

1 A Memory Scheduling Infrastructure for 2 Multi-core Systems with Re-programmable Logic

3 Denis Hoornaert ✉

4 Technical University of Munich, Germany

5 Shahin Roozkhosh ✉

6 Boston University, USA

7 Renato Mancuso ✉

8 Boston University, USA

9 — Abstract —

10 The sharp increase in demand for performance has prompted an explosion in the complexity of
11 modern multi-core embedded systems. This has led to unprecedented temporal unpredictability
12 concerns in Cyber-Physical Systems (CPS). On-chip integration of programmable logic (PL) alongside
13 a conventional Processing Systems (PS) in modern Systems-on-Chip (SoC) establishes a genuine
14 compromise between specialization, performance, and re-configurability. In addition to typical
15 use-cases, it has been shown that the PL can be used to observe, manipulate, and ultimately manage
16 memory traffic generated by a traditional multi-core processor.

17 This paper explores the possibility of PL-aided memory scheduling by proposing a Scheduler In-
18 the-Middle (SchIM). We demonstrate that the SchIM enables transaction-level control over the main
19 memory traffic generated by a set of embedded cores. Focusing on extensibility and reconfigurability,
20 we put forward a SchIM design covering two main objectives. First, to provide a safe playground
21 to test innovative memory scheduling mechanisms; and second, to establish a transition path from
22 software-based memory regulation to provably correct hardware-enforced memory scheduling. We
23 evaluate our design through a full-system implementation on a commercial PS-PL platform using
24 synthetic and real-world benchmarks.

25 **2012 ACM Subject Classification** Computer systems organization → Real-time system architecture

26 **Keywords and phrases** MPSoC, FPGA, Memory Scheduling

27 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

28 **Funding** *Denis Hoornaert*: Denis Hoornaert was supported by an Alexander von Humboldt Profes-
29 sorship endowed by the German Federal Ministry of Education and Research.

30 *Renato Mancuso*: The material presented in this paper is based upon work supported by the
31 National Science Foundation (NSF) under grant number CCF-2008799. Any opinions, findings, and
32 conclusions or recommendations expressed in this publication are those of the authors and do not
33 necessarily reflect the views of the NSF.

34 **1** Introduction

35 It is undeniable that the massive increase in expectation on the performance of next-generation
36 cyber-physical systems has deeply impacted the way we design modern embedded and real-
37 time systems. High-resolution, high-bandwidth sensors such as lidars, and depth cameras on
38 the one hand, and data-intensive processing workload such as machine-learning applications
39 on the other hand, have exacerbated the push for high-performance embedded platforms.
40 Following this performance *moving target*, chip manufactures have significantly scaled up
41 clock speeds, CPU count, and heterogeneity. For instance, the on-chip integration of powerful
42 graphic processing units (GPUs) has been the characterizing factor in the NVIDIA Tegra
43 series of embedded systems-on-a-chip (SoC).



© Omitted for review;
licensed under Creative Commons License CC-BY 4.0
42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:22



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

44 In this context, an embedded architectural paradigm that is surging in popularity among
45 manufacturers, researchers, and industry practitioners is the PS-PL organization. This
46 class of embedded platforms integrates on the same die (1) traditional full-speed embedded
47 CPUs and (2) programmable logic constructed using field-programmable gate array (FPGA)
48 technology. This organization naturally defines two macro-domains, namely the Processing
49 System (PS) and the Programmable Logic (PL), hence the name. PS-PL platforms establish a
50 good trade-off between specialization, raw performance, and mission-specific re-configurability.
51 The current generation of commercially available PS-PL platforms is dominated by ARM-
52 based products offered by, most notably, Intel [12] and Xilinx [38]. A pilot large-scale,
53 high-performance PS-PL system is the Enzian platform [3] being rolled out by ETH Zurich¹.
54 Furthermore, a RISC-V-based solution has been recently made available by Microsemi with
55 their PolarFire SoC [18].

56 From a real-time perspective, the co-existence of traditional CPUs and a tightly-coupled
57 block of PL has more profound implications than expected. Clearly, it is possible to define
58 custom accelerators in PL and to relieve the main CPUs of some of the heavy data-processing
59 workload. However, more interestingly, recent studies have highlighted the possibility of using
60 the PL also as a way to manage the memory traffic originated from the main CPUs [13, 29].
61 Such a possibility opens the doors for memory traffic inspection and control at the level
62 of individual transactions; which in turn promises to unlock provable determinism for the
63 real-time workload.

64 In this paper, we embrace the concept of PL-aided memory traffic management and propose
65 an infrastructure to develop, test and evaluate memory scheduling policies. Specifically, we
66 propose a component, called the Scheduler In-the-Middle—or SchIM, for short—that can
67 be instantiated in the PL to enforce a set of configurable scheduling policies on individual
68 memory transactions generated by the CPUs in the PS.

69 The overarching goal of the proposed SchIM is twofold. First, we want to provide a
70 playground for researches to test promising novel memory scheduling ideas for multi-core
71 platforms, much like LITMUS^{RT} [7] fostered research on CPU scheduling techniques. Second,
72 we want our SchIM to act as an intermediate stepping stone for industrial applications where
73 strong determinism over memory performance is required. The SchIM can be used to analyze
74 the behavior of realistic workload in a multitude of what-if memory management use-cases.
75 We note that such kind of analysis was previously possible only through full-system simulation
76 or by synthesizing the entire SoC on FPGA—that is, with a soft-core implementation.

77 In short, this paper makes the following contributions. (1) We demonstrate that a
78 configurable module could be interposed between the cores and the memory controller to
79 perform transaction-level scheduling in commercial PS-PL platforms; (2) we propose a
80 design for a memory scheduling infrastructure that focuses on extensibility and runtime
81 reconfigurability; (3) we address important issues to correctly account and regulate CPU-
82 generated traffic when a shared last-level cache is present; (4) we design and implement two
83 pilot memory scheduling policies as a proof-of-concept on the potential of our SchIM; and (5)
84 we perform a full system integration and implementation on a commercial PS-PL embedded
85 platform to evaluate the behavior of the SchIM with synthetic and realistic workload.

¹ Also see <http://enzian.systems/>

2 Related Work

There is a broad consensus that memory resources represent the main performance bottleneck in modern multi-core processors. The observation has sparked a host of research works addressing the problem from multiple angles [17]. In this context, the works representing the inspiration for our SchIM fall in two macro-categories, namely **hardware-based** and **software-based** techniques for main memory traffic management.

The first category includes a large body of works aimed at achieving better and/or more predictable performance by advancing novel hardware redesigns. The works in [22–24] strive to construct high-performance and fair memory schedulers. The addition of software-controlled memory deadlines and transactional semantics were explored in [33] and [10], respectively. Next, the work by Åkesson et al. [1, 2] and Paolieri et al. [25] attains timing predictability through careful scheduling of SDRAM commands. Finally, the MEDUSA DRAM controller [9, 34] implements a two-tiers scheduler at the DRAM controller to ensure predictability when accessing memory areas where access time strongly impact application performance. Finally, the hardware designs proposed in [8, 26, 43] put their emphasis on main memory bandwidth partitioning; clever dynamic pipelining is further explored in [20] to better balance average performance and determinism.

Among the software-based techniques are the mechanisms that stemmed from MemGuard, originally proposed in [42] and that rely on broadly available performance counters to regulate the bandwidth extracted by individual CPUs. Later extensions to jointly consider regulation and cache partitioning [39] and to expose control over memory bandwidth as a lockable resource [40] were proposed. Software-based memory throttling has also been implemented at the hypervisor-level [21, 30]. Remarkably, the work in [30] combines regulation mechanisms for CPU and embedded accelerators through the ARM QoS extensions [4].

In addition to the two categories surveyed above, perhaps the most closely related works are those that explored memory isolation techniques in PS-PL platforms. The work in [11] demonstrated that the PL-side can be used to define private memory storage, control, and bus units to strongly isolate high-criticality workload. A number of techniques developed as part of the FRED framework [6] put an emphasis on memory traffic arbitration and management for in-PL accelerators [27, 28]. The AXI HyperConnect [27] is perhaps the component most similar to the SchIM in terms of high-level design. However, both are substantially different as the SchIM is designed to manage embedded CPUs' memory traffic.

Compared to the literature reviewed above, what sets this work apart are the following aspects. (1) Our SchIM applies to existing PS-PL commercial systems without introducing any hardware modification; (2) it allows management in the PL of memory traffic originated by the embedded CPUs residing in the PS; (3) it provides the framework to test the feasibility and performance of custom memory scheduling policies; and (4) it is designed such that multiple schedulers can coexist, be activated, and configured at runtime.

3 Background Concepts

In this section, we introduce some fundamental concepts necessary to understand the overall system design and the class of platforms targeted by this work.

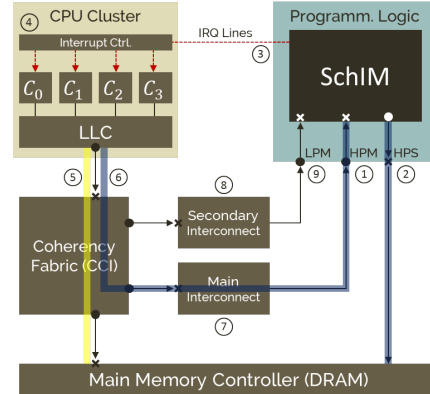
3.1 Hybrid Multi-Core Platforms with Programmable Logic

This work targets the aforementioned class of embedded multi-core platforms with programmable logic—i.e., PS-PL platforms. In such platforms, the PS encompasses a multi-core

130 processor with a multi-level cache hierarchy and a main memory (DRAM) controller. A
 131 simplified block diagram for a reference PS-PL organization is illustrated in Fig. 1. The
 132 figure considers a platform with four CPUs denoted as $C_0, C_1, C_2,$ and C_3 .

133 A key feature in PS-PL platforms is
 134 the presence of high-performance commu-
 135 nication channels between the two do-
 136 mains. These come in the form of
 137 data exchange interfaces and interrupt
 138 lines. Data exchange channels follow a
 139 master-slave paradigm. Specifically, high-
 140 performance masters (HPM, Fig. 1①) and
 141 high-performance slaves (HPS, Fig. 1②)
 142 send and receive transactions to and from
 143 the PL, respectively. Additionally, there ex-
 144 ist programmable interrupt request (IRQ)
 145 lines (see Fig. 1③) that can be driven by
 146 the PL and are connected to the interrupt
 147 controller (Fig. 1④) inside the PS. As we
 148 discuss in Section 5.7, the presence of PS-PL
 149 interrupt lines is crucial to building PL-assisted memory traffic regulation.

150 Note also that there might exist PS-PL data ports that are routed through a secondary
 151 interconnect (Fig. 1⑧). These can generally sustain less throughput compared to HPS ports;
 152 hence we refer to them as low-performance masters (LPM, Fig. 1⑨). LPM ports are useful
 153 to perform memory-mapped configuration of PL modules.




■ Figure 1 PS-PL interconnect block diagram.

154 3.2 Programmable Logic In-the-Middle

155 In this work, we leverage the ability to route main memory traffic originated by the CPUs
 156 through the PL. This technique is known as Programmable Logic In-the-Middle, or PLIM
 157 for short. PLIM was originally proposed in [29]. To fully grasp how PLIM can be achieved,
 158 one needs to understand how memory accesses are routed in PS-PL platforms.

159 Any CPU-generated memory access that results in an LLC miss is routed directly to
 160 main memory if its physical address falls within the aperture, say the address range $[A, B]$
 161 handled by the DRAM controller. We refer to this as the *normal route*, depicted in Fig. 1⑤
 162 and highlighted in yellow.

163 Conversely, generic memory access resulting from an LLC cache miss will be sent on an
 164 HPM port if the corresponding physical address falls within another range, say $[C, D]$. One
 165 can then insert (1) a lightweight layer of virtualization to map all the physical addresses
 166 of a guest OS to the PL, i.e., to fall in the range $[C, D]$; and (2) an address translator in
 167 the PL that re-bases request physical addresses to access main memory and relays back the
 168 data payload to the requesting CPU(s). In other words, one can find a constant k such that
 169 $C = A + k$. Then, the translator in the PL, upon receiving any request at address $x \in [C, D]$
 170 will issue a main memory request at the address $(x - k)$ through the HPS port and provide
 171 the response to the CPU. The PLIM technique introduces a secondary memory route for
 172 reaching the DRAM, called the *PL loop-back*, or simply *loop-back*, which is highlighted in
 173 blue in Fig. 1⑥. Memory transactions on the loop-back route typically traverse the main
 174 interconnect, as depicted in Fig. 1⑦. The advantage of PLIM is that transactions on the
 175 loop-back route can be inspected, blocked, re-routed, and in general managed by custom
 176 re-programmable logic. Importantly, switching from the direct to the loop-back route can



177 be done dynamically at runtime so that the overhead of PLIM can be avoided if deemed
 178 detrimental for the application under analysis.

179 In this paper, we leverage the PLIM approach to perform memory scheduling, hence, we
 180 call our module the Scheduler In-the-Middle, or SchIM for short.

181 3.3 Advanced eXtensible Interface (AXI)

182 The vast majority of PS-PL platforms currently available are ARM-based. This is also the
 183 case for the platform we used for our evaluation, namely the Xilinx Zynq UltraScale+ MPSoC.
 184 Thus, we briefly introduce the communication protocol used for on-chip communication
 185 in ARM-based SoCs, namely the Advanced eXtensible Interface (AXI). The AXI is an
 186 open specification bus protocol [5] used for high-bandwidth data exchanges between on-chip
 187 subsystems — such as cache controllers, memory controllers, DMAs, PL modules. It is also
 188 used in the PS-PL platforms of reference to exchange data on the HPM and HPS ports.

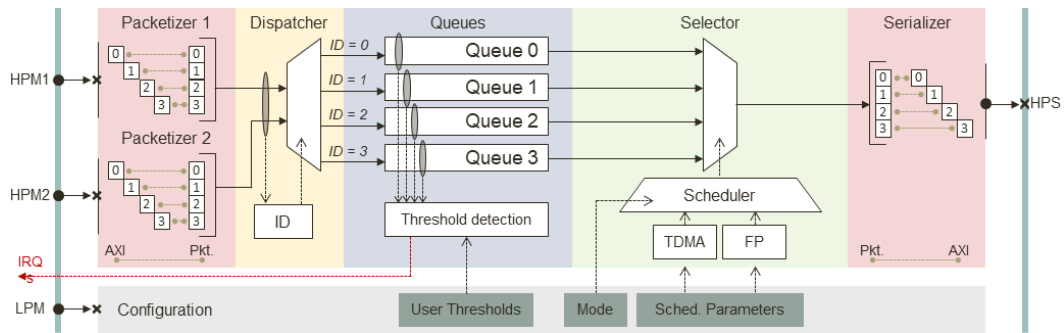
189 The AXI protocol is based on the master-slave duality. A master AXI interface can
 190 initiate transactions toward a connected slave interface. The latter responds master-initiated
 191 requests. Masters and the slaves communicate with each other through five different channels
 192 named AW (address write), W (write), B (write acknowledgment), AR (address read) and R
 193 (read), as illustrated in Fig. 2a.

194 A write transaction begins with an address phase ① where the channel AW is used to
 195 transmit the transaction’s meta-data, such as the destination address, the transaction ID,
 196 and the cacheability attributes the type/length of the burst, and so on. Upon completing
 197 this phase, follows the data phase ②, which consists of the transmission of the data payload
 198 to be written through the W channel. The response phase ③ concludes a successful write
 199 transaction and occurs on the B channel.

200 The transmission of a read transaction is carried out in a similar way. The address phase
 201 ① is transmitted through the equivalent AR channel and is directly followed by the data
 202 phase ②. A response initiated by the slave follows where the read data is transferred over the
 203 R channel. The protocol is asynchronous because different phases of different transactions
 204 can interleave on any AXI bus segment. Hence, multiple outstanding transactions can be
 205 emitted by a single master and the receipt of out-of-order responses is possible.

206 4 Design Goals and Overview

207 In this section, we introduce the proposed SchIM design and describe the overarching goals
 208 of this work. We then provide a bird’s-eye view of the SchIM organization and principles of
 209 operation.



■ **Figure 3** SchIM internal organization connected to the PS via the HPM, LPM and HPS ports.

210 4.1 Design Goals

211 As briefly surveyed in Section 2, there have been numerous proposals for better memory
 212 controllers and approaches to manage memory traffic in modern multi-core embedded
 213 platforms. With respect to the existing literature, the purpose of this work is twofold. First,
 214 we want to demonstrate that scheduling CPU-originated memory traffic at the granularity
 215 of individual transactions is possible in PS-PL platforms. Second, and more importantly,
 216 we want to provide an infrastructure that is generic and extensible enough for the broader
 217 research community to adopt and foster a new chapter on PL-assisted memory scheduling.
 218 With this in mind, we establish the following goals.

219 **Extensible memory scheduling infrastructure.** First and foremost, the SchIM has
 220 been designed with modularity and extensibility in mind. We separate the functionalities
 221 that concern handling, queuing, selection, and forwarding of memory requests inside our
 222 infrastructure. Moreover, we design our SchIM to be able to support multiple memory
 223 scheduling policies simultaneously. A simple, standardized interface is provided to define new
 224 memory scheduling policies without impacting the design of the rest of the SchIM. We discuss
 225 in Section 5.5 the generic interface provided by the SchIM to implement a new memory
 226 scheduling policy.

227 **Runtime configuration and transparency.** We want the SchIM to be a robust
 228 supporting infrastructure to evaluate, compare, and contrast memory scheduling policies.
 229 As such, we strive to provide (1) runtime reconfigurability and (2) operational transparency.
 230 It is possible to rapidly identify desirable configuration parameters by allowing memory
 231 scheduling policies to be switched at runtime. Besides, an adopted policy can be tuned
 232 according to the workload criticality and memory intensiveness. For this purpose, the SchIM
 233 exposes a memory-mapped configuration interface where all the operational parameters can
 234 be changed at runtime. At the same time, we want to ensure that the applications and the
 235 (real-time) operating system under analysis need not be modified to use the SchIM. Hence,
 236 we propose using a thin virtualization layer to selectively route memory traffic through the
 237 SchIM without changes to the binary of OS kernel and applications.

238 **Realistic performance with experimental policies.** One of the limiting factors of
 239 research on memory scheduling policies is the ability to construct evidence of performance
 240 improvements with the realistic workload. Proposing a new memory scheduling policy is
 241 traditionally done with either a simulated setup or with a full-system soft-core implementation.
 242 Both cases have their drawbacks. The former gives a great deal of flexibility but achieving
 243 clock-level accuracy requires simulating many components the SoC whose details might not
 244 be publicly available. In addition, simulated setups that propose custom hardware designs

cannot be directly adopted on real platforms without being first synthesized in hardware. Full soft-core-based SoC implementations suffer from two shortcomings. First, they run at relatively low frequencies and thus can extract only a fraction of the available DRAM bandwidth. Secondly, they are typically based on processors IPs that do not feature the same Instructions Set Architecture (ISA) as widely available COTS, which further limits the practical impacts of these works.

As reported in [1], re-routing the traffic of the core cluster through the PL-side comes at a cost in terms of extra latency and reduced bandwidth. Nonetheless, as PS-PL platforms mature and the interplay of PL and memory resources improves, a SchIM-like design could be the way to go for mission-reconfigurable, upgradable embedded systems.

4.2 Design Overview

As previously mentioned, the SchIM leverages the PLIM approach. CPU-originated main memory transactions are re-routed through the programmable logic and scheduled by the SchIM according to a flexible and configurable policy. The result is that the timing of memory transactions generated by real-time applications can be carefully determined and reasoned upon. Because the SchIM follows a PLIM approach, transactions can be selectively sent to the SchIM for scheduling. However, it is always possible to dynamically exclude the SchIM and route transactions directly to the main memory. Toward this paper's incentive, we consider a setup in which SchIM handles all the CPU-generated memory transactions.

Fig. 1 provides an overview of the location of the SchIM in the reference platform, while its internal organization is visible in Fig. 3. Application memory requests reach the SchIM the aforementioned HPM ports. Without loss of generality, we consider a SchIM instance with two arrival lanes, which are labeled as HPM_1 and HPM_2 in Fig. 3. The SchIM then forwards the received transactions towards main memory through the HPS interface. A more detailed view of the SchIM module is provided in Fig. 3 where the same convention is used to identify input and output ports. In addition, as shown in Fig. 3, a fourth LPM port is used to configure the SchIM from the PS.

The SchIM is composed of a number of sub-modules grouped into three different domains, namely (i) the *interfacing domain*, (ii) the *queuing domain*, and (iii) the *scheduling domain*.

The interfacing domain encompasses the sub-modules to interface the core logic of the SchIM with the rest of the system using the AXI protocol. This is comprised of three sub-modules. These are (i) the *packetizer(s)*, (ii) the *serializer*, and (iii) the previously mentioned *configuration* interface.

The PS-facing end of the **packetizer** offers an AXI slave port to accept new incoming transactions. Upon receipt, this module transforms each transaction into an equivalent *packet* that can be queued and scheduled by SchIM. Packetization of AXI transactions is necessary to be able to store transactions that are serial by nature. A standard AXI transaction is composed of one address phase (AR or AW channel) followed by a data phase (R or W channel), which can be itself composed of multiple successive bursts.

In many ways, the **serializer** is the dual module of the packetizer. Its purpose is to transform the packets that encode CPU-generated memory requests back into AXI-compliant transactions. As such, the serializer offers a master port to the rest of the system to be routed to the main memory controller.

The queuing domain handles how packets are stored between receipt and re-transmission. This domain is comprised of (i) the *dispatcher* module, (ii) the *transaction queues*, and (iii) the *selector* module.

291 The use of **multiple transaction queues** is necessary to differentiate the traffic of the
292 CPUs and perform scheduling. As such, the SchIM associates a queue to each of the active
293 cores — four in the platform of reference. The queues implemented in the SchIM not only act
294 as a holding space for in-flight memory transactions. They also (a) provide information to
295 the scheduling domain regarding their current state, and (b) they can generate a congestion
296 control signal to the associated CPU core.

297 Congestion control is vital because memory transactions originated at the LLC controller
298 follow the same route to the SchIM regardless of the originating CPU. The total number of
299 outstanding transactions that the cores can emit exceeds the queuing elements' capacity on
300 the loop-back route. Hence, priority inversion arises if a low-priority CPU's memory traffic
301 is (temporarily) held. Latter is due to the uncontrolled queue buildup, which provokes a
302 head-of-line blockage. Importantly, what described is true also for the normal route and it is
303 a direct consequence of the best-effort nature of traditional multi-core memory buses. The
304 SchIM allows the user to specify a configurable threshold on the occupancy of the queues
305 that, when reached, issues a regulation signal to the corresponding CPU. We describe in
306 greater detail how congestion control was implemented on the target platform in Section 5.7.

307 As suggested by Fig. 3, transactions are categorized and enqueued based on the source of
308 traffic. The **dispatcher** module performs the matching between an incoming transaction
309 and the destination queue. Similarly, transactions are dequeued by the **selector** module and
310 sent directly to the output of the SchIM following the scheduling domain's resolutions.

311 **The scheduling domain** encompasses all the sub-modules that enable arbitration of
312 transactions issued by the different cores of the PS. The modules in this domain are intended
313 to be generic for extensibility, albeit the first set of two template schedulers is provided as
314 a proof of concept. The scheduling policies currently implemented in the SchIM are Fixed
315 Priority (FP) and Time Division Multiple Access (TDMA). Each of the parameters required
316 by the implemented policies — such as the priorities and the periods — can be adjusted at
317 runtime via the configuration interface.

318 The FP scheduler allows associating a priority value to each of the transaction queues.
319 Pending transactions at the queues are then forwarded out of the SchIM following the
320 user-defined priority order. The TDMA scheduler allows associating a transmission time slot
321 to each of the queues expressed in PL clock cycles. The module then builds a schedule by
322 concatenating the per-core slots so that only pending transactions from one queue at a time
323 are forwarded by the SchIM.

324 **5 SchIM Design and Implementation**

325 A full-system implementation was carried out on a Xilinx ZCU102 development system,
326 which is based on a Xilinx Zynq UltraScale+ XCZU9EG PS-PL SoC. The PS comprises four
327 ARM Cortex-A53 CPUs that share a unified 1 MB LLC. The PS includes a DDR4-2666
328 controller connected to a 4 GB DDR4 memory module. There are two high-performance
329 master interfaces (HPM1 and HPM2); and a third interface routed through the low power
330 domain (LPM). The PL is capable of driving up to 16 interrupt requests lines towards the
331 PS interrupt controller. We hereby provide key details on the operation of our SchIM in the
332 target platform. These include complementary software stack, memory traffic accounting,
333 regulation to prevent head-of-line blocking, and programming model.

334 5.1 Software Stack

335 As mentioned in Section 4.1, we want to ensure that the SchIM can be used with no
336 modification to the OS and the applications under analysis. For this reason, we rely on a
337 thin virtualization layer that can be used to redirect memory traffic from the direct route to
338 the loop-back route (see Section 3.2). For this purpose, we use the open-source Jailhouse [16]
339 partitioning hypervisor² Jailhouse does not boot the target machine. Instead, it relies on a
340 standard Linux kernel to perform the initial boot sequence. When enabled from a Linux
341 driver, Jailhouse dynamically virtualizes the original OS. In line with its partitioning-only
342 philosophy, Jailhouse has a small footprint and enforces virtualization-aided partitioning of
343 essential resources like CPUs, interrupts, main memory, I/O devices. It does not perform
344 any virtual-CPU scheduling.

345 Following Jailhouse’s nomenclature, a resource partition is called a *cell*, while guest OS’s
346 are referred to as *inmates*. An inmate can be either a bare-metal application, an RTOS
347 or a full-fledged OS like Linux. Jailhouse uses ARM hardware Virtualization Extensions
348 (VE) to offer a set of Intermediate Physical Address (IPA) to its inmates that is compatible
349 with the way they have been compiled. Jailhouse then maps IPA ranges of different cells to
350 configurable Physical Addresses (PAs) — stage-2 translation. By changing the configured
351 stage-2 mapping, it is possible to dynamically re-route via the loop-back the memory traffic
352 generated by each inmate.

353 As described below, some modifications were necessary to the mainline Jailhouse code for
354 our full system implementation³.

355 5.2 Altered communication scheme

356 In order to achieve the objective of re-ordering transactions, one must alter the standard AXI
357 communication scheme explained in the Section 3.3. To this end, the SchIM is interposed
358 between the master (HPM) and the slave (HPS) as depicted in Fig. 2b. As shown in Fig. 2b,
359 only the phases initiated by the masters (i.e., address phase on AW and AR and the data
360 phase on W) are intercepted for re-ordering by the SchIM. The introduction of the SchIM
361 has a direct consequence on the overall communication scheme. Unlike the response phases
362 on channels R and B that remain unchanged, the address and write data phases are handled
363 following a store-and-forward scheme. Consequently, a write transaction will start exactly
364 as in the standard AXI scheme with its address phase ① and data phase ②. These two
365 transactions are buffered within the SchIM’s queues (③) and only relayed following the
366 internal memory scheduler’s logic. This release of the transaction leads to the initialization
367 of two new addresses and data phase ④, and ⑤. Finally, the response phase ⑥ goes directly
368 from the slave to the master without being intercepted. For read transactions, the same
369 modifications apply to the address phase ① which is buffered (②) for some time before
370 being re-emitted in ③. Just like for write acknowledgments writing, the read response phase
371 ④ is not intercepted by the SchIM.

372 5.3 Queuing Domain

373 At the heart of the queuing domain, lies the queues. They work as FIFOs. However, instead
374 of inserting the new data at the back of the queue, the new data is always inserted as close

² The source code is available at <https://github.com/siemens/jailhouse.git>.

³ The modified Jailhouse sources are available at <https://github.com/rntmancuso/jailhouse-rt>.

375 as possible to the front of the queue. This mechanism helps avoiding gaps within the queues
 376 prevents the loss of few clock cycles that would be required to move the data from the back
 377 to the front. From the authors' experiments, saving clock cycles in SchIM is vital to keep
 378 the final bandwidth as high as possible.

379 Furthermore, the queues have been designed to deal with three constraints. Firstly, the
 380 queues store both read and write packets such that the order at which transactions arrived
 381 is guaranteed. This implies that all the queue slots have the same size regardless of whether
 382 they contain read or write packets. Secondly, due to the altered communication scheme (see
 383 Section 5.2), each slot needs to be large enough to store both the address phase payload and
 384 the corresponding data of an AXI write transaction (678 bits). The depth of each queue is
 385 determined by considering the worst-case scenario. The latter consists of having to handle
 386 the maximum number of outstanding read and write transactions simultaneously. Our SchIM
 387 instance on the considered Xilinx UltraScale+ platform was configured with queues that are
 388 16 slots in-depth. Indeed, the HPM ports in this platform cannot handle more than eight
 389 transactions of each type [37].

390 5.4 LLC-SchIM Interface and Traffic Accounting

391 As illustrated in Fig. 1, the considered system features an LLC shared between the four cores
 392 of the PS. For a non-cacheable read (resp., write) memory access, which CPU represents
 393 the source of the traffic is carried in the ID bits of the corresponding AR (resp., AW) AXI
 394 transaction. But for cacheable memory accesses, which is the norm for application workload,
 395 this is not the case. This is mainly because cache controllers typically use a write-back
 396 strategy. In this case, a read or write cache miss causes up to two events: (1) a cache refill
 397 and (2) a cache eviction. The cache refill is carried out with a read AXI transaction. If
 398 the line being evicted was previously written (dirty), then the eviction causes a write AXI
 399 transaction. It follows that, while read AXI transactions have an easily identifiable source,
 400 write transactions do not. Indeed, a CPU x might be causing the eviction of a line previously
 401 allocated and modified by CPU y . Hence, accounting (and scheduling) the resulting write
 402 transaction as if it originated from CPU x would be incorrect.

403 To ensure fair accounting for both read and write traffic, we rely on cache partitioning
 404 through coloring. As studied in a number of previous works, cache coloring is easy to
 405 implement at the hypervisor level [15, 21, 32]. In our system setup, we leverage the support
 406 Jailhouse already provides. The standard support has been extended to support booting
 407 a Linux inmate over colored memory. Cache partitioning allows us to establish a 1-to-1
 408 relationship between any read/write transaction traversing the SchIM and the originating
 409 CPU. Moreover, with cache coloring in place, the SchIM uses the color bits in the address
 410 of the memory transactions (AR and AW channels) — instead of the AXI ID bits — to
 411 differentiate between the traffic of the various cores.

412 Finally, recall that the SchIM forwards transactions between HPM and HPS ports. These
 413 ports follow the asynchronous AXI protocol that allows issuing multiple outstanding AR and
 414 AW transactions. The protocol dictates that any outstanding transaction must have a unique
 415 AXI ID. This property is crucial to be able to match received responses with outstanding
 416 requests. Unfortunately, a potential mismatch between the bit-width of the AXI ID emitted
 417 at the HPM ports and the bit-width of AXI ID accepted by the HPS ports. For instance, in
 418 the platform of reference, the HPMS emit 16-bit AXI IDs, while the HPS AXI ID bit-width
 419 is 6 bits. Therefore, the SchIM also acts as an AXI ID translator.

420 5.5 Scheduling Interface and Implemented Policies

421 All the memory schedulers included in the scheduling domain share a common interface to
422 ease the integration of a new scheduler. In terms of input signals, a generic scheduler module
423 must define (1) a manual reset signal that can be triggered through the configuration port;
424 (2) a vector of bits where each bit indicates whether the associated queue is empty; and (3) a
425 signal indicating if the last scheduled transaction has been consumed. Alongside these inputs,
426 the scheduling modules also have access to all the configuration registers listed in Table 1.
427 In terms of outputs a SchIM scheduler must define (1) a signal to the selector indicating
428 the queue considered for scheduling; and (2) a signal stating whether the current scheduling
429 decision is valid. We hereby review the initial set of memory scheduling policies implemented
430 in the SchIM.

431 5.5.1 Fixed Priority

432 The FP scheduling module aims at enforcing strict prioritization of cores' memory traffic.
433 The priority ordering is explicitly defined by the user through the configuration port. While
434 the SchIM instance used in this paper only has four queues, 16 different levels of priority
435 are offered as the considered platform supports up to 16 different colors. This is useful if an
436 hypervisor that supports vCPU scheduling is used. In this case, the SchIM allows assigning
437 different priorities to different partitions sharing the same physical CPU. The core-to-priority
438 assignment must be strict, meaning that two cores cannot be assigned the same priority.

439 The FP scheduling module only needs two pieces of information. That is (1) the priority
440 associated with each queue and (2) whether a given queue contains at least one buffered
441 transaction. The module logic always selects the queue with the highest priority. Lower
442 priority queues are considered when higher priority queues do not have transactions. This is
443 done by internally setting the user-defined priority of a queue as 0 when the corresponding
444 queue is empty.

445 5.5.2 Time Division Multiple Access

446 The TDMA memory scheduler is a non-work conserving policy that operates by defining a
447 per-core time *slot* during which the core has exclusive access to main memory. The slots are
448 expressed in PL clock cycles, to maximize granularity. The configuration port can be used to
449 specify and change the slots specifications at runtime.

450 The implementation of the module uses a counter register to track the time elapsed in
451 the current TDMA primary frame — defined as the sum of all the cores' slots. It is reset
452 to 0 at the beginning of a new major frame. Using the time-tracking register, the module
453 determines to which core the current slot belongs, and forwards the information to the queue
454 selector. This is done by summing up the length of all the previous slots, and determining if
455 the current time falls within the interval of the considered core's slot.

456 5.6 Programming Model

457 The parameters that compose the programming interface of the SchIM are summarized in
458 Table 1. The `base` address referenced in the table can be set when the SchIM is deployed in
459 the PL. By default, this is set to `0x80000000`. All the parameter registers are 32 bit wide,
460 except for the priorities of the FP scheduler. In this case, the priority values are encoded
461 using 8 bits. The last “Mode” register allows a user to select the active memory scheduler.

■ **Table 1** Available SchIM configuration registers.

Parameter	Associated Core	Address
TDMA slots	C_0	base+0x00
	C_1	base+0x04
	C_2	base+0x08
	C_3	base+0x0C
User Thresholds	C_0	base+0x10
	C_1	base+0x14
	C_2	base+0x18
	C_3	base+0x1C
FP Priorities	C_0 C_1 C_2 C_3	base+0x20
Reserved		
Mode	N/A	base+0x38

462 5.7 PL-to-PS Feedback

463 Each of the HPM ports interfacing the PS and the PL sides (HPM1 and HPM2) have two
 464 dedicated queues for read and write transactions. Since transactions are being buffered inside
 465 SchIM as well as in these port buffers, head-of-line blocking can happen. Head-of-the-line
 466 blocking is harmful for performance; and can cancel out the benefits of transaction scheduling
 467 performed by the SchIM. For instance, in the case of a non work-conserving policy (e.g.,
 468 TDMA), if the HPM port queue gets filled with transaction coming for the same core, no
 469 other transaction will be able to reach the SchIM and thus be considered for scheduling. This
 470 implies that no transaction would be scheduled until the end of the active core’s TDMA slot.
 471 On the other hand, for work-conserving policies (e.g., FP) in the presence of head-of-line
 472 blocking, the decisions being taken by SchIM would directly depend on the order at which
 473 transactions are emitted by the HPM port buffer.

474 In both cases, one must prevent the cores from saturating the HPM port buffers. In
 475 order to avoid such situation, we implemented a feedback scheme aimed at slowing down
 476 the cores when necessary. As we mentioned in the context of Fig. 3, the SchIM’s queues are
 477 associated a programmable threshold. Whenever the queue occupancy reaches (or exceeds)
 478 the associated threshold, a per-core interrupt line is asserted from the PL to the PS side.
 479 When received, the interrupt is treated by the platform software as a *fast interrupt request*
 480 (FIQ) and directly handled by the hypervisor—invisible to any guest OS. The advantage of
 481 using FIQs instead of regular IRQs is the significantly reduced handling latency [31]. Minor
 482 modifications to the TrustZone monitor were necessary to correctly configure FIQ handling.
 483 To minimize overhead, the installed FIQ handler only executes two assembly instructions.
 484 These are (1) a **dsb** memory barrier that stops the core until all the outstanding memory
 485 transactions have been completed, and (2) a **eret** instruction to exit the FIQ context. There
 486 is not need to save/restore any register because FIQs have banked syndrome/status registers
 487 and because no general purpose register is modified in the handler.

488 Ideally, the available space in the HPM buffers should be shared evenly between the cores.
 489 Since each HPM port has a buffer with a depth of 8+8 transactions, each core should occupy
 490 at most 2 slots in each buffer. Unfortunately, our experiments highlighted that the control
 491 over amount of transactions buffered by each core is imperfect. Often times, the selected
 492 threshold is exceeded by up to two transactions. This is the main reason why we propose
 493 a dual-ported SchIM which uses both the available HPM ports. Indeed, by assigning two

494 cores on each of the ports, the ideal threshold on maximum amount buffered transactions
495 can be doubled. The increase provides enough room to compensate for imperfections in the
496 micro-regulation performed with PL-to-PS FIQ delivery.

497 **6** Evaluation


498 The present section aims at evaluating the behavior of the SchIM on the target platform, its
499 overhead and benefits. First, in subsection 6.1, we review our experimental setup. Thereafter,
500 we assess the overhead introduced by the SchIM in Section 6.2. Section 6.3 explores the
501 impact of the PL-to-PS feedback on the control and the performance. In Section 6.4, an
502 in-depth analysis of the SchIM’s behavior is presented. Finally, an evaluation of the temporal
503 behavior of a set of real-world benchmarks operating through the SchIM is provided in
504 Section 6.5.

505 **6.1** Experimental Setup

506 The SchIM has been evaluated using synthetic benchmarks (or *Memory Bombs*), real
507 benchmarks selected from the San Diego Vision Benchmark Suite (SD-VBS) [35] and a
508 combination of the two. Specifically, seven memory-intensive benchmarks have been selected,
509 i.e. *stitch*, *texture synthesis*, *disparity*, *tracking*, *localization*, *mser* and *sift*. For our runs, we
510 have considered all the intermediate input sizes ranging from SQCIF (128×196 pixels) to
511 VGA (640×480 pixels). When running any benchmark, we use the cache coloring mechanism
512 implemented in the Jailhouse hypervisor [32] to partition the LLC evenly amongst the 4 cores
513 and to prevent our measurements from being affected by inter-core cache interference. As a
514 result, each benchmark operates on 1/4 of the total cache space—256 KB. As extensively
515 discussed in [14, 41], it is also important to avoid inter-core DRAM bank conflicts, which
516 can cause the arbitrary re-ordering of transactions originating from different cores. This is
517 accomplished by (1) configuring the DRAM controller to disable DRAM bank interleaving;
518 and (2) by performing static cache bleaching [11, 29] at the SchIM’s output to re-compact
519 accesses to colored pages into contiguous DRAM accesses. In this platform, there are a
520 total of 16 DRAM banks of 256 MB each. Thanks to bleaching, we can assign the full size
521 of 4 banks (i.e., 1 GB) to each core, instead of being restricted to only 1/4 of that due to
522 non-overlapping color and bank address bits.

523 To evaluate the capabilities of the SchIM, two memory routes for the traffic generated
524 by the cores are compared. The first serves as baselines, whereas, the last one is the one
525 under analysis and involves the SchIM module. The first path consists in the cores directly
526 accessing the main memory. As illustrated in Fig. 1, the traffic simply goes through the
527 *Main Interconnect* before arriving at the DDR controller. This path is referred to as the
528 *normal route*. Secondly, we consider the case where the SchIM module is deployed and in use
529 to schedule memory traffic generated by the CPUs in the PL. Cores 0 and 1 target HPM1
530 aperture, while cores 2 and 3 target HPM2. In our analysis, the SchIM is used in all the
531 available modes, i.e., FP and TDMA.

532 Note that in the case of the *normal route*, combining both a strict cache partitioning and
533 strict bank partitioning could not be applied. In fact, as a direct consequence of the address
534 coloring and in the absence of a bleacher, only 1/16 of each 1 GB wide memory allocation
535 can be used by each core. The resulting reduced space of 64 MB is not enough for running
536 Linux. Consequently, in the case of the *normal route*, the cores have been split into two
537 groups of two, where each group targets independent sets of banks. This configuration allows
538 the cache to be partitioned using address coloring.



■ **Figure 4** Bandwidth in MBps for different path under increasing set of cores contending.


539 6.2 Platform Capabilities and performance degradation

540 Intuitively and as discussed in [29], redirecting the traffic coming from the cores to the PL
 541 side incurs a performance hit. In spite of the lower frequency at which the SchIM operates
 542 (250 MHz), the theoretical throughput when using both the HPM lanes should be around
 543 8 GBps. We observe, however, that the achievable throughput through the HPM ports is
 544 a fraction of what we measured by accessing the main memory through the *normal route*
 545 (2116.5 MBps and 1207.41 MBps for solo and full contention by 3 other cores, respectively).
 546 We further provide a discussion on the bandwidth drop when transactions are routed through
 547 the PL in Section. For the sake of completeness, we quantify in Fig. 4 the maximum
 548 bandwidth achieved through the PL — and hence through the SchIM. Nevertheless, it is
 549 important to remember that the absolute figures are strictly platform dependent.


550 In Fig. 4, we have computed the throughput of one *core under analysis*, here core 0 (noted
 551 C_0) when a synthetic memory-intensive application is deployed on an increasing number
 552 of cores denoted with the same notation. The first bar cluster (“Normal”) refers to the
 553 throughput measured via the normal route. The other two clusters capture the observed
 554 bandwidth when traffic is routed through and managed by the SchIM. One cluster is provided
 555 for each of the implemented memory scheduling policies, namely — from left to right — FP
 556 and TDMA. As expected, there is a sharp reduction (around 75%) in terms of absolute
 557 bandwidth. Importantly, however, two aspects need to be highlighted. First, the bandwidth
 558 achieved through the SchIM is still remarkably high and allows studying the behavior of the
 559 realistic workload under custom memory scheduling policies, which is the primary goal of
 560 this research. Second, it emerges that the implemented FP and TDMA policies are capable
 561 of protecting the core under analysis from inter-core interference, while this is not the case
 562 when going through the normal route

563 6.3 PL-to-PS feedback performance impact

564 As mentioned in Section 5.7, the PL-to-PS feedback enables our SchIM to regulate the HPM
 565 ports buffer occupancy to prevent head-of-line blocking. Since this feedback directly throttles
 566 the desired core, the selection of an adequate threshold is important to preserve the balance
 567 between control and performance. Therefore, in Fig. 5, we have explored the sensitivity to
 568 the threshold for each of the proposed schedulers under different levels of contention. The
 569 thresholds in use range from 1 to 8 and even include the case where the feedback mechanism



(a) Threshold-Bandwidth relationship curves for the FP scheduler



(b) Threshold-Bandwidth relationship curves for the TDMA scheduler

■ **Figure 5** Figures showing the impact of the threshold in use on the final bandwidth experienced by the cores for the offered schedulers

570 is disabled (noted *NA*). The contention is created by up to four co-running cores emitting
 571 write transactions. For each parameter applied to a scheduler (i.e., fixed priority or TDMA
 572 slot), the co-running cores are assigned the most demanding parameters available (i.e., the
 573 highest priority for FP or the biggest TDMA slot).

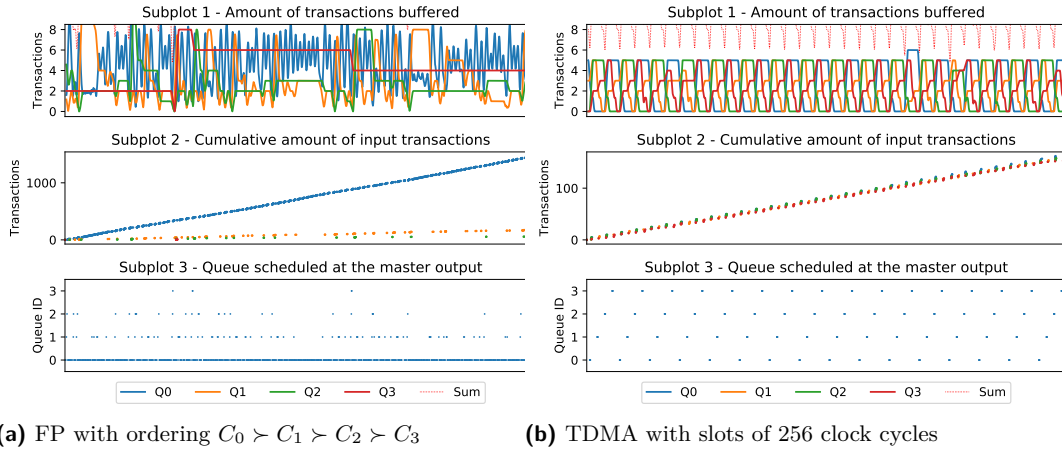
574 In the case of the FP scheduler (Fig. 5a), one can observe that when running alone, the
 575 threshold has no influence on the throughput. However, as soon as co-runners are added, the
 576 cores start to experience a decrease in throughput. Fig. 5b shows that the TDMA scheduler
 577 is not impacted considerably by the threshold with respect to the throughput. Globally, the
 578 scheduler manages to preserve a constant throughput regardless of the contention and the
 579 assigned slot.

580 Nonetheless, under high contention, one can observe that the throughput of each core is
 581 affected. The fourth inset of Fig. 5a and Fig. 5b illustrate the importance of the threshold and
 582 the PL-to-PL feedback mechanism as a considerable drop of throughput can be observed
 583 for the highest priority of FP and for a TDMA period of 32.

584 Considering these experiments, setting the threshold to four for all the schedulers seems
 585 to bring the best trade-off between control and performance. However, this value cannot be
 586 blindly applied to all cases as this experiment is performed for a sequential and contiguous
 587 access pattern.

588 6.4 Internal Behaviour of SchIM

589 The next objective is to verify the correct behavior of the schedulers at the granularity of
 590 a clock cycle by observing the inputs, the outputs and the internal signals and registers
 591 of the SchIM module. This is made possible thanks to the *Integrated Logic Analyzer* (or
 592 ILA) provided by Xilinx [36]. The latter IP can be directly implemented on the PL side,
 593 alongside the SchIM, and is able to probe the signals and to store them in a local memory.
 594 For this experiment, a group of relevant internal signals have been probed and captured
 595 during a window of 16384 contiguous clock cycles. Then, the information has been extracted
 596 by post-processing the data. To characterize the behavior of the two different policies, the



■ **Figure 6** Trace snapshots of SchIM for FP (6a), TDMA (6b)


597 ILA has been instrumented to collect (i) the amount of transactions being buffered in the
 598 queues at each clock cycle (inset 1 in Fig. 6a and Fig. 6b) (ii) the rate at which queues receive
 599 new transactions from the cores cluster (inset 2 in Fig. 6a and Fig. 6b) and (iii) the queues
 600 ID of each transaction forwarded by the SchIM module (inset 3 in Fig. 6a and Fig. 6b).

601 For the Fixed Priority trace snapshot displayed in Fig. 6a, the following strict priority
 602 ordering has been considered: $C_0 \succ C_1 \succ C_2 \succ C_3$ where the \succ operator means that the
 603 left argument has a strictly higher priority than the right argument. In this experiment,
 604 a regulation threshold of 3 for each core has been used. As emphasized by the inset 2 in
 605 Fig. 6a, the FP scheduler is able to prioritize the traffic of one core at the expense of the
 606 others according to the priorities assignment. Furthermore, one can observe that the rate at
 607 which the queues receive new transactions from their associated core is proportional to the
 608 priority level in the priority ordering. Finally, the third inset in Fig. 6a confirms the correct
 609 behavior of the FP policy. One can see that the cores with the highest priority also feature
 610 the highest density of transactions at the output of the SchIM.

611 The trace snapshot displayed in Fig. 6b has been obtained by configuring the SchIM
 612 module in TDMA mode. For the sake of clarity, a slot of 256 clock cycles has been set for each
 613 core. Besides, the threshold of each core has been set to 4 to create sharp transitions. The
 614 insets 2 and 3 of Fig. 6b clearly show the behavior expected from a TDMA schedule. In fact,
 615 one can clearly see in the latter that transactions originating from one core are only being
 616 repeated out of the SchIM module during a well-defined and periodic time slot of 256 clock
 617 cycles. In the inset 2 of Fig. 6b, we can observe a similar pattern, with transactions arriving
 618 only during the TDMA slot associated with their queue (and indirectly core). Globally, the
 619 rate at which queues receive transactions is steady and constant.

620 6.5 Memory Isolation

621 On the platform considered for this set of experiments, the Xilinx ZCU102 development board,
 622 we denote three main sources of inter-core performance interference: (1) cache contention,
 623 (2) DRAM bank conflicts, and (3) the congestion and saturation of the memory controller.
 624 Despite their orthogonality, the two first sources are tackled respectively via the integration
 625 of page coloring in the hypervisor and static bleaching in the SchIM. On the other hand,
 626 since the SchIM provides fine-grained control over the timing and ordering of transactions



■ **Figure 7** Normalized execution time for each benchmark and input size for *Solo* and *Stress*. Each column denotes a given benchmark of the SD-VBS suite, while each row denotes a specific input size (in increasing order from top to bottom).

627 originating from the application cores as they reach the memory controller. Thus, the SchIM
 628 brings memory bandwidth management into the PL, and provides not only regulation but a
 629 generic infrastructure to experiment with custom bandwidth management techniques, both
 630 work-conserving and non-work-conserving.

631 The evaluation setup considered for this experiment is identical to the one presented in
 632 Section 6.1. The routes going through the PL and using our SchIM (i.e., FP and TDMA)
 633 benefit from both cache partitioning and bank partitioning. On the other hand, the *normal*
 634 *route* uses cache partitioning and sees its cores divided into two sets targeting each a different
 635 group of private banks.

636 To evaluate the capability of our SchIM with respect to its ability to ensure performance
 637 isolation between the cores, a set of experiments involving SD-VBS benchmarks were designed.
 638 Here, we compare the execution time of an application on a given core when running alone
 639 (referred to as *Solo*) and when running alongside interfering synthetic benchmarks (write
 640 memory bombs) on all the other cores (referred to as *Stress*). For each combination of a
 641 route to main memory (i.e., the *normal route* or the *SchIM route*) and scheduler, the result
 642 obtained for *Stress* is normalized with respect to the equivalent configuration in *Solo*. The
 643 results obtained on the considered benchmarks are listed in Fig. 7. The results in the Fig. 7
 644 are the aggregation (arithmetic average) of 30 different runs in the same configuration. Each
 645 bar cluster of the Fig. 7 insets represents one of the aforementioned configuration for *Solo*
 646 and *Stress*. The height of each bar denotes its normalized execution time.

647 For this set of experiments, the FP scheduler was configured such that the core under
 648 analysis (i.e., the one running the benchmark) has the highest priority and a threshold of 8.
 649 The other cores are assigned lower priorities and thresholds matching their priority order

650 (i.e., 4, 2, 1). Under TDMA scheduling, the core under analysis has a slot of 512 clock cycles
 651 and a threshold of 14 while the co-runners are assigned slots of 32 and 16 clock cycles with
 652 thresholds of 4 and 1.

653 The *normal route* is used as a baseline for this experiment because no scheduling is
 654 performed in this configuration. The Fig. 7 highlights the sensitivity of both *disparity* and
 655 *mser* to inter-core interference on the *normal route*. This is especially the case for large
 656 input sizes such as *cif* and *vga*. On the other hand, *texture synthesis* and *localization* do
 657 not suffer from inter-core interference. Globally, the TDMA scheduler always manages to
 658 preserve the isolation of the core, having execution times under *Stress* similar or smaller than
 659 the *normal route*. This is particularly visible for *qcif*, *cif* and *vga* input sizes of *disparity*
 660 and *mser*. Similarly, the FP scheduler is also capable of ensuring sound isolation of the core
 661 under analysis.

662 **7 Discussion and Limitations**

663 By design, the PLiM module proposed in this paper, the SchIM, centralizes the memory
 664 traffic and its scheduling. A centralized design makes sense on the specific target platform
 665 because there exist only one memory controller and thus a single path between the LLC and
 666 the DRAM controller. In systems where multiple paths between the processing units and the
 667 memory controllers exist, for instance when multiple controllers and channels are present, a
 668 decentralized design is to be preferable to better exploit the available memory parallelism.
 669 In such platforms, a possible avenue could be instantiating multiple SchIM modules, roughly
 670 one per channel, and introducing appropriate out-of-band signaling between the modules for
 671 coordination off the critical path.

672 As we mentioned in Section 6.1, our setup includes the Jailhouse partitioning hypervisor.
 673 While the SchIM module does not strictly require the PS side to use a hypervisor, Jailhouse
 674 has been extensively used for the evaluation as it provides convenient features to control
 675 physical memory allocation. For instance, the support for page coloring has been used to
 676 both partition the LLC space and to easily identify the owner of each memory transactions
 677 in the SchIM (as presented in Section 5.4). However, instead of enforcing cache partitioning,
 678 one could instead identify the ownership of memory transactions by extracting a different
 679 subset of address bits. For instance, if the physical memory allocated to different partitions
 680 is not interleaved, then the most significant bits of the address can be used to perform
 681 traffic accounting. In addition, the IPA address virtualization is convenient to transparently
 682 redirect the memory traffic of the application partitions through the PL side, even if they
 683 are initially booted through the normal route. Finally, the cores throttling mechanism (see
 684 Section 5.7) via the FIQs can be implemented at EL3 (Secure Monitor) or in the individual
 685 guest OS's instead (EL1). Implementing FIQ handling in the hypervisor (EL2), however,
 686 has the advantage of not requiring any change in the guest OS's, as well as not requiring a
 687 full switch into secure mode compared to an implementation at EL3.

688 On the same note, provided that the FIQ lines are not used by the inmates, the feedback
 689 regulation mechanism is entirely transparent to the guest OS's (or even for bare-metal
 690 applications) and introduces minimum overhead. The Linux kernel do not use FIQs, and
 691 the same goes for typical RTOS's. Nonetheless, it must be acknowledged that defining a
 692 FIQ handler to be used for CPU throttling might interfere with (and be interfered by) the
 693 latency of FIQ handling in guest OS's that rely on the same functionality. This is mainly
 694 because FIQ handling is non-preemptive. We also recognize that the PL-to-PS feedback
 695 mechanism is relatively coarse. Inset 1 of Fig. 6b highlights this problem. Even though

696 all the queues have been assigned a threshold of 4, the threshold is often exceeded. The
697 worst-case being queue 3 exceeding the threshold by 2 on the right-hand side of the plot.
698 This problem can be attributed to the reaction time of the FIQ routine, and to the fact that
699 jumping to the FIQ handler itself might cause a few memory transactions depending on the
700 cache state. Currently, the thresholds used for FIQ-based regulation require to be fine-tuned
701 manually by the user. Future extensions of the SchIM will explore the implementation of
702 schedulers capable of dynamically adapting the thresholds to maximize performance and
703 improve isolation.

704 The loss in bandwidth caused by routing transactions through the PL is important and a
705 serious drawback against the adoption of the SchIM. Our experiments in Section 6.2 have
706 shown that rerouting the traffic through the PL has a cost. As illustrated in Fig. 4, up to
707 2100 MBps can be extracted from the *normal route* whereas any route through the PL only
708 achieves around 320 MBps. In contrast, a back-of-the-envelope calculation reveals that for
709 a PL operating at 250 MHz (the SchIM frequency), and with a bus width of 128 bits, a
710 full-duplex throughput of approximately 3.7 GBps can be sustained. This calculation is in
711 line with the reported throughput in an experiment conducted in [19], in which PL-originated
712 transactions targeting the DRAM passed through the one of the HP ports. This suggests
713 that the PL-to-DRAM route can sustain a much higher throughput than what has been
714 experimentally observed in our evaluation setup, where transactions originate from the PS
715 side. In light of these considerations, we can conclude that the source of the bandwidth
716 loss can be imputed to the bus segments connecting the CPU cluster to the HPM ports.
717 A focused study is necessary to narrow down the exact reason for the performance drop.
718 Nonetheless, vendor-imposed bandwidth throttling, PS-to-PL clock-domain crossing delays,
719 and shallow FIFOs at the HPM ports and/or at the main PS-side interconnect represent
720 plausible reasons. We anticipate that due to the platform-specific nature of this issue, the
721 raw performance of the SchIM will substantially vary across different SoCs.

722 **8 Conclusion**

723 In the present article we introduced the SchIM, a memory transactions scheduler framework
724 that can be integrated with commercially available platforms featuring a tightly coupled
725 processing system and programmable logic. A full-system implementation in a commercially
726 available PS-PL platform has been detailed, which encompasses the accompanying software
727 stack and the platform-specific integration steps have been detailed in as well as advanced
728 scheduling techniques are few of many possible future directions.

729 Through a set of experiments, we assessed the capabilities of the framework and demon-
730 strated the correct behavior of the proposed scheduling policies, namely Fixed Priority, Time
731 Division Multiple Access and Traffic Shaping. Finally, we showed using a suite of real-world
732 benchmarks that the SchIM is capable of enforcing strong temporal isolation despite heavy
733 memory contention.

734 The authors see the proposed SchIM as a stepping stone to propose, test and validate novel
735 memory scheduling policies to be tested on embedded platforms with realistic performance
736 and complex workload. For this reason, the SchIM has been designed to be open-source and
737 with extensibility in mind. Especially, we strongly envision that the SchIM could represent a
738 stepping-stone toward profile-based memory traffic scheduling.

739 — References

- 740 1 B. Akesson. *Predictable and composable system-on-chip memory controllers*. PhD thesis, 2010.
741 doi:10.6100/IR658012.
- 742 2 B. Akesson, K. Goossens, and M. Ringhofer. Predator: a predictable SDRAM memory
743 controller. In *Proceedings of the 5th IEEE/ACM international conference on Hardware/software
744 codesign and system synthesis*, pages 251–256, 2007.
- 745 3 G. Alonso, T. Roscoe, D. Cock, M. Ewaida, Kaan Kara, Dario Korolija, D. Sidler, and
746 Ze ke Wang. Tackling hardware/software co-design from a database perspective. In *Conference
747 on Innovative Data Systems Research (CIDR)*, Amsterdam, Netherlands, Jan. 2020.
- 748 4 ARM. ARM® CoreLink™ QoS-400 Network Interconnect Advanced Quality of Service, 2013.
749 Accessed on 09.01.2020.
- 750 5 ARM. AMBA AXI and ACE Protocol Specification. Technical report, 2019. URL: https://static.docs.arm.com/ihi0022/g/IHI0022G_amba_axi_protocol_spec.pdf.
- 751 6 A. Biondi, A. Balsini, M. Pagani, E. Rossi, M. Marinoni, and G. Buttazzo. A framework for
752 supporting real-time applications on dynamic reconfigurable FPGAs. In *2016 IEEE Real-Time
753 Systems Symposium (RTSS)*, pages 1–12, 2016. doi:10.1109/RTSS.2016.010.
- 754 7 J. M. Calandrino, H. Leontyev, A. Block, U. C. Devi, and J. H. Anderson. LITMUS^{RT}
755 : A testbed for empirically comparing real-time multiprocessor schedulers. In *2006 27th
756 IEEE International Real-Time Systems Symposium (RTSS'06)*, pages 111–126, 2006. doi:
757 10.1109/RTSS.2006.27.
- 758 8 F. Farshchi, Qijing Huang, and H. Yun. BRU: Bandwidth regulation unit for real-time
759 multicore processors. *2020 IEEE Real-Time and Embedded Technology and Applications
760 Symposium (RTAS)*, pages 364–375, 2020.
- 761 9 F. Farshchi, P. Kumar, R. Mancuso, and H. Yun. Deterministic Memory Abstraction and Sup-
762 porting Multicore System Architecture. In Sebastian Altmeyer, editor, *30th Euromicro Confer-
763 ence on Real-Time Systems (ECRTS 2018)*, volume 106 of *Leibniz International Proceedings in
764 Informatics (LIPIcs)*, pages 1:1–1:25, Barcelona, Spain, July 2018. Schloss Dagstuhl–Leibniz-
765 Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2018/9001>,
766 doi:10.4230/LIPIcs.ECRTS.2018.1.
- 767 10 C. Ferri, A. Marongiu, B. Lipton, R. Bahar, T. Moreshet, L. Benini, and M. Herlihy. SoC-TM:
768 integrated HW/SW support for transactional memory programming on embedded MPSoCs. In
769 *Proceedings of the seventh IEEE/ACM/IFIP international conference on Hardware/software
770 codesign and system synthesis*, pages 39–48, 2011.
- 771 11 G. Gracioli, R. Tabish, R. Mancuso, R. Mirosanlou, R. Pellizzoni, and M. Caccamo. Designing
772 mixed criticality applications on modern heterogeneous MPSoC platforms. In *31st Euromicro
773 Conference on Real-Time Systems (ECRTS 2019)*. Schloss Dagstuhl–Leibniz-Zentrum fuer
774 Informatik, 2019.
- 775 12 Intel, Corp. Intel’s Stratix 10 FPGA: Supporting the smart and connected revolution,
776 October 2016. Accessed on 09.01.2020. URL: [https://newsroom.intel.com/editorials/
777 intels-stratix-10-fpga-supporting-smart-connected-revolution/](https://newsroom.intel.com/editorials/intels-stratix-10-fpga-supporting-smart-connected-revolution/).
- 778 13 A. K. Jain, S. Lloyd, and M. Gokhale. Microscope on memory: MPSoC-enabled computer
779 memory system assessments. In *2018 IEEE 26th Annual International Symposium on Field-
780 Programmable Custom Computing Machines (FCCM)*, pages 173–180, 2018. doi:10.1109/
781 FCCM.2018.00035.
- 782 14 H. Kim, D. de Niz, B. Andersson, M. Klein, O. Mutlu, and R. Rajkumar. Bounding memory
783 interference delay in COTS-based multi-core systems. In *2014 IEEE 19th Real-Time and
784 Embedded Technology and Applications Symposium (RTAS)*, pages 145–154, 2014. doi:10.
785 1109/RTAS.2014.6925998.
- 786 15 H. Kim and R. Rajkumar. Real-time cache management for multi-core virtualization. In *2016
787 International Conference on Embedded Software (EMSOFT)*, pages 1–10, 2016.
- 788 16 J. Kiszka, V. Sinitin, H. Schild, and contributors. Jailhouse Hypervisor. Accessed on
789 09.01.2020. URL: <https://github.com/siemens/jailhouse>.
- 790

- 791 17 C. Maiza, H. Rihani, J. Rivas, J. Goossens, S. Altmeyer, and R. Davis. A Survey of Timing
792 Verification Techniques for Multi-Core Real-Time Systems. *ACM Comput. Surv.*, 52(3), June
793 2019. doi:10.1145/3323212.
- 794 18 Microsemi — Microchip Technology Inc. PolarFire SoC - Lowest Power, Multi-Core RISC-
795 V SoC FPGA, July 2020. Accessed on 09.01.2020. URL: [https://www.microsemi.com/
796 product-directory/soc-fpgas/5498-polarfire-soc-fpga](https://www.microsemi.com/product-directory/soc-fpgas/5498-polarfire-soc-fpga).
- 797 19 S. Min, S. Huan, M. El-Hadedy, J. Xiong, D. Chen, and W. Hwu. Analysis and optimization of
798 I/O cache coherency strategies for SoC-FPGA device. In *2019 29th International Conference
799 on Field Programmable Logic and Applications (FPL)*, pages 301–306, 2019. doi:10.1109/
800 FPL.2019.00055.
- 801 20 R. Mirosanlou, M. Hassan, and R. Pellizzoni. DRAMbulism: balancing performance and
802 predictability through dynamic pipelining. In *2020 IEEE Real-Time and Embedded Technology
803 and Applications Symposium (RTAS)*, pages 82–94, 2020. doi:10.1109/RTAS48715.2020.
804 00–15.
- 805 21 P. Modica, A. Biondi, G. Buttazzo, and A. Patel. Supporting temporal and spatial isolation
806 in a hypervisor for ARM multicore platforms. In *2018 IEEE International Conference on
807 Industrial Technology (ICIT)*, pages 1651–1657, 2018.
- 808 22 O. Mutlu and T. Moscibroda. Stall-time fair memory access scheduling for chip multiprocessors.
809 In *40th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2007)*,
810 pages 146–160. IEEE, 2007.
- 811 23 O. Mutlu and T. Moscibroda. Parallelism-aware batch scheduling: Enhancing both performance
812 and fairness of shared DRAM systems. In *2008 International Symposium on Computer
813 Architecture*, pages 63–74. IEEE, 2008.
- 814 24 K. Nesbit, N. Aggarwal, J. Laudon, and J. Smith. Fair queuing memory systems. In *2006
815 39th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'06)*, pages
816 208–222. IEEE, 2006.
- 817 25 M. Paolieri, E. Quinones, F. Cazorla, and M. Valero. An analyzable memory controller for
818 hard real-time CMPs. *IEEE Embedded Systems Letters*, 1(4):86–90, 2009.
- 819 26 N. Rafique, W. Lim, and M. Thottethodi. Effective management of DRAM bandwidth in
820 multicore processors. In *16th International Conference on Parallel Architecture and Compilation
821 Techniques (PACT 2007)*, pages 245–258. IEEE, 2007.
- 822 27 F. Restuccia, A. Biondi, M. Marinoni, G. Cicero, and G. Buttazzo. AXI HyperConnect: A
823 predictable, hypervisor-level interconnect for hardware accelerators in FPGA SoC. In *2020
824 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6, 2020. doi:10.1109/
825 DAC18072.2020.9218652.
- 826 28 F. Restuccia, M. Pagani, A. Biondi, M. Marinoni, and G. Buttazzo. Is your bus arbiter really
827 fair? restoring fairness in AXI interconnects for FPGA SoCs. *ACM Trans. Embed. Comput.
828 Syst.*, 18(5s), October 2019. doi:10.1145/3358183.
- 829 29 S. Roozkhosh and R. Mancuso. The potential of programmable logic in the middle: Cache
830 bleaching. In *26th IEEE Real-Time and Embedded Technology and Applications Symposium
831 (RTAS 2020)*, Sydney, Australia, April 2020.
- 832 30 P. Sohal, R. Tabish, U. Drepper, and R. Mancuso. E-WarP: a system-wide framework for
833 memory bandwidth profiling and management. In *41st IEEE Real-Time Systems Symposium
834 (RTSS 2020)*, Houston, TX, USA, Dec. 2020.
- 835 31 ST Microelectronics Inc. Real-time performance using FIQ interrupt handling in SPEAr
836 MPUs, January 2010. Accessed on 10.01.2020.
- 837 32 M. Solieri T. Kloda, R. Mancuso, N. Capodieci, P. Valente, and M. Bertogna. Deterministic
838 Memory Hierarchy and Virtualization for Modern Multi-Core Embedded Systems. In *25th
839 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2019)*, pages
840 1–14, Montreal, Canada, April 2019. doi:10.1109/RTAS.2019.00009.

- 841 33 H. Usui, L. Subramanian, K. Chang, and O. Mutlu. Dash: Deadline-aware high-performance
842 memory scheduler for heterogeneous systems with hardware accelerators. *ACM Transactions*
843 *on Architecture and Code Optimization (TACO)*, 12(4):1–28, 2016.
- 844 34 P. Valsan and H. Yun. MEDUSA: A predictable and high-performance DRAM controller
845 for multicore based embedded systems. In *2015 IEEE 3rd International Conference on*
846 *Cyber-Physical Systems, Networks, and Applications*, pages 86–93. IEEE, 2015.
- 847 35 S. K. Venkata, I. Ahn, D. Jeon, A. Gupta, C. Louie, S. Garcia, S. Belongie, and M. B. Taylor.
848 SD-VBS: The san diego vision benchmark suite. In *2009 IEEE International Symposium on*
849 *Workload Characterization (IISWC)*, pages 55–64, 2009.
- 850 36 Xilinx. Integrated Logic Analyzer v6.2 LogiCORE IP Product Guide. Technical report,
851 2016. URL: [https://www.xilinx.com/support/documentation/ip_documentation/ila/v6_](https://www.xilinx.com/support/documentation/ip_documentation/ila/v6_2/pg172-ila.pdf)
852 [2/pg172-ila.pdf](https://www.xilinx.com/support/documentation/ip_documentation/ila/v6_2/pg172-ila.pdf).
- 853 37 Xilinx. Zynq UltraScale+ Device Technical Reference Manual. Technical re-
854 port, 2019. URL: [https://www.xilinx.com/support/documentation/user_guides/](https://www.xilinx.com/support/documentation/user_guides/ug1085-zynq-ultrascale-trm.pdf)
855 [ug1085-zynq-ultrascale-trm.pdf](https://www.xilinx.com/support/documentation/user_guides/ug1085-zynq-ultrascale-trm.pdf).
- 856 38 Xilinx, Inc. Zynq UltraScale+ MPSoC - All Programmable Heterogeneous MPSoC, August
857 2016. Accessed on 09.01.2020. URL: [https://www.xilinx.com/products/silicon-devices/](https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html)
858 [soc/zynq-ultrascale-mpsoc.html](https://www.xilinx.com/products/silicon-devices/soc/zynq-ultrascale-mpsoc.html).
- 859 39 M. Xu, L. T. X. Phan, H. Choi, Y. Lin, H. Li, C. Lu, and I. Lee. Holistic resource allocation
860 for multicore real-time systems. In *2019 IEEE Real-Time and Embedded Technology and*
861 *Applications Symposium (RTAS)*, pages 345–356, 2019. doi:10.1109/RTAS.2019.00036.
- 862 40 H. Yun, W. Ali, S. Gondi, and S. Biswas. BWLOCK: A Dynamic Memory Access Control
863 Framework for Soft Real-Time Applications on Multicore Platforms. *IEEE Transactions on*
864 *Computers*, 66(7):1247–1252, 2017.
- 865 41 H. Yun, R. Mancuso, Z. P. Wu, and R. Pellizzoni. Palloc: DRAM bank-aware memory
866 allocator for performance isolation on multicore platforms. In *2014 IEEE 19th Real-Time*
867 *and Embedded Technology and Applications Symposium (RTAS)*, pages 155–166, 2014. doi:
868 10.1109/RTAS.2014.6925999.
- 869 42 H. Yun, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha. MemGuard: Memory bandwidth
870 reservation system for efficient performance isolation in multi-core platforms. In *2013 IEEE*
871 *19th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 55–64,
872 2013.
- 873 43 Y. Zhou and D. Wentzlaff. MITTS: Memory inter-arrival time traffic shaping. *ACM SIGARCH*
874 *Computer Architecture News*, 44(3):532–544, 2016.