp, #0x60 2 stp x20, x19, [sp, #32] 3 stp fp, lr, [sp, #16] 4 bl kpacd_pac_24 // PAC [sp+24] 5 6 // function body omitted 7 8 bl kpacd_aut_24 // AUT [sp+24] 9 ldp fp, lr, [sp, #16] 10 ldp x20, x19, [sp, #32] 11 add sp, sp, #0x60 12 ret </pre>
(a)	(b)

Listing 2. Example function from Memcached patched by *libkpac* for *kpacd* invocation. The `kpacd_{pac,aut}_24` trampoline operates on the return address at the offset 24 from the SP (a) Before patching (b) After patching.

not the case, especially in end-user environments. Thus, we propose a binary-compatible method of adding the software-emulated PA to the programs already compiled for ARM PA by providing a run-time library (*libkpac*). *libkpac* patches the program at load time and can be applied selectively to the whole system or single application processes.

Technically, *libkpac* is injected by setting the `LD_PRELOAD` environment variable, which causes the system’s dynamic loader to additionally load the library and execute its constructor function. The `LD_PRELOAD` mechanism provides for maximum flexibility on the user’s side. For example, the user might run one instance with KPAC support to improve the security and another one without, for the performance-sensitive activities.

The constructor function parses the memory map of the process, exposed by Linux in the *procfs* file system, and takes note of the executable memory areas in the address space. It then iterates over these areas and searches for the `PACIASP` and `AUTIASP` instructions. At these places, the code needs to be patched to invoke KPAC by either the synchronous *SVC* or the asynchronous *RCSC* mechanism.

***SVC-Only Mode:*** In this mode, the `PACIASP` and `AUTIASP` instructions are simply replaced by their respective *SVC* equivalents. As this takes only a single opcode and clobbers the same set of registers (just the LR register), this is trivially possible in all the cases.

***kpacd and kpacpl Modes:*** Invoking *kpacd* or *kpacpl* via their shared-memory interface requires inserting additional branches to a subfunction, which is more complicated and not (safely) possible in all the cases. The general idea of patching a function for such invocation is demonstrated in Listing 2. Fundamentally, it is not possible to just replace `PACIASP` and `AUTIASP` by a call to the *kpacd/kpacpl* invocation, as this would overwrite the return address stored in the LR to be protected. Instead, the invocations have to be put at the end of the prologue (beginning of the epilogue), when LR has been saved onto the stack. The required space for these calls is created by shifting the stack frame (de)allocation sequences into the `PACIASP/AUTIASP` instructions. However, as compilers might do arbitrary things in their function pro/epilogues (e.g., reordering), the patching falls back to the *SVC* mechanism if no familiar stack frame (de)allocation sequence is detected. When instrumenting binaries for *kpacpl*, the *SVC* fallback uses the accelerator for QARMA computation from the kernel space.

Upon invocation, the actual return address to be en/decoded resides on the stack, but at an varying offset that depends on the function-specific stack frame. To deal with this, *libkpac* provides trampoline functions for the offsets from 0 to 504 bytes (at machine word granularity). As the AArch64 BL instruction can jump only in the region of  $\pm 128$  MiB, *libkpac* furthermore places the trampolines in neighboring address-space holes in the case of large text sections. For example, this is required to fully patch the Chromium’s 161.02 MiB executable segment.

## VI. EVALUATION

In our evaluation, we 1) demonstrate the latency of a single PAC/AUT transaction; 2) show that our approaches efficiently implement PA in software; 3) illustrate the multicore scalability using the memory caching system Memcached; and 4) showcase the ease use of the binary-compatible approaches using the *Chromium browser*.

We have integrated our mechanism into Linux 6.1 on the three systems: 1) Xilinx Zynq UltraScale+ ZCU102 evaluation board with XCZU9EG MPSoC at 1.2 GHz; 2) Raspberry Pi 4 single-board computer with Broadcom BCM2711 at 1.8 GHz; and 3) Gigabyte R152-P31 rack server with 80-core Ampere Altra Q80-30 CPU at 3 GHz. The ZCU102 evaluation board allows us to directly compare our approaches with the PAC-PL reimplementations as the SoC features an FPGA fabric on the chip. The Raspberry Pi resembles a typical medium-end hardware used in embedded appliances. This does obviously not hold for the 80-core Ampere Altra/Memcached setup, which we include for the sole purpose of stressing the multicore scalability of our approach.

As the baseline, we chose to run the targets without any enabled PA. This corresponds to our measurements on the Apple’s M1 Ultra (see Table I), which did not yield any measurable overhead for the native PA.

### A. Cost of User–Kernel Interaction and Hashing

First, we evaluate the cost of the user kernel interaction methods introduced in Section IV by measuring the end-to-end latency of the transactions on the mentioned systems. Simultaneously, we demonstrate the high overhead of the QARMA hashing algorithm when implemented in software and motivate the choice of SipHash for fast hashing. Table II demonstrates the 99th percentile latencies for the PA requests in clock cycles over 32 million samples. The cycles are measured using the PMU’s cycle counter `PMCCNTR_ELO`. Consequently, the *kpacd* spin loop does not use WFE, as it is undefined whether the counter continues to increment in low-power state [4].

The measurement in the *None* row of Table II does not perform any hashing and thus represents the raw communication overhead for the *SVC*- and *RCSC*-based transactions. Comparing the raw communication overhead across the systems, both the ZCU102 and Raspberry Pi 4 require over 2000 cycles for an NOP system call. The *kpacd* request on these systems is much faster, by factor 4.99 on the Raspberry Pi 4 and by factor 9.31 on the ZCU102 evaluation board.

TABLE II  
99TH PERCENTILE ROUND-TRIP LATENCY OF PA REQUESTS IN CLOCK CYCLES. (A) AMPERE ALTRA Q80, 3 GHz. (B) BROADCOM BCM2711, 1.8 GHz. (C) XILINX XCZU9EG, 1.2 GHz

(A)			(B)		
Hashing alg.	<i>svc</i>	<i>kpacd</i>	Hashing alg.	<i>svc</i>	<i>kpacd</i>
None	389	476	None	2095	420
SipHash	434	488	SipHash	2170	508
QARMA	3903	3947	QARMA	6779	5332

(C)			
Hashing algorithm	<i>kpacpl</i>	<i>svc</i>	<i>kpacd</i>
None	643	2020	217
SipHash	—	2122	339
QARMA	650	11608	8231

On the Xilinx ZCU102 evaluation board, a *kpacd* transaction takes only 217 clock cycles, which amounts to 180.8 ns. On the same system, a round trip to the PL takes 643 cycles or 535.8 ns. In contrast, the Ampere Altra Q80 system shows a different picture: a system call is 18.28% faster than a round trip over the shared memory and takes only 389 clock cycles or 129.7 ns. This demonstrates that high-end systems might implement the system calls more efficiently and could profit from the fast software PA without an additional accelerator core.

Moving onto the hashing algorithms, the ARM-suggested QARMA cipher comes with an overhead of at least 3000 cycles on all the systems even when subtracting the raw communication costs. For instance, a system call calculating the QARMA cipher takes 9.67  $\mu$ s on the ZCU102 evaluation board. Having an FPGA at its disposal, this is the only system providing a low latency for QARMA with *kpacpl* taking 650 clock cycles to authenticate a pointer. In fact, the QARMA calculation costs are fully amortized by the communication overhead, as the round trip latency is nearly equal to the latency without hashing. This stems from the mismatch in the clock frequencies between the FPGA and the host CPU, together with the costs for the data transfer.

SipHash offers a practical alternative to the QARMA cipher. Subtracting the communication overhead, its calculation takes roughly 100 cycles on all the systems. This results in the round-trip latency of 339 cycles (282.5 ns) on the ZCU102 evaluation board, 434 cycles (144.7 ns) on the Ampere Altra machine via *svc*, and 508 cycles (282.2 ns) on the Raspberry Pi 4. Therefore, we use SipHash in the remainder of evaluation.

## B. Approach Comparison

For the comparison of the KPAC approaches, we chose the CortexSuite [23]. It is a representative embedded workload, as it consists of machine learning, computer vision, language processing, and IoT tasks. As the baseline, we compile all the benchmarks without any protection using GCC 12.2 with the `-O2` optimization level. For the static instrumentation, the benchmarks are compiled with our compiler plugin with the

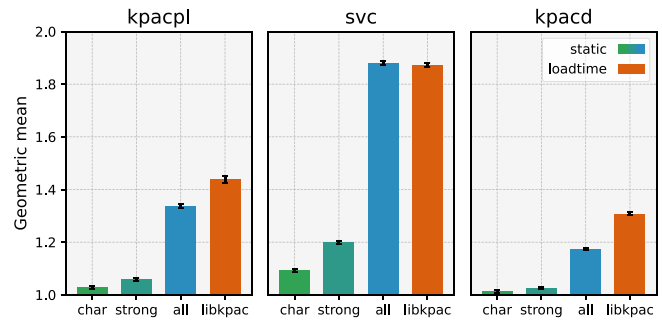


Fig. 4. Geometric mean of CortexSuite benchmark run durations normalized to the baseline run on Xilinx ZCU102.

same optimization level.<sup>2</sup> For the binary-compatible *libkpac* evaluation, the compilation flags are complemented with `-mbranch-protection=pac-ret` to add the return-address protection using ARMv8.3-A PA. The benchmarks are executed on the Xilinx Zynq UltraScale+ ZCU102 evaluation board, allowing us a direct comparison to the PL-based approach *kpacpl*.

*Static Instrumentation:* The advantage of instrumenting applications statically using the compiler plugin lies in the ability to mitigate the performance overhead using the protection scope heuristics introduced in Section V-C. Thus, we evaluate the three protection scopes and compare the communication approaches to each other. Fig. 4 provides a high-level overview of the overhead over all the CortexSuite benchmarks (summarized using the geometric mean) and Table III breaks down the figure for the individual benchmarks.

The highest overhead is measurable for the most secure protection scope *all*, where all the nonleaf functions are protected. The *svc*-based instrumentation has the highest average overhead of 1.88 $\times$  and is outperformed by *kpacpl* with the average overhead of 1.34 $\times$ . The *kpacd* approach leads in this category with the average overhead of 1.17 $\times$ .

The protection scope *strong* reduces the overhead to a lower figure for all the approaches, while keeping the security guarantees high as the likelihood of a stack-buffer overflow occurring in a function that never exposes the addresses to its stack is exceedingly low. There, *kpacd* still outperforms all the approaches with an average overhead of 1.88 $\times$ . The worst-case overhead for *kpacd-strong* is observed in the *sphinx* benchmark with 1.20 $\times$  or 20%. The respective *kpacpl* figures are 1.06 $\times$  for the average and 1.48 $\times$  for the worst case in *sphinx*. Even for *svc*, which is inherently suboptimal on this system due to the cost of the system calls, the average overhead is reduced to 1.20 $\times$  with the worst case of 3.34 $\times$ .

Depending on the application, the security can be traded for the performance by reducing the protection scope to the *char*: functions that allocate the character arrays on the stack. Note that, this is the default mode of GCC's stack protector and the only one evaluated by Serra and associates for PAC-PL originally. This reduces the worst-case overhead to 1.26 $\times$  for *kpacpl*, 2.25 $\times$  for *svc*, and 1.11 $\times$  for *kpacd*. On average, *char* yields the lowest overhead regardless of the approach.

<sup>2</sup>This excludes optimizations incompatible with current compiler plugin prototype discussed in Section V-C.

TABLE III  
 RUN TIMES OF CORTEXSUITE BENCHMARKS NORMALIZED TO THE BASELINE RUN WITHOUT PROTECTION ON XILINX ZCU102

Benchmark	Baseline run duration [s]	<i>kpacpl</i>				<i>svc</i>				<i>kpacd</i>			
		<i>char</i>	<i>strong</i>	<i>all</i>	<i>libkpac</i>	<i>char</i>	<i>strong</i>	<i>all</i>	<i>libkpac</i>	<i>char</i>	<i>strong</i>	<i>all</i>	<i>libkpac</i>
lda	18.25	1.01	1.01	4.37	7.53	1.04	1.04	17.38	17.37	1	1	2.38	5.37
sphinx	12.39	1.26	1.48	3.21	3.6	2.25	3.34	11.78	11.43	1.11	1.2	1.87	2.3
rbm	21	1	1.08	1.17	1.17	1	1.4	1.8	1.8	1	1.03	1.06	1.07
srr	29.12	1.01	1.04	1.04	1.11	1.01	1.16	1.16	1.16	1	1.02	1.02	1.08
svd3	14.55	1	1	1.02	1.03	1	1	1.14	1.14	1	1	1.01	1.01
motion-est.	9.42	1.03	1.04	1.05	1.03	1.03	1.08	1.13	1.1	1.02	1.03	1.03	1.01
spectral	6.96	1	1	1	1	1	1	1	1	1	1	1	1
pca	3.29	1	1	1	1.01	1	1	1	1.01	1	1	1	1
liblinear	25.75	1	1	1	1	1	1	1	1	1	1	1	1
kmeans	33.74	1	1	1	1	1	1	1	1	1	1	1	1
Geometric mean		1.03	1.06	1.34	1.44	1.09	1.2	1.88	1.87	1.01	1.03	1.17	1.31

TABLE IV  
 LOAD-TIME STATISTICS FROM LIBKPCAC PATCHING ROUTINE

Benchmark	Text section size [KiB]	Patched locations	Patching time [ $\mu$ s]	
			<i>svc</i>	<i>kpacd/kpacpl</i>
lda	9.11	53/54	455	640
sphinx	284.65	1529/1555	7928	8275
rbm	3.46	18/26	288	469
srr	12.45	30/39	490	654
svd3	25.15	145/145	872	1049
motion-est.	4.01	20/22	300	491
spectral	5.71	17/18	377	559
pca	5.40	18/18	293	496
liblinear	37.90	119/128	1060	1262
kmeans	2.09	6/8	286	473
Average	38.99	97.12%	1235	1437

651 *Load-Time Instrumentation:* Next, we examine the binary-  
 652 compatible approach based on the load-time patching using  
 653 *libkpac*. In terms of the security, load-time patching is tan-  
 654 tamount to the protection scope *all*, as GCC hardens all the  
 655 nonleaf functions with PACIASP/AUTIASP instructions.

656 The average overheads are 1.44 and  $1.31\times$  for *kpacpl* and  
 657 *kpacd*, respectively. The worst overhead can be observed in  
 658 *svc*-only mode (*svc-libkpac*), where the load-time instrumen-  
 659 tation yields  $1.87\times$  overhead on average. The is slightly better  
 660 than the respective static approach (*svc-all* with  $1.88\times$ ) as  
 661 the dynamic instrumentation (unlike the compiler plugin) does  
 662 not require disabling any optimizations. Here, *kpacd-libkpac*  
 663 represents the middle ground between *kpacd-all* and *svc-all*,  
 664 since *libkpac* replaces PACIASP/AUTIASP conservatively  
 665 with a call to the optimized *kpacd* routine, resorting to  
 666 costly *svc* where no familiar prologue/epilogue sequences are  
 667 recognized (due to the instruction reordering).

668 Table IV provides the additional statistics on the load-  
 669 time patching. Overall, *libkpac* in the *kpacd/kpacpl* modes  
 670 manages to successfully patch 97.12% of prologues and  
 671 epilogues in CortexSuite. The time required for patching does  
 672 not exceed 10 ms for any of the benchmarks and correlates  
 673 roughly with the size of the executable segment. The required  
 674 average time per KiB is 32  $\mu$ s for *svc*-only and 37  $\mu$ s for  
 675 *kpacd/kpacpl*.

### C. Case Study: Memcached

676

677 Despite the fact that *kpacd* requests do not require  
 678 synchronization with the other threads, there is a risk of  
 679 high contention on a single service core when serving the  
 680 multiple application cores. To accommodate highly parallel  
 681 applications, the KPAC kernel allows flexibly configuring the  
 682 amount of *kpacd* service cores and the mapping to the  
 683 application cores that they serve. The synthetic benchmarks  
 684 from the CortexSuite are single-threaded and do not assess the  
 685 multithreading aspect of KPAC. Hence, we deploy Memcached  
 686 1.6.22 on a Gigabyte R152-P31 rack server featuring the  
 687 Ampere Altra Q80-30 at 3 GHz with 80 Neoverse-N1 cores  
 688 and 256 GiB of DRAM. The machine offers uniform memory  
 689 access (UMA) latencies for all the cores.

690 We chose Memcached for several reasons. The code base  
 691 is written in plain C, making it easy to deploy it with custom  
 692 CFI techniques like our software-emulated PA. Furthermore,  
 693 benchmarking tools are readily available. Also, Memcached  
 694 is a realistic use-case for the PA as it is used in security  
 695 relevant environments, for example in combination with an  
 696 LDAP service for the user authentication.

697 *Workload:* Memcached server is compiled with the default  
 698 compiler flags, including the  $-O2$  optimization level. It is  
 699 complemented with PA-based return address protection and  
 700 ran with 32 threads pinned to 32 CPU cores. For the client side,  
 701 we use the *memtier\_benchmark* [30], which is developed by  
 702 Redis specifically for benchmarking the key-value databases.  
 703 The Memtier benchmark starts 32 threads on another 32 cores  
 704 of the same machine. Each threads opens 50 connections  
 705 (resulting in 1600 active connections) and records latencies  
 706 of SET and GET requests with the default SET:GET ratio of  
 707 1:10 for 100 s. We vary the amount of *kpacd* service threads  
 708 on the remaining 16 cores of the machine.

709 *Results:* Fig. 5 displays the average latency in milliseconds  
 710 as well as the 99th-percentile tail latency for the baseline  
 711 reference without PA, *svc*-, and *kpacd*-instrumented runs  
 712 (statically via the compiler plugin and load time via *libkpac*).  
 713 For *kpacd*, the latencies are measured for different amounts  
 714 of the service cores (*kpacd-x*).

715 The baseline average and tail latencies are 1.09 and 1.66 ms,  
 716 respectively. When using one service core there is a high  
 717 latency for all the protection scopes except for *char*. In fact,



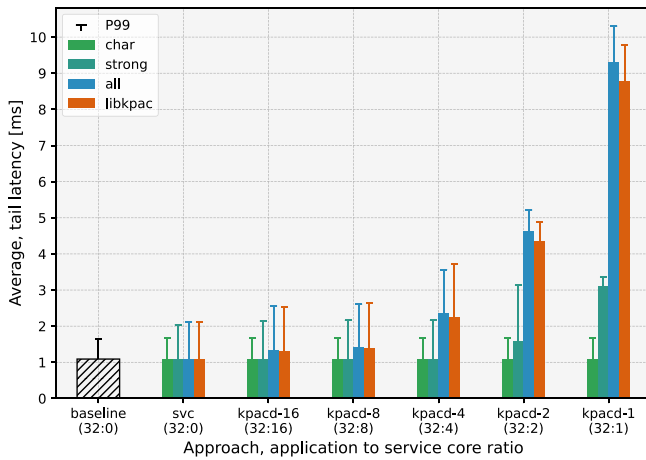


Fig. 5. Average latencies and 99th percentiles for the Memtier benchmark with 0–16 service cores (*svc* and *kpacd-x*).

718 due to the low amount of protected functions, *char* shows  
 719 no measurable change for all the approaches both in average  
 720 and tail latencies. For all, the average latency increases almost  
 721 tenfold to 9.30 ms. This figure halves as we double the amount  
 722 of service cores: it amounts to 4.62 ms (4.24 $\times$ ) for the two  
 723 service cores and 2.37 ms (2.17 $\times$ ) for the four service cores.  
 724 Distributing the load over eight *kpacd* threads and above,  
 725 they are no longer saturated, and the latency does not exceed  
 726 1.43 ms (1.31 $\times$  of the baseline). The figures are similar to  
 727 all for *libkpac* instrumented *kpacd* experiments as *libkpac*  
 728 manages to patch 91.78% of locations with a *kpacd* invocation.  
 729 Looking at the *strong* protection scope, the average latency  
 730 increase is low for the four service cores and above. For two  
 731 service cores the increase is 1.45 $\times$  or 1.58 ms.

732 Interestingly, *svc*-instrumented servers (including *all* and  
 733 *libkpac* variants) demonstrate the same average latency as  
 734 the baseline run. The tail latency, however, shows a minor  
 735 increase of 28.52% for *libkpac* and 23.20% for *strong*.  
 736 This stems from the architecture of the used machine. As  
 737 demonstrated in Section VI-A, the Ampere Altra Q80 CPU  
 738 features particularly the fast system calls. Combined with  
 739 the fact that the *svc* approach does not induce contention  
 740 in multithreaded scenarios, this results in fast return-address  
 741 protection. This highlights that the underlying hardware and  
 742 architecture needs to be taken into account when applying  
 743 the mechanism. In this case, the *svc-libkpac* approach can  
 744 be easily applied to an already compiled Memcached server  
 745 (e.g., from the distribution’s repository) without significantly  
 746 affecting the performance of the database. This comes at a  
 747 cost of relatively short 1.55 ms patching time for the 147 KiB  
 748 executable segment of the Memcached binary and all the  
 749 libraries it links with.

#### 750 D. Case Study: Chromium

751 Investigating the binary-compatible approach further, we  
 752 concentrate on its ease of use with the already existing  
 753 software and toolchains. We demonstrate this by applying *svc*-  
 754 and *kpacd*-based PA using *libkpac* to the *Chromium browser*.  
 755 For this we use the Raspberry Pi 4 single-board computer

756 featuring a quad-core Cortex-A72 64-bit SoC clocked at  
 757 1.8 GHz. We chose this system, as it represents a small  
 758 lightweight ARM desktop PC.

759 The *Chromium browser* is a long-existing project with a  
 760 large code base, leading the browser market with the usage  
 761 share of 63% on all the platforms (September 2023) [31].  
 762 Moreover, the binary Chromium package of the AArch64  
 763 Debian distribution is already hardened with the ARMv8.3-A  
 764 return address protection using PACIASP/AUTIASP instruc-  
 765 tions. However, this protection has no effect on systems  
 766 without PA. This makes Chromium a prime target for our  
 767 evaluation, as we want to showcase the ease of use and the low  
 768 adoption hurdle for the end-users. Security in Web browsers  
 769 is highly relevant in general, as the users use them for the  
 770 online banking, healthcare information, and a wide range of  
 771 other sensitive tasks. The Chromium project itself states that  
 772 around 70% of their security bugs are the memory safety  
 773 problems [32]. The severity of those bugs would be alleviated  
 774 by enabling PA.

775 *Workload:* To stress test our mechanism and give an  
 776 intuition on how usable the binary-compatible approaches  
 777 are for the user-oriented applications, we execute the  
 778 *Speedometer 2.1*, *Jetstream 2.1*, and *Motionmark 1.2* browser  
 779 benchmarks from the WebKit’s Browserbench suite [33].  
 780 Speedometer emulates user input by adding, changing and  
 781 removing to-do items from a Web application, evaluating the  
 782 browser’s responsiveness. Jetstream, on the other hand consists  
 783 of the Web assembly and JavaScript benchmarks (64 in total),  
 784 which are then scored using the geometric mean. These  
 785 benchmarks consist of several cryptography algorithms, data  
 786 processing tasks, parsers, and so on. The third Browserbench  
 787 benchmark, Motionmark, puts the browser’s graphics engine  
 788 to the test by animating complex scenes.

789 *Results:* Our library manages to patch 95:85% of prologues  
 790 and epilogues of the *Chromium browser* binary with the  
 791 optimized *kpacd* invocations. As this binary is quite large  
 792 (161.02 MiB executable segment), we need 549.89 ms to patch  
 793 it. This corresponds to the rate of 3.42 ms/MiB. Extending the  
 794 patching onto the libraries that link with *Chromium browser*,  
 795 the patching takes 571.13 ms. Out of those libraries, only  
 796 *libffmpeg.so* (part of the Chromium package) and *libgnutls.so*  
 797 (distribution’s version) are compiled with ARMv8.3-A PA. All  
 798 the experiments ran without any changes to the source code  
 799 and without any crashes or errors.

800 Fig. 6 displays the reached scores for both the bench-  
 801 marks. The optimized *kpacd*-based instrumentation reduces  
 802 the scores by a factor of 4.43 $\times$  for the speedometer benchmark  
 803 and by a factor of 3.54 $\times$  for the Jetstream benchmark. On  
 804 the other hand, the *svc* approach reduces the scores by a  
 805 factor of 10.63 and 7.13 $\times$  for the speedometer and Jetstream,  
 806 respectively. The Motionmark benchmark shows only minor  
 807 difference between the three browser variants, with a worst-  
 808 case score reduction of 14.05% for *svc*.

809 The results are consistent with the transaction latencies  
 810 measured for this system in Section VI-A, where performing a  
 811 system call calculating SipHash PAC has the quadruple latency  
 812 of an RCSC transaction. The high overall overhead can be  
 813 attributed to the fact that the browsers spend significant amount




Fig. 6. Scores as reported by the Browserbench benchmarks for *libkpac*-instrumented browser. Higher scores are better.

814 of their time interpreting the JavaScript code. If one of the  
 815 interpreter’s hot functions is nonleaf and thus authenticates  
 816 its return address, this results in a high performance impact  
 817 when compared to the CortexSuite and Memcached figures.  
 818 However, even with *svc-libkpac*, the browser remains usable  
 819 and responsive, which suggests restricting this protection  
 820 technique for the security-critical applications.

## 821 VII. DISCUSSION

822 *General Applicability:* Given that the vast majority of  
 823 even recent ARMv8.3 designs do not yet include the PA  
 824 extensions, its efficient software-based emulation in the kernel  
 825 will probably be useful for several years but (hopefully)  
 826 eventually become obsolete. However, the four techniques and  
 827 their tradeoffs presented in this article for such emulation are  
 828 not restricted to PA. They could most probably be applied  
 829 also to the future security/safety-related ISA extensions. The  
 830 RCSC mechanism is furthermore usable for the easy offloading  
 831 of any kind of performance- or security-critical service to  
 832 a dedicated core. By its CPU-local page-tables, it provides  
 833 seamless integration into multithreaded applications without  
 834 extra synchronization efforts.

835 *Hardware Costs:* Offloading kernel tasks to the dedicated  
 836 cores has been shown to be effective in improving the  
 837 performance in many settings [29], but to the best of our  
 838 knowledge not yet as an alternative to a relatively simple  
 839 FPGA-based solution. We consider this as a question  
 840 of pragmatics. Technically, (i.e., with respect to the HW  
 841 overhead), the FPGA-based solution is undoubtedly a lot  
 842 cheaper. However, actual availability and market prices often  
 843 tell a different story. While SoCs, including an FPGAs are  
 844 still a development niche, multicore CPUs are prevalent on  
 845 the market and benefit from competitive pricing and mass  
 846 production for procurement, the SoC including an extra core  
 847 is often cheaper than the one with the FPGA. Besides, these  
 848 multicore CPUs are rarely utilized to their full potential due  
 849 to the limited parallelism within the software. This warrants  
 850 considering dedicating one or multiple cores to a service like  
 851 KPAC for increased security or performance or employing  
 852 them instead of an FPGA accelerator. In the end, the question  
 853 of spending an extra core or not comes down to the actual  
 854 performance cost tradeoff, as developers can always opt for  
 855 one of the synchronous emulation variants.

## VIII. RELATED WORK

856 *Software-Based Pointer Protection:* Before ARMv8.3-A PA,  
 857 the idea of adding a cryptographic MAC to the code pointers  
 858 has been explored in a technique called *CCFI* [34]. To keep  
 859 the key secret, CCFI reserves 11 XMM registers on x86-  
 860 64, which constitutes a change to the ABI, requiring the  
 861 recompilation of the program and all its dependencies. Its  
 862 predecessor, *PointGuard* [35] introduces a compiler extension  
 863 that instruments programs to encrypt the pointers when storing  
 864 them into memory using simple XOR with a key stored in  
 865 the same address space. Another approach, called *CPI* [36],  
 866 protects pointers by storing them in a secret location along  
 867 with metadata. However, Evans [37] demonstrated an attack  
 868 that is able to bypass CPI and argued that the security  
 869 mechanisms relying on the information hiding are ineffective.

870 Compared to these approaches, KPAC maintains higher  
 871 security guarantees by computing the cryptographic signature  
 872 in the kernel space, which allows it to reliably hide the secret  
 873 key from the attackers.  
 874

875 *Applications of PACs:* The ARM PA mechanism is not  
 876 limited to the return address protection. In the recent years,  
 877 many CFI mechanisms employing PA codes in the user space  
 878 have been proposed. Liljestrand et al. [24] have presented  
 879 several works on this subject. *PARTS* [24] is an instrumentation  
 880 framework, which extends the set of protected pointers to the  
 881 local, global, and static pointers as well as the pointers in C  
 882 structures. *PCan* [38] revisits the concept of stack canaries by  
 883 dynamically generating their value for each function call using  
 884 PA instructions, eliminating the need to hide their value in  
 885 memory. *PACStack* [39] upgrades the return address protection  
 886 by cryptographically binding its value to all the previous  
 887 return addresses in the call stack, preventing the pointer reuse  
 888 attacks. *PTAuth* [40], *PACMem* [41], and *CryptSan* [42] are  
 889 the sanitizers that detect the spatial and temporal memory  
 890 bugs by leveraging PACs. Schilling, Nasahl, and colleagues  
 891 utilize PACs to thwart not only the software but also the fault  
 892 attacks by 1) ensuring CFI at the basic block granularity [43]  
 893 and 2) protecting indirect branches by encoding them at  
 894 compile time and verifying them at run time [44]. The work  
 895 of Fanti et al. [45] generalizes PA by protecting not only the  
 896 pointers but all the spilled registers.

897 Given the scarcity of systems with hardware PA, many  
 898 of these works have resorted to emulating PA instructions  
 899 using a rudimentary XOR “encryption” as a proof-of-concept.  
 900 With KPAC, all these PA-based techniques could be seam-  
 901 lessly integrated with our cryptographically secure approaches,  
 902 extending the CFI guarantees beyond the backward edge  
 903 protection for the systems without hardware-assisted PA.

904 *Dedicated Service Cores:* Several other works have explored  
 905 the possibility of dedicating CPU cores of the system for  
 906 some specific service in order to avoid the context switching  
 907 overhead. For example, Lozi et al. [29] replaced lock acquisi-  
 908 tions with remote calls to a dedicated core executing a critical  
 909 section and observe the performance increase, attributing it  
 910 to the data locality. The technique of offloading network  
 911 packet processing to a separate CPU core has been repeatedly  
 912 proposed since the early days of consumer grade multicore

913 CPUs [46], [47]. *IsoStack* [48] and *Shenago* [49] imple-  
 914 ment that kind of network stack and demonstrates significant  
 915 performance improvements. A similar technique has also been  
 916 successfully applied to speed up virtualization [50], [51].

917 The novelty of our approach lies in the idea of modifying  
 918 the virtual memory layer to present each application thread  
 919 with the page private to the CPU core it is running on. This  
 920 forms a framework for the secure communication with the  
 921 dedicated service core without requiring any synchronization.

## 922 IX. CONCLUSION

923 ARMv8.3-A pointer authentication is a promising CFI  
 924 mechanism, which is expected to gain more traction in the fol-  
 925 lowing years. It provides significant security gains for minimal  
 926 performance impact, owing to its hardware implementation.  
 927 However, CPUs implementing this feature are still rare and  
 928 we face many systems without PA support.

929 In this work, we explore how PA can be emulated in  
 930 software, while maintaining the low performance overhead.  
 931 For this, we extended the Linux kernel with a PA service  
 932 and exposed two communication interfaces for the user appli-  
 933 cations: 1) the classical synchronous system call and 2) a  
 934 shared memory page for the asynchronous communication.  
 935 We also investigated two instrumentation methods for the  
 936 existing applications: 1) statically, by recompiling them with  
 937 our compiler plugin and (2) in the ARMv8.3-ABI-compatible  
 938 way, by patching them at load time. In the static case, we  
 939 employ several heuristics inspired by GCC's stack protector  
 940 feature [13] to limit protection to the vulnerable functions,  
 941 offering a flexible balance between the performance and  
 942 security.

943 We combined all the aspects into a total of eight approaches  
 944 and evaluated their run-time impact using the CortexSuite  
 945 benchmarks and the Memcached key-value database. We also  
 946 assessed the ease of use in the end-user environments by  
 947 applying our approaches to the *Chromium browser* without  
 948 recompilation. For the best of our approaches, we observed  
 949 low overheads: a worst-case run duration increase of 20% for  
 950 the CortexSuite benchmarks when using *kpacd-strong* and a  
 951 modest 29% increase in tail latency for the Memcached with  
 952 *svc-libkpac*.

953 The source code and evaluation artifacts are available at:  
 954 <https://github.com/luhsra/kpac>

## 955 ACKNOWLEDGMENT

956 The authors thank the anonymous reviewers for their feed-  
 957 back and fruitful comments.

## 958 REFERENCES

959 [1] Intel® 64 and IA-32 Architectures Software Developer's Manual,  
 960 Combined Volumes: 1–4, Intel, Santa Clara, CA, USA, Apr. 2022.  
 961 [2] T. Garrison (Intel, Santa Clara, CA, USA). *Intel CET Answers Call to*  
 962 *Protect Against Common Malware Threats*. 2020, [Online]. Available:  
 963 [https://www.intel.com/content/www/us/en/newsroom/opinion/intel-cet-](https://www.intel.com/content/www/us/en/newsroom/opinion/intel-cet-answers-call-protect-common-malware-threats.html)  
 964 [answers-call-protect-common-malware-threats.html](https://www.intel.com/content/www/us/en/newsroom/opinion/intel-cet-answers-call-protect-common-malware-threats.html).  
 965 [3] V. Shanbhogue, D. Gupta, and R. Sahita, "Security analysis of processor  
 966 instruction set architecture for enforcing control-flow integrity," in *Proc.*  
 967 *8th Int. Workshop Hardw. Archit. Support Security Privacy*, 2019,  
 968 pp. 1–11, doi: 10.1145/3337167.3337175.  
 969 [4] *Arm Architecture Reference Manual for A-Profile Architecture*,  
 970 Arm Limited, Cambridge, U.K., 2022.

[5] R. Roemer, E. Buchanan, H. Shacham, and S. Savage, "Return- 971  
 oriented programming: Systems, languages, and applications," *ACM* 972  
*Trans. Inf. Syst. Security*, vol. 15, no. 1, pp. 1–34, Mar. 2012, 973  
 doi: 10.1145/2133375.2133377. 974

[6] H. Shacham, "The geometry of innocent flesh on the bone: 975  
 Return-into-Libc without function calls," in *Proc. 14th ACM Conf.* 976  
*Comput. Commun. Security*, New York, NY, USA, 2007, pp. 552–561, 977  
 doi: 10.1145/1315245.1315313. 978

[7] T. Bletsch, X. Jiang, V. W. Freeh, and Z. Liang, "Jump-oriented 979  
 programming: A new class of code-reuse attack," in *Proc. 6th ACM* 980  
*Symp. Inf. Comput. Commun. Security*, New York, NY, USA, 2011, 981  
 pp. 30–40, doi: 10.1145/1966913.1966919. 982

[8] T. H. Dang, P. Maniatis, and D. Wagner, "The performance cost 983  
 of shadow stacks and stack canaries," in *Proc. 10th ACM Symp.* 984  
*Inf. Comput. Commun. Security*, Singapore, 2015, pp. 555–566, 985  
 doi: 10.1145/2714576.2714635. 986

[9] C. Zou, X. Wang, Y. Gao, and J. Xue, "Buddy stacks: Protecting 987  
 return addresses with efficient thread-local storage and runtime re- 988  
 randomization," *ACM Trans. Softw. Eng. Methodol.*, vol. 31, no. 2, 989  
 pp. 1–37, Apr. 2022, doi: 10.1145/3494516. 990

[10] C. Zou, Y. Gao, and J. Xue, "Practical software-based shadow stacks 991  
 on x86-64," *ACM Trans. Archit. Code Optim.*, vol. 19, no. 4, pp. 1–26, 992  
 Dec. 2022, doi: 10.1145/3556977. 993

[11] M. Abadi, M. Budiu, U. Erlingsson, and J. Ligatti, "Control-flow 994  
 integrity: Principles, implementations, and applications," in *Proc. 12th* 995  
*ACM Conf. Comput. Commun. Security*, New York, NY, USA, 2005, 996  
 pp. 340–353, doi: 10.1145/1102120.1102165. 997

[12] N. Burow, X. Zhang, and M. Payer, "SoK: Shining light on shadow 998  
 stacks," in *Proc. IEEE Symp. Security Privacy*, San Francisco, CA, USA, 999  
 2019, pp. 985–999, doi: 10.1109/SP.2019.00076. 1000

[13] (GNU project, Boston, MA, USA). *Using the GNU Compiler* 1001  
*Collection (GCC), Version 12.1*. 2022. [Online]. Available: 1002  
<https://gcc.gnu.org/onlinedocs/gcc-12.1.0/gcc/> 1003

[14] "The LLVM project clang 14.0.0 documentation." 2022. [Online]. 1004  
 Available: <https://releases.llvm.org/14.0.0/tools/clang/docs/> 1005

[15] J. Corbet. "The rest of the 6.6 merge window." Sep. 2023. [Online]. 1006  
 Available: <https://lwn.net/Articles/943245/> 1007

[16] (Microsoft, Redmond, WA, USA). *MWC 2022: The Next Microsoft* 1008  
*Pluton Device + PAC Technology*. 2022. [Online]. Available: 1009  
[https://blogs.windows.com/windowsexperience/2022/02/28/mwc-2022-](https://blogs.windows.com/windowsexperience/2022/02/28/mwc-2022-the-next-microsoft-pluton-device-pac-technology/) 1010  
[the-next-microsoft-pluton-device-pac-technology/](https://blogs.windows.com/windowsexperience/2022/02/28/mwc-2022-the-next-microsoft-pluton-device-pac-technology/) 1011

[17] G. Serra, P. Fara, G. Cicero, F. Restuccia, and A. Biondi, "PAC-PL: 1012  
 Enabling control-flow integrity with pointer authentication in FPGA SoC 1013  
 platforms," in *Proc. 28th IEEE Real-Time Embed. Technol. Appl. Symp.*, 1014  
 2022, pp. 241–253, doi: 10.1109/RTAS54340.2022.00027. 1015

[18] R. Avanzi et al., "The tweakable block cipher family QARMAv2," 1016  
*Cryptol. ePrint Arch., IACR, Bellevue, WA, USA, Rep. 2023/929*, 2023. 1017  
 [Online]. Available: <https://eprint.iacr.org/2023/929> 1018

[19] "Linux 5.0 changelog." 2018. [Online]. Available: [https://cdn.kernel.org/](https://cdn.kernel.org/pub/linux/kernel/v5.x/ChangeLog-5.0) 1019  
[pub/linux/kernel/v5.x/ChangeLog-5.0](https://cdn.kernel.org/pub/linux/kernel/v5.x/ChangeLog-5.0) 1020

[20] "ARMv8.3 pointer authentication in xnu." 2021. [Online]. Available: 1021  
[https://opensource.apple.com/source/xnu/xnu-7195.50.7.100.1/doc/pac.](https://opensource.apple.com/source/xnu/xnu-7195.50.7.100.1/doc/pac.md) 1022  
[md](https://opensource.apple.com/source/xnu/xnu-7195.50.7.100.1/doc/pac.md) 1023

[21] "Linux 5.7 changelog." 2020. [Online]. Available: [https://cdn.kernel.org/](https://cdn.kernel.org/pub/linux/kernel/v5.x/ChangeLog-5.7) 1024  
[pub/linux/kernel/v5.x/ChangeLog-5.7](https://cdn.kernel.org/pub/linux/kernel/v5.x/ChangeLog-5.7) 1025

[22] "Apple pointer authentication guidelines." 2023. [Online]. Available: 1026  
[https://developer.apple.com/documentation/security/preparing\\_your\\_](https://developer.apple.com/documentation/security/preparing_your_app_to_work_with_pointer_authentication) 1027  
[app\\_to\\_work\\_with\\_pointer\\_authentication](https://developer.apple.com/documentation/security/preparing_your_app_to_work_with_pointer_authentication) 1028

[23] S. Thomas et al., "CortexSuite: A synthetic brain benchmark suite," in 1029  
*Proc. IEEE Intl. Symp. Workload Characterization*, 2014, pp. 76–79, 1030  
 doi: 10.1109/IISWC.2014.6983043. 1031

[24] H. Liljestrand, T. Nyman, K. Wang, C. C. Perez, J.-E. Ekberg, and 1032  
 N. Asokan, "PAC it up: Towards pointer integrity using ARM pointer 1033  
 authentication," in *Proc. 28th USENIX Conf. Security Symp.*, 2019, 1034  
 pp. 177–194. 1035

[25] Y. Wang, J. Wu, T. Yue, Z. Ning, and F. Zhang, "RetTag: 1036  
 Hardware-assisted return address integrity on RISC-V," in *Proc. 15th* 1037  
*Eur. Work. Syst. Security*, New York, NY, USA, 2022, pp. 50–56, 1038  
 doi: 10.1145/3517208.3523758. 1039

[26] J.-P. Aumasson and D. J. Bernstein, "SipHash: A fast short-input PRF," 1040  
 in *Proc. Int. Conf. Cryptol.*, 2012, pp. 489–508. 1041

[27] H. Härtig, M. Hohmuth, J. Liedtke, S. Schönberg, and J. Wolter, 1042  
 "The performance of  $\mu$ -kernel-based systems," in *Proc. 16th ACM* 1043  
*Symp. Oper. Syst. Princ.*, New York, NY, USA, 1997, pp. 66–77, 1044  
 doi: 10.1145/269005.266660. 1045

[28] D. Wentzlaff and A. Agarwal, "Factored operating systems (fos): The 1046  
 case for a scalable operating system for multicores," *ACM SIGOPS Oper.* 1047  
*Syst. Rev.*, vol. 43, pp. 76–85, Apr. 2009. doi: 10.1145/1531793.1531805. 1048

- 1049 [29] J.-P. Lozi, F. David, G. Thomas, J. Lawall, and G. Muller, "Remote core  
1050 locking: Migrating critical-section execution to improve the performance  
1051 of Multithreaded applications," in *Proc. USENIX Annu. Tech. Conf.*,  
1052 2012, p. 6.
- 1053 [30] "Memtier benchmark on Github." RedisLabs. 2024. [Online]. Available:  
1054 [https://github.com/RedisLabs/memtier\\_benchmark](https://github.com/RedisLabs/memtier_benchmark)
- 1055 [31] (StatCounter, Dublin, Ireland). *Browser Market Share Worldwide*. 2023.  
1056 [Online]. Available: <https://gs.statcounter.com/browser-market-share>
- 1057 [32] "The chromium project." 2023. [Online]. Available: [https://www.  
1058 chromium.org/Home/chromium-security/memory-safety/](https://www.chromium.org/Home/chromium-security/memory-safety/)
- 1059 [33] "WebKit's browserbenchmarks." 2023. [Online]. Available: [https://  
1061 browserbench.org](https://<br/>1060 browserbench.org)
- 1061 [34] A. J. Mashtizadeh, A. Bittau, D. Boneh, and D. Mazières, "CCFI:  
1062 Cryptographically enforced control flow integrity," in *Proc. 22nd ACM  
1063 SIGSAC Conf. Comput. Commun. Security*, New York, NY, USA, 2015,  
1064 pp. 941–951, doi: [10.1145/2810103.2813676](https://doi.org/10.1145/2810103.2813676).
- 1065 [35] C. Cowan, S. Beattie, J. Johansen, and P. Wagle, "PointGuard™:  
1066 Protecting pointers from buffer overflow vulnerabilities," in *Proc. 12th  
1067 USENIX Security Symp.*, 2003, pp. 90–104.
- 1068 [36] V. Kuznetsov, L. Szekeres, M. Payer, G. Candea, R. Sekar, and D. Song,  
1069 "Code-pointer integrity," in *Proc. 11th USENIX Symp. Oper. Syst. Design  
1070 Implement.*, Broomfield, CO, USA, 2014, pp. 147–163.
- 1071 [37] I. Evans et al., "Missing the point(er): On the effectiveness of  
1072 code pointer integrity," in *Proc. IEEE Symp. Security Privacy*, 2015,  
1073 pp. 781–796, doi: [10.1109/SP.2015.53](https://doi.org/10.1109/SP.2015.53).
- 1074 [38] H. Liljestrand, Z. Gauhar, T. Nyman, J.-E. Ekberg, and N. Asokan,  
1075 "Protecting the stack with PACed canaries," in *Proc. 4th Work.  
1076 System Softw. Trusted Execution*, New York, NY, USA, 2019, pp. 1–6,  
1077 doi: [10.1145/3342559.3365336](https://doi.org/10.1145/3342559.3365336).
- 1078 [39] H. Liljestrand, T. Nyman, L. J. Gunn, J.-E. Ekberg, and N. Asokan,  
1079 "PACStack: An authenticated call stack," in *Proc. 30th USENIX Security  
1080 Symp.*, 2021, pp. 357–374.
- 1081 [40] R. M. Farkhani, M. Ahmadi, and L. Lu, "PTAuth: Temporal memory  
1082 safety via robust points-to authentication," in *Proc. 30th USENIX  
1083 Security Symp.*, 2021, pp. 1037–1054.
- 1084 [41] Y. Li et al., "PACMem: Enforcing spatial and temporal  
1085 memory safety via ARM pointer authentication," in *Proc. ACM  
1086 SIGSAC Conf. Comput. Commun. Security*, 2022, pp. 1901–1915,  
1087 doi: [10.1145/3548606.3560598](https://doi.org/10.1145/3548606.3560598).
- [42] K. Hohentanner, P. Zieris, and J. Horsch, "CryptSan: Leveraging  
1088 ARM pointer authentication for memory safety in C/C++," in  
1089 *Proc. 38th ACM/SIGAPP Symp. Appl. Comput.*, 2023, pp. 1530–1539,  
1090 doi: [10.1145/3555776.3577635](https://doi.org/10.1145/3555776.3577635). 1091
- [43] R. Schilling, P. Nasahl, and S. Mangard, "FIPAC: Thwarting fault-and  
1092 software-induced control-flow attacks with ARM pointer authentication,"  
1093 in *Proc. 13th Int. Work. Constr. Side Channel Anal. Secure Design*, 2022,  
1094 pp. 100–124, doi: [10.1007/978-3-030-99766-3\\_5](https://doi.org/10.1007/978-3-030-99766-3_5). 1095
- [44] P. Nasahl, R. Schilling, and S. Mangard, "Protecting indirect  
1096 branches against fault attacks using ARM pointer authentication,"  
1097 in *Proc. IEEE Int. Symp. Hardw. Oriented Security Trust*, 2021,  
1098 pp. 68–79, doi: [10.1109/HOST49136.2021.9702268](https://doi.org/10.1109/HOST49136.2021.9702268). 1099
- [45] A. Fanti, C. C. Perez, R. Denis-Courmont, G. Roascio, and J. Ekberg,  
1100 "Toward register spilling security using LLVM and ARM pointer authen-  
1101 tication," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 41,  
1102 no. 11, pp. 3757–3766, Nov. 2022, doi: [10.1109/TCAD.2022.3197511](https://doi.org/10.1109/TCAD.2022.3197511). 1103
- [46] M. Rangarajan, A. Bohra, K. Banerjee, E. V. Carrera, R. Bianchini, and  
1104 L. Iftode, "TCP servers: Offloading TCP processing in Internet servers.  
1105 Design, implementation, and performance," Dept. Arts Sci. Comput.  
1106 Sci., Rutgers Univ., New Brunswick, NJ, USA, Rep. DCS-TR-481,  
1107 2002. 1108
- [47] T. Brecht, G. J. Janakiraman, B. Lynn, V. A. Saletore, and  
1109 Y. Turner, "Evaluating network processing efficiency with processor  
1110 partitioning and asynchronous I/O," in *Proc. EuroSys Conf.*, 2006,  
1111 pp. 265–278, doi: [10.1145/1217935.1217961](https://doi.org/10.1145/1217935.1217961). 1112
- [48] L. Shalev, J. Satran, E. Borovik, and M. Ben-Yehuda, "IsoStack—highly  
1113 efficient network processing on dedicated cores," in *Proc. USENIX Annu.  
1114 Tech. Conf.*, 2010, pp. 1–14, doi: [10.5555/1855840.1855845](https://doi.org/10.5555/1855840.1855845). 1115
- [49] A. Ousterhout, J. Fried, J. Behrens, A. Belay, and  
1116 H. Balakrishnan, "Shenango: Achieving high CPU efficiency  
1117 for latency-sensitive datacenter workloads," in *Proc. 16th  
1118 USENIX Symp. Netw. Syst. Design Implement.*, 2019,  
1119 pp. 361–378. 1120
- [50] J. Liu and B. Abali, "Virtualization Polling Engine (VPE): Using dedi-  
1121 cated CPU cores to accelerate I/O virtualization," in *Proc. 23rd Int. Conf.  
1122 Supercomput.*, 2009, pp. 225–234, doi: [10.1145/1542275.1542309](https://doi.org/10.1145/1542275.1542309). 1123
- [51] A. Landau, M. Ben-Yehuda, and A. Gordon, "SplitX: Split  
1124 guest/Hypervisor execution on multi-core," in *Proc. 3rd Work I/O  
1125 Virtualization*, 2011, pp. 1–7. 1126