

# Static Hardware Partitioning on RISC-V: Shortcomings, Limitations, and Prospects

Ralf Ramsauer\*, Stefan Huber\*, Konrad Schwarz†, Jan Kiszka† and Wolfgang Mauerer\*†

\*{*ralf.ramsauer, stefan.huber, wolfgang.mauerer*}@oth-regensburg.de

*Technical University of Applied Sciences Regensburg, Germany*

†{*konrad.schwarz, jan.kiszka*}@siemens.com

*Siemens AG, Corporate Research, Munich, Germany*

**Abstract**—On embedded processors that are increasingly equipped with multiple CPU cores, static hardware partitioning is an established means of consolidating and isolating workloads onto single chips. This architectural pattern is suitable for mixed-criticality workloads that need to satisfy both, real-time and safety requirements, given suitable hardware properties.

In this work, we focus on exploiting contemporary virtualisation mechanisms to achieve freedom from interference respectively isolation between workloads. Possibilities to achieve temporal and spatial isolation—while maintaining real-time capabilities—include statically partitioning resources, avoiding the sharing of devices, and ascertaining zero interventions of superordinate control structures.

This eliminates overhead due to hardware partitioning, but implies certain hardware capabilities that are not yet fully implemented in contemporary standard systems. To address such hardware limitations, the customisable and configurable RISC-V instruction set architecture offers the possibility of swift, unrestricted modifications.

We present findings on the current RISC-V specification and its implementations that necessitate interventions of superordinate control structures. We identify numerous issues adverse to implementing our goal of achieving zero interventions respectively zero overhead: On the design level, and especially with regards to handling interrupts. Based on micro-benchmark measurements, we discuss the implications of our findings, and argue how they can provide a basis for future extensions and improvements of the RISC-V architecture.

## 1. Introduction

To reduce hardware costs and the overall complexity of the increasing amount of system-on-chips (SoCs) shipped in many mass-products for industry and consumers, consolidation of such systems onto a single chip [1] is desired. This also applies to high- and mixed-criticality systems that at least partially contain safety and real-time critical components. To ensure safety and security, uncontrolled interference between components must be avoided.

In general, isolation is required to guarantee non-interference between components. A proven means for the latter is (a) static partitioning [2, 3, 4] of resources (*i.e.*,

no sharing of devices or cores between isolated components), and (b) zero interventions/overhead caused by the superordinate control structure (*e.g.*, hypervisor) during operation [3]. This prevents a domain from inadvertently violating restrictions imposed by the control structure, and thus from interfering with other critical components.

By exploiting virtualisation technologies of modern CPUs to achieve the required isolation, we can obtain isolation guarantees directly from the hardware. This allows for independently running general purpose (GPOS) and real-time (RTOS) operating systems, as well as bare-metal applications on a SoC.

The absence of hypervisor activity is almost a necessary precondition for freedom from interference [5] with regards to real-time capabilities and decreases the amount of hypervisor code requiring validation or certification as needed for safety-critical environments.

During the implementation of our hypervisor *Jailhouse* [3]—which strives for zero-overhead virtualisation—on x86 [6] and ARM [7], we have learned that real-world issues of hardware devices are often underestimated. Yet they counter our design goal of eliminating virtualisation overheads. As safety-critical real-time systems are only a subordinate aspect in the revenue of chip vendors, it is hard to convince them to implement required architectural changes. On customisable architectures, such as ARM, substantial license fees may occur.

Given that the free and open RISC-V architecture supports hardware-level virtualisation [8], it is in an appropriate initial design stage [9], and may be a suitable candidate for Hardware/Software Co-Design [10] activities that aim at zero-overhead virtualisation, especially since custom designs do not need to deal with conflicting requirements imposed by 3rd parties in control of the architectural specification.

To identify limitations of the current RISC-V hardware specifications detrimental to our design goal, we have ported the Jailhouse hypervisor to this emerging architecture. This sheds light on shortcomings and limitations of the architecture, and gives prospects on possibly required improvements for zero-trap virtualisation.

The rest of this paper is structured as follows: In Section 2, we present related work. Afterwards, we briefly introduce properties of the hypervisor extensions of the

RISC-V architecture, and architecture and rationale of the Jailhouse hypervisor in Section 3. In Section 4, we describe a set of fundamental micro-benchmarks. We quantify the impact of virtualisation overhead on the NOEL-V platform, a synthesisable VHDL model of a six-core 64-bit processor that implements the RISC-V architecture, and supports hypervisor extensions [11]. We discuss our findings in Section 5, and conclude in Section 6.

## 2. Related Work

Before the final ratification of the hypervisor extensions for the RISC-V architecture end of 2021, Pinto *et al.* implemented the preliminary draft of the extension (v0.61) on a FPGA [9]. As their focus is on implementing the hypervisor extensions and customising hardware components for virtualisation, they left out time-consuming hardware optimisations, such as optimisations for address translations. They ported their hypervisor *Bao* to RISC-V to identify limitations and factors of latency overhead, and benchmark their implementation. As they identified interrupt handling as a major issue, they optimised hardware design with respect to eliminating the necessity of hypervisor activity for guest interrupt delivery. This results in major latency decrease for IRQ handling. The RISC-V community is still in the process of researching the implications of trap-free wired interrupt handling for virtual machines [12]. Performance was tested on Firesim [13], a cycle-exact simulator running at 3200MHz simulation clock, using *mibench embedded benchmark suite: automotive subset*, which is frequently used in this area. The benchmarks compare the performance of bare-metal vs. their enhanced hypervisor implementation. They further investigated performance deviations when using cache colouring mechanisms, a technique that can be used for decreasing inter-virtual-machine interference.

In contrast to Pinto *et al.*, we conduct our benchmarks and measurements on real hardware, a synthesised RISC-V processor on a Xilinx Virtex UltraScale+ VCU118 FPGA. We therefore use the final ratified specification of the hypervisor extensions.

In [14], Caforio *et al.* present VOSySmonitoRV, a mixed-criticality solution that aims for implementing static hardware partitioning on RISC-V, by intentionally abstaining from hypervisor extension. They justify their design decision with lack of hardware that supports virtualisation extensions. In contrast, they exploit the machine mode privilege level of the platform to implement means for partitioning, such as memory area isolation or interrupt delivery. As hypervisor extensions are ratified now, future hardware implementations will likely optimise for ratified extensions rather than for custom exploitation of the machine mode privilege level.

## 3. Architecture

### 3.1. Jailhouse: Concepts and Rationale

Safely running real-time workloads of mixed criticality on multi-core systems [15] next to Linux is a common

industrial requirement in many domains. Contemporary multi-core platforms typically feature more CPU cores—Hardware Threads (HARTs) in RISC-V language—than workloads, and critical tasks can be exclusively assigned to dedicated, isolated CPU cores. Linux, together with its feature-rich ecosystem, can then execute uncritical tasks on the remaining CPU cores.

Embedded virtualisation is a promising approach for implementing safe isolation of different workloads. Execution domains, including Linux, run as guests of a hypervisor. This approach is, for example, implemented by XtratuM[16], NOVA [17], and PikeOS [18].

Static hardware partitioning is a special case of embedded virtualisation; it *exclusively* assigns hardware resources to compute domains. Exclusive assignment of hardware resources includes exclusive assignment of physical CPU cores to logical domains. Hence, static hardware partitioning assumes that available computational resources exceed required computational power. Consequently, no scheduler is required by the hypervisor, which avoids scheduling overhead. Virtualisation extensions ensure safe cross-domain isolation.

Our approach is based on Jailhouse, a thin Linux-based partitioning hypervisor that targets real-world systems. Motivated by the exokernel concept [19], our aim is to reduce the hypervisor to a minimum level of abstraction. Our goal is to minimise the hypervisor’s interaction with guests, with the intention of preserving key quality parameters of any guest software regardless of if it is executed natively, or under the presence of a control structure. With this approach, guests inherit real-time guarantees of the underlying hardware by design. Besides unavoidable hardware overhead due to the virtualisation of the system (*e.g.*, second level page table translation [20]), no further software-induced overhead due to the existence of a VMM occurs during operation.

A small code base is a precondition for certifiability for critical environments. The reduction of guest interaction ensures the maintenance of the platform’s real-time capabilities by design—if no interceptions take place, the hypervisor cannot introduce increased latencies.

Running Linux in uncritical partitions of the system is a requirement for many real-world use cases. Therefore, we partition a booted Linux system, instead of booting Linux on a partitioned system. This offloads complex hardware initialisation to Linux, and ensures a small code base of the hypervisor, as only a few platform specific drivers are required (during the operational phase, Linux is lifted into the state of a virtual machine).

To create new isolated domains, specific hardware resources (*e.g.*, CPUs, memory, peripheral devices) are offlined and removed from Linux. The hypervisor is called to create a new domain has raw access to these resources. Secondary real-time operating systems, including Linux, or even bare-metal applications can be loaded into the domains. Jailhouse does not paravirtualise any resources as it exclusively assigns resources to computing domains.

The hypervisor shall only be active during its boot phase (the initialisation of the hypervisor) and during the partitioning phase (creation, initialisation and boot of new

domains). During the operational phase (system is partitioned, and all partitions are running), the goal is no further action by the hypervisor.

## 3.2. RISC-V Platform Virtualisation

**3.2.1. Virtualisation Architecture.** While the RISC-V platform is designed to be fully virtualisable even without dedicated virtualisation extensions (*i.e.*, via trap-and-emulate mechanisms), the hypervisor extension, which allows for executing most instructions of virtual guests natively, has recently been ratified by *RISC-V International*. RISC-V implements three basic privilege modes: (1) Machine-Mode (M-Mode) where usually the Supervisor Binary Interface (SBI)—a BIOS-like firmware—resides, (2) Supervisor-Mode (S-Mode), typically used for the privileged operating system and (3) User-Mode (U-Mode) for unprivileged user-level applications. When the hypervisor extension is active, the S-Mode is utilised by the hypervisor and called *Hypervisor extended-Supervisor (HS-Mode)*. Guests run in Virtualised Supervisor (VS-Mode), which provides shadows of key registers to minimise interventions by the hypervisor.

**3.2.2. Memory Management.** The memory management unit (MMU) is virtualisation aware: page tables are resolved transparently for guests using a second translation stage for guest physical memory to host physical memory conversion. No hypervisor intervention is needed for page table walks and modifications. The two-stage address translation process [21] does reduce performance, especially on TLB misses. However, TLB misses can be reduced by using huge pages in the second *G-Stage* translation level. As the MMU counterpart for IO-devices (IOMMU) is still under specification process [22], its desired memory protection features and virtualisation capabilities are lacking for devices that use direct memory access (DMA) features. This makes direct assignment of such devices to a guest—via techniques like interrupt remapping—impossible, and guests cannot use such devices without heavy hypervisor intervention.

**3.2.3. Interrupt Controller.** An architectural weakness of current generation RISC-V is the *Platform Level Interrupt Controller (PLIC)*, which is the first generation standard interrupt controller.

The RISC-V hypervisor extension defines an *interrupt pass-through* mechanism. The intention is to allow interrupt requests to raise exceptions in guests without mediation by the hypervisor. As such, this feature is highly desirable for a hypervisor focused on the real-time domain.

Unfortunately, having been developed earlier than the hypervisor extension, neither the *Core Local Interrupt Controller (CLINT)* nor PLIC allow for direct interrupt remapping to the guest. Any interrupt (timer, software, external) therefore first arrives at the hypervisor (timer and software interrupts even make another detour through M-mode), before having to be injected into the targeted domain. Due to further design misconceptions of the interrupt controllers, additional hypervisor intervention is needed for the guest to mark

an interrupt as being handled (*claim*) respectively handled (*complete*). This is because registers associated with claim and complete are memory mapped for multiple HARTs on the same memory page, which means we cannot rely on the MMU for access control. Therefore a superordinate control structure is needed to prevent a cell from (un-)intentionally interfering with any other cell’s interrupts. These problems heavily affect interrupt latency and thus real-time capabilities.

**3.2.4. Hyperthreading.** As RISC-V does not support hyperthreading, there are way less possibilities for malicious inter-cell interference than on platforms like Intel (*e.g.* spectre, meltdown). However, last level caches are shared and not yet partitionable, which opens up the possibility for influencing other cell latencies via cache pollution [9].

## 4. Evaluation

### 4.1. Benchmarking Setup

To test performance implications of the hypervisor overhead on *real hardware*, we use the *Xilinx Virtex UltraScale+ VCU118 FPGA*, using the NOEL-V [23] bitstream, which is a synthesisable VHDL model of a RISC-V processor that implements Hypervisor Extensions. H-Extensions and the interrupt controllers follow the final, ratified specifications. While there is an open-source bitstream available, we used the commercial one, which supports performance optimisations and L1 and last level caches (LLCs). The NOEL-V has six HARTs, each of which has a dedicated L1 cache, while sharing a common LLC. HARTs and caches run at 100MHz. For static hardware partitioning, we use Jailhouse as hypervisor.

We perform micro-benchmarks to quantify additional overheads resp. latencies due to the existence of the hypervisor. All micro-benchmarks are conducted in the following measurement scenarios:

- (A) As bare-metal application without an underlying hypervisor,
- (B) with Jailhouse in a static partitioned execution domain (parallel to Linux),
- (C) As (B), but with additional load in the Linux partition.

In scenario (A), we measure the *baseline* of the raw system, that is, overheads and latencies without the existence of a hypervisor. (A) represents the *raw noise* of the platform that we cannot fall below. Scenario (B) represents the base overhead that exists due to the existence of the hypervisor. Finally, scenario (C) simulates conditions in a real asymmetric multiprocessing (AMP) environment: arbitrary load on neighbouring execution domain to stress shared system components, such as caches or system buses.

For our micro-benchmarks, we implemented our own minimalist operating system, which is publicly available as Open Source Software.<sup>1</sup>

1. Refer to <https://github.com/lfd/grinch>. We call it the *Grinch*, as it benchmarks NOEL-V, which—apart from being a RISC-V implementation—is also French for *Christmas*.




Figure 1. Illustration of the cross-systems code path for our benchmarks. Dots/squares mark traps; squares are unavoidable, while dots arise from HV interaction. For the IPI round trip measurement (teal), the path is also traversed in backward direction, as indicated by the loop. The additional load in the Linux domain to perturb the measurement is optional, and only generated in scenario (C).

We selected micro-benchmarks to measure relevant code paths where the hypervisor has to intervene active in typical real-time scenarios, such as cyclic timer interrupts, IPIs, external interrupts and frequent firmware calls, such as those used in RISC-V for remote fences.

## 4.2. Benchmark #1—IRQ Reinjection

As mentioned before, any IRQ on RISC-V is received in S-mode and re-injected into VS-mode. Basically, there are 3 types of IRQs: Timer, IPI, and External interrupts (peripheral devices). External IRQs are managed by the interrupt controller (*i.e.*, PLIC).

In our first benchmark, we investigate timers (shown in Figure 1 in dotted ochre), as they do not need interaction with the PLIC, but still need to be injected by the hypervisor. Further, while reading the current timer value can be done without hypervisor interaction, programming the timer requires interaction with the SBI [24], which results in moderation by the hypervisor. Any SBI call must be moderated by the hypervisor to ensure that the call has no cross-domain effects (*e.g.*, CPU offlining, which is conducted via SBI, must not affect a neighbored domain). Typically, the overhead that is required for setting the timer only plays a subordinate role, as the time that is required to set the timer vs. its expiration time are significantly apart. However, for the sake of completeness, we quantify all overheads in Fig. 2.

The essential measurement is the timer *jitter*: The difference between scheduled and actual arrival time of the timer IRQ in (V)S-Mode. In a virtualised scenario, the hypervisor receives the timer (2)—again via detour through the SBI (3)—and directly injects it by setting the corresponding pending-bit (2)–(1). When the timer arrives, our benchmark will set the next timer expiration time to a point in future. This will automatically clear the pending flag [24].

## 4.3. Benchmark #2—IPI Round Trip Time

As IPI Round Trip Time (RTT), shown in Figure 1 in solid teal, we define the time that is required for sending an IPI to a secondary target HART, and back. The only task of the target is to send the IPI back to the initial sender. We chose this measurement, as IPIs are frequently used by operating systems for signalling and synchronisation purposes in real-time contexts.

On RISC-V, IPIs are raised on the platform via SBI, where the target is specified as an argument. The SBI call must be intercepted by the hypervisor (2) as the domain membership of the target must be verified. After verification, the IPI is propagated to the firmware (3), where it is finally sent. From now on, the sender actively polls for the returning IPI. On receiver side, the IPI first arrives at the hypervisor (6)—again via detour through SBI—which injects the IPI into the guest (7)–(8) by setting the appropriate pending bit. The guest software actively polls on the pending bit, so once the guest sees the IPI, it sends an IPI back to the sender. The same path is traversed backwards again: Moderation of the IPI, arrival at the sender in the hypervisor, re-injection.

In total, four hypervisor interceptions are required for the IPI RTT measurement: moderation for sender, arrival at receiver, moderation for receiver, arrival at sender.

## 4.4. Benchmark #3—PLIC Emulation

The PLIC interrupt controller offers no virtualisation possibilities. Furthermore, the memory layout of the PLIC is unfavourably organised (*e.g.*, cross-hart configuration interfaces reside on the same memory page). This requires that accesses to the PLIC must be completely emulated.

The PLIC processes arriving external IRQs as follows:

- Physical arrival: set the external IRQ pending bit: (a)
- Interruption of S-mode: (b)
- *Claiming* the IRQ (*i.e.*, read from PLIC register): (c)–(d)
- *Acknowledgement (complete)* of the IRQ (*i.e.*, write to a PLIC register): (e)–(f).

Under the presence of a hypervisor, the IRQ, shown in Figure 1 in dash-dotted black, first arrives in HS-Mode. The hypervisor re-injects the external interrupt to its guest, which will be interrupted.<sup>2</sup> The guest claims the IRQ by reading the PLIC claim/complete register, which requires hypervisor moderation, as well as the acknowledgement of the IRQ. The time required for moderation can be found in Fig. 2.

## 4.5. Benchmark #4—Synchronous Traps

Synchronous traps, shown in Figure 1 in dashed grey, arise when certain privileged instructions are executed from less privileged modes like VS-mode. The processor traps into higher privileged modes, where the instructions are handled (*i.e.*, for permission checks). We measure the overhead of the remote fence (*rfence*) firmware call, which is frequently used

2. This is the first trap, comparable with the arrival of a timer interrupt.




Figure 2. Measurement Results (notice the double logarithmic axes): CPU cycles taken for our benchmarks—comparing performance bare-metal vs. with hypervisor (with and without load on other cores).

to enforce ordering constraints on memory operations. It has to detour through the SBI (3). The call from VS-Mode to SBI is moderated by the hypervisor by trapping into S-Mode (2). We measure synchronous traps cycle-precise using the *rdcycle* instruction before and after the trap.

## 5. Discussion

In agreement with findings by Pinto *et al.* [9], our measurements show that real-time oriented virtualisation on RISC-V—without hardware optimisations—comes at a huge cost. The highest cost is the unavoidable virtualisation of the PLIC interrupt controller. Regarding complete PLIC moderation, for example, access to the register that only requires three cycles on bare metal may grow up to 16,000 cycles when the neighbouring Linux domain is under load.

Another source of unnecessary overhead and complexity is the design of software interrupts usually used for IPIs via the CLINT. Interrupt injection and handling has to take a detour via both, S and M-Mode. This causes unnecessary mode-changes, and thus increases latency. Additionally, IPIs cannot be distinguished: An IPI is implemented as doorbell interrupt, and unlike with other platforms (*e.g.*, ARM), it is not possible to directly obtain an *IPI number*. This necessitates storing and reading the required information at a shared memory location, which causes additional overhead.

Naturally, there are unavoidable sources of virtualisation overhead. Firstly, cell management (*i.e.*, creating, starting, stopping and destroying cells) involves hypervisor activity, yet only during system initialisation and not during the operational phase. Secondly, additional translation for guest virtual to guest physical addresses causes temporal overhead, even if it does not trigger traps. A two-stage translation

process can double TLB pressure, which is another source of overhead. However, it can be mitigated by using huge pages for large subsequent memory areas in the second stage.

A shortcoming of the RISC-V architecture is the current generation interrupt controller PLIC. The RISC-V hypervisor extension defines an *interrupt pass-through* mechanism with the intention to allow interrupt requests to raise exceptions in guests without intermediation by the hypervisor. As such, this feature is highly desirable for a hypervisor focused on the real-time domain, having been developed earlier than the hypervisor extension, the PLIC does not support this feature.

To handle an external interrupt, three traps in the hypervisor are necessary. This situation is supposed to be improved by the RISC-V *Advanced Interrupt Architecture* (AIA) [12], which is currently under specification. It addresses current limitations, and includes optimisations for virtualisation. Both, the PLIC and the CLINT will be superseded by advanced versions (APLIC respectively ACLINT), and an additional interrupt controller, the *Incoming Message-Signalled Interrupt Controller* (IMSIC). One such controller to handle MSI interrupts will be available per HART. MSIs will be generated by writing into the IMSIC’s dedicated *files* for M/S/VS-contexts.

An IOMMU is currently under specification, and will allow for direct device assignment into guests. Device MSIs can be directly (*i.e.*, without hypervisor intervention) forwarded to VS-mode with the IOMMU by writing into a dedicated VS-file of the corresponding IMSIC of the target HART.

The APLIC however, will only support virtualisation partly, and wired interrupt handling will still require trap-and-emulate mechanisms (*APLIC direct-mode*), as the RISC-V community still discusses about the need and implications of directly forwarding these kind of interrupts directly into



guests without hypervisor supervision [12]. However, on systems that implement both, APLIC and IMSIC(s), the APLIC can be configured to translate wired interrupts into MSIs, which enables direct forwarding into VS-mode in conjunction with an IMSIC (*APLIC MSI-mode*).

The IMSIC can also be used as an alternative to the the (A)CLINT for sending IPIs between virtual HARTs respectively *cells*, yet without hypervisor intervention, again by writing directly into the VS-file of the targeted HART. Although the ACLINT will feature a separate device for S-Mode software interrupts (*SSWI*) to allow direct sending of IPIs without detour into M-Mode, this will be still not possible for VS-mode respectively virtual machines.

APLIC and ACLINT will not be fully optimised for virtualisation, yet most of the problems can mostly be handled with trickery in conjunction with the IMSIC.

## 6. Conclusion

Our work shows that static hardware partitioning on contemporary RISC-V hardware comes at a high cost. Targeted hardware optimisations for reducing hypervisor activity have high potential for eliminating overheads in static partitioned scenarios. Changes to the interrupt architecture (*e.g.*, direct interrupt handling) influence desired performance parameters more than state-of-the-art improvements of traditionally well optimised mechanisms (*e.g.*, MMU 2-Stage Translation). This underlines, once more, the importance of Hardware/Software Co-Design as an important means for swift design and implementation of useful hardware enhancements. The RISC-V ecosystem is an optimal environment for such endeavours.

Although the use-cases presented here are tailored—and crucial—to embedded systems, the implications of our measurements are valid for throughput-oriented general purpose systems, with significant possible performance gains.

In future work, we will investigate the advanced interrupt architecture of the RISC-V ecosystem.

## References

- [1] Manfred Broy. “Challenges in Automotive Software Engineering”. In: *28th ICSE*. ACM Press, 2006.
- [2] José Martins, Adriano Tavares, Marco Solieri, Marko Bertogna, and Sandro Pinto. “Bao: A lightweight static partitioning hypervisor for modern multi-core embedded systems”. In: *NG-RES*. 2020.
- [3] Ralf Ramsauer, Jan Kiszka, Daniel Lohmann, and Wolfgang Mauerer. “Look Mum, no VM Exits! (Almost)”. In: *Proc. of the 13th OSPERT*. 2017.
- [4] Hao Li, Xuefei Xu, Jinkui Ren, and Yaozu Dong. “ACRN: a big little hypervisor for IoT development”. In: *Proc. of the 15th ACM SIGPLAN/SIGOPS Int. Conf. on VEE*. 2019.
- [5] *ISO 26262: Road vehicles – Functional safety*. International Organization for Standardization.
- [6] Rich Uhlig, Gil Neiger, Dion Rodgers, et al. “Intel virtualization technology”. In: *Computer* (2005).
- [7] Prashant Varanasi and Gernot Heiser. “Hardware-supported virtualization on ARM”. In: *Proc. of the 2nd APSys*. 2011.
- [8] Andrew Waterman, Krste Asanović, and Jon Hauser. *The RISC-V Instruction Set Manual Volume 2: Privileged Architecture Version 20211203*. Tech. rep. University of California at Berkeley, 2021.
- [9] Bruno Sá, José Martins, and Sandro Emanuel Salgado Pinto. “A first look at RISC-V virtualization from an embedded systems perspective”. In: *IEEE Transactions on Computers* (2021).
- [10] Wayne H. Wolf. “Hardware-Software Co-Design of Embedded Systems”. In: *Proc. of the IEEE* (1994).
- [11] J. Andersson. “Development of a NOEL-V RISC-V SoC Targeting Space Applications”. In: *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. IEEE Computer Society, 2020.
- [12] *RISC-V Advanced Interrupt Architecture Specification Draft*. 2022.
- [13] Sagar Karandikar, Howard Mao, Donggyu Kim, et al. “FireSim: FPGA-accelerated cycle-exact scale-out system simulation in the public cloud”. In: *Proc. of the ACM/IEEE 45th ISCA*. IEEE, 2018.
- [14] Flavia Caforio, Pierpaolo Iannicelli, Michele Paolino, and Daniel Raho. “VOSySmonitoRV: a mixed-criticality solution on Linux-capable RISC-V platforms”. In: *Proc. of 10th MECO*. IEEE, 2021.
- [15] Lukas Bulwahn. “Is Linux Kernel Development Good Enough to Make Your Life Depend on it? Progress on Procedures & Methods to Qualify the Linux Kernel Development Process”. In: *ELCE17*. 2017.
- [16] Alfons Crespo, Ismael Ripoll, and Miguel Masmano. “Partitioned Embedded Architecture based on Hypervisor: The XtratuM approach”. In: *Proc. of the 8th EDCC*. IEEE, 2010.
- [17] Udo Steinberg and Bernhard Kauer. “NOVA: a microhypervisor-based secure virtualization architecture”. In: *Proc. of the 5th EuroSys*. ACM, 2010.
- [18] Robert Kaiser and Stephan Wagner. “Evolution of the PikeOS microkernel”. In: *First International Workshop on Microkernels for Embedded Systems*. 2007, p. 50.
- [19] Dawson R Engler, M Frans Kaashoek, et al. *Exokernel: An operating system architecture for application-level resource management*. Vol. 29. 5. ACM, 1995.
- [20] Ulrich Drepper. “The Cost of Virtualization: Software Developers Need to Be Aware of the Compromises They Face When Using Virtualization Technology.” In: *Queue* 6.1 (Jan. 2008), pp. 28–35.
- [21] Ulrich Drepper. “The Cost of Virtualization”. In: *Queue* (2008).
- [22] *RISC-V IOMMU Task Group*. 2022.
- [23] Cobham Gaisler. *NOEL-V Processor*. 2022.
- [24] *RISC-V SBI Specification*. 2022.