

Predictable Computer Architecture

The Foundation of Hard Real-Time Systems

Antti Nurmi
Tampere University
antti.nurmi@tuni.fi

Outline

- ▶ Terminology
- ▶ Unpredictable Computers
- ▶ Our Research

Terminology

What are Real-Time Systems?

(1/5)

- ▶ Computing systems with **explicit timing requirements**.

What are Real-Time Systems?

(2/5)

- ▶ Consider:
 - “System X should process as many FPS as possible.”
- vs.
 - “System Y must respond to incoming event Z within 10 μ s.”

What are Real-Time Systems?

(3/5)

- ▶ Two broad categories for seriousness of the system timing requirements [1]:
 - Soft:
 - Missing deadlines degrades performance.
 - Inconvenient rather than fatal.
 - E.g.: Video streaming for **sports events**.

What are Real-Time Systems?

(4/5)

- ▶ Hard:
 - Missing deadlines is a system failure.
 - Expensive things break, people die.
 - E.g.: Video streaming for **remote brain surgery**.

What are Real-Time Systems?

(5/5)

- ▶ Related: Mixed-Criticality Systems.
- ▶ Consider an airplane as one system:¹
 - Flight controls
 - In-flight information



¹Daniel Schwen, CC BY-SA 4.0 <<https://creativecommons.org/licenses/by-sa/4.0/>>, via Wikimedia Commons

Real-Time Operating Systems (RTOS)

(1/2)

- ▶ A category of operating system specifically designed for implementing real-time systems.
 - Generally more lightweight than full general-purpose operating systems (GPOS), closer to bare-metal.
 - Examples: Zephyr, FreeRTOS, many proprietary solutions.
 - **Linux is not an RTOS** [2].

- Typical RTOS utilities:
 - Task/thread creation & management
 - (Preemptive) context switching
 - Prioritized scheduling
 - Synchronization & shared resource control
 - Mutexes, locks, semaphores, etc.
 - Timing utilities
 - Timers, SysTicks
 - Inter-process communication (IPC)

Unpredictable Computers

What makes a computer (un)predictable? (1/2)

- ▶ Fundamentally, digital computers are perfectly deterministic.
- ▶ The most significant advances in computer architecture aim to improve **average-case** performance.
 - Caches, multicores, branch prediction, etc.

What makes a computer (un)predictable? (2/2)

- ▶ Real-time systems need to be designed around **worst-case** performance and predictable timing behaviour [3].
- ▶ Many average-case optimizations are directly detrimental to worst-case performance.

Example: Memory Hierarchies

(1/5)

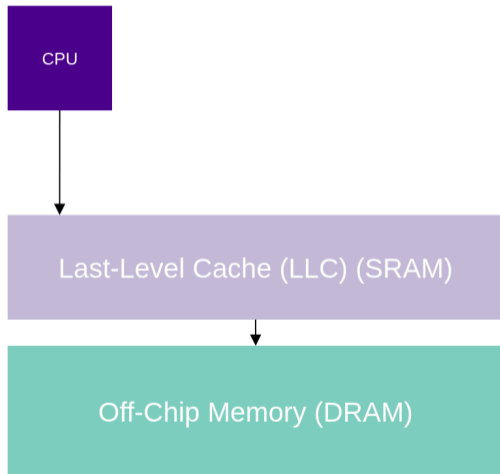
- ▶ **How long does a generic CPU memory access take?**
- ▶ Assume an SRAM access latency of 3 cycles.
- ▶ Here: 3 cycles, **always**.



Example: Memory Hierarchies

(2/5)

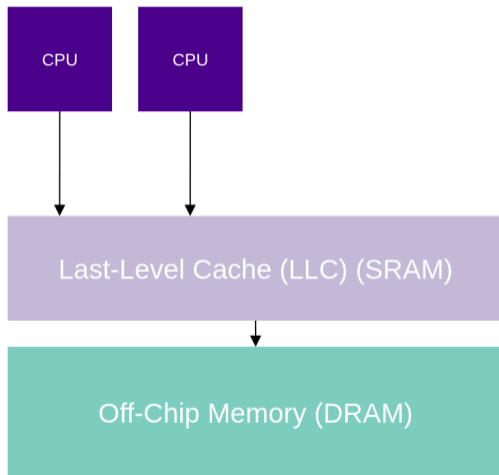
- ▶ What if your system needs a lot of memory? (e.g, for running Linux?)
- ▶ Assume DRAM accesses (inc. LLC refill) take from 1000 to 5000 cycles.
- ▶ LLC hit: 3 cycles
- ▶ LLC miss: 1000-5000 cycles



Example: Memory Hierarchies

(3/5)

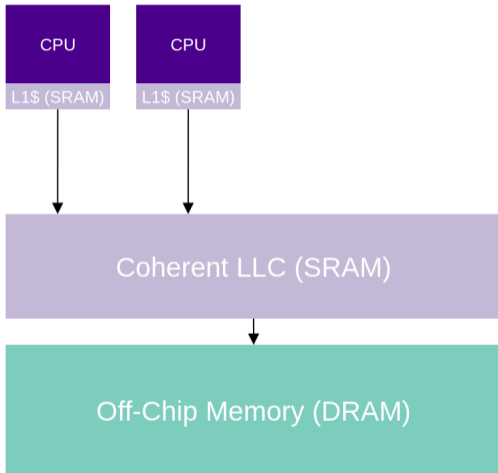
- ▶ What about multicores?
- ▶ Similar latency, but more cache misses because of conflicting access patterns.
- ▶ Hardware resource contention can be catastrophic for time-predictability [4].



Example: Memory Hierarchies

(4/5)

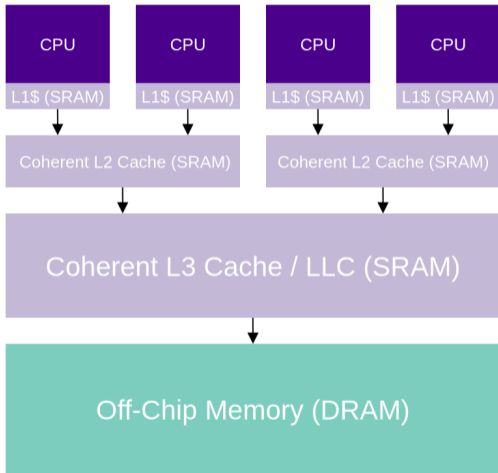
- ▶ Adding L1 caches for better cache hit ratio.
- ▶ Coherency required from LLC.
- ▶ Access times now:
 - L1 hit: 3 cycles
 - LLC hit, L1 refill: 100 cycles
 - LLC miss, refill: 1000-5000 cycles



Example: Memory Hierarchies

(5/5)

- ▶ Adding more cores, more cache layers.
- ▶ Access times now:
 - L1 hit: 3 cycles
 - L2 hit, L1 refill: 100 cycles
 - LLC hit, L2 refill: 300 cycles
 - LLC miss, refill: 1000-5000 cycles



Example: Branch Prediction

(1/2)

- ▶ Long instruction pipelines depend on **branch prediction** to maintain performance.
- ▶ Can be extremely simple or extremely complex.
 - Simple: Always guess yes, ~70 % accuracy.
 - Complex: TAGE [5], neural networks [6], > 95 % accuracy.
- ▶ **But**, branch prediction is inherently probabilistic.

Example: Branch Prediction

(2/2)

- ▶ Proving time-predictability for systems with branch prediction/speculative execution is an active research area [7].

$$\begin{aligned}
 &pre \sqsubseteq_S PC \sqsubseteq_S IF \sqsubseteq_S ID \sqsubseteq_S IS \sqsubseteq_S \{ALU, MUL_1, LSU, CSR, DIV\} \sqsubseteq_S \{MUL_2, LU, SU\} \sqsubseteq_S CO \sqsubseteq_S ST \sqsubseteq_S post \\
 &cycle(c)(i) := \begin{cases} (c.nstg(i), c.nlat(i)) & : c.ready(i) \wedge c.free(c.nstg(i)) \\ (c.stg(i), c.nent(i)) & : otherwise \end{cases} & c.isnext(s, i) := c.stg(i) = s \wedge \forall j < i. c.stg(j) \sqsubseteq_S s \\
 &c.nent(i) := \begin{cases} c.cnt(i) - 1 & : c.cnt(i) > 0 \\ 0 & : otherwise \end{cases} & c.nlat(i) := \begin{cases} memlat_y(i) & : c.nstg(i) = IF \wedge \neg ichit(i) \\ memlat_d(i) & : (c.nstg(i) = LU \wedge \neg dchit(i)) \\ & \vee c.nstg(i) = ST \\ ezlat(i) & : c.nstg(i) = DIV \\ 0 & : otherwise \end{cases} \\
 &c.pending(i, op) := \exists j < i. opc(j) = op \wedge c(j) \sqsubseteq_P (lstg(op), 0) \\
 &c.nstg'(i) := \begin{cases} c.stg(i) \neq pre \wedge \neg c.pending(i, branch) \wedge pwrapng(i) & : post \\ c.nstg'(i) & : otherwise \end{cases} \\
 &c.nstg(i) := \begin{cases} PC & : c.stg(i) = pre \\ IF & : c.stg(i) = PC \\ ID & : c.stg(i) = IF \\ IS & : c.stg(i) = ID \\ LSU & : c.stg(i) = IS \wedge opc(i) \in \{load, store, atomic\} \\ LU & : c.stg(i) = LSU \wedge opc(i) = load \\ SU & : c.stg(i) = LSU \wedge opc(i) \in \{store, atomic\} \\ MUL_1 & : c.stg(i) = IS \wedge opc(i) = mul \\ MUL_2 & : c.stg(i) = MUL_1 \\ DIV & : c.stg(i) = IS \wedge opc(i) = div \\ CSR & : c.stg(i) = IS \wedge opc(i) = csr \\ ALU & : c.stg(i) = IS \wedge opc(i) \notin \{load, store, atomic, mul, div, csr\} \\ CO & : c.stg(i) \in \{ALU, MUL_2, DIV, CSR, LU, SU\} \\ ST & : c.stg(i) = CO \wedge opc(i) \in \{store, atomic\} \\ post & : (c.stg(i) = CO \wedge opc(i) \notin \{store, atomic\}) \vee (c.stg(i) = ST) \end{cases} & lstg(op) := \begin{cases} LU & : op = load \\ ST & : op = store \\ ST & : op = atomic \\ IS & : op = mul \\ DIV & : op = div \\ CO & : op = csr \\ ALU & : op = branch \end{cases} \\
 &c.ready(i) := (c.stg(i) \neq pre \wedge \neg c.pending(i, branch) \wedge pwrapng(i)) \\
 &\quad \vee (c.cnt(i) = 0 \wedge c.isnext(c.stg(i), i)) \\
 &\quad \wedge (c.stg(i) = PC \Rightarrow \{ichit(i) \\
 &\quad \quad \vee (\neg c.pending(i, branch) \wedge \neg c.pending(i, load) \wedge \neg c.pending(i, store) \wedge \neg c.pending(i, atomic))\}) \\
 &\quad \wedge (c.stg(i) = IS \Rightarrow \{opc(i) \notin \{load, store, atomic\} \Rightarrow \neg c.pending(i, csr) \\
 &\quad \quad \wedge (opc(i) \in \{mul, div\} \Rightarrow \neg c.pending(i, div)) \\
 &\quad \quad \wedge (\forall j < i.dep(i, j) \Rightarrow c.stg(j) \sqsubseteq_S CO)\}) \\
 &\quad \wedge (c.stg(i) = LSU \Rightarrow \{opc(i) \in \{store, atomic\} \wedge \neg c.pending(i, atomic) \\
 &\quad \quad \vee (opc(i) = load \wedge (\neg c.pending(i, store) \wedge \neg c.pending(i, atomic))\}) \\
 &c.free(s) := s \in \{ALU, MUL_1, CSR, MUL_2, CO, post\} \\
 &\quad \vee (s \in \{IF, IS, LSU, SU\} \wedge c.slot(s)) \\
 &\quad \vee (s \in \{PC, ID, DIV, LU, ST\} \wedge (\neg \exists j. c.stg(j) = s) \vee (\exists j. c.stg(j) = s \wedge c.ready(j) \wedge c.free(c.nstg(j)))) \\
 &\quad \vee (\exists i. c.stg(i) = s \wedge pwrapng(i) \wedge \neg c.pending(i, branch)) \\
 &c.slot(IF) := (\#\{j | c.stg(j) = IF\} < iq_size) \vee c.free(ID) \wedge \forall j. c.stg(j) = IF \Rightarrow c.cnt(j) = 0 \\
 &c.slot(IS) := \#\{j | IS \sqsubseteq_S c.stg(j) \sqsubseteq_S CO\} < iq_size \vee (\exists j'. c.isnext(CO, j') \wedge c.ready(j') \wedge (opc(j') \in \{store, atomic\} \Rightarrow c.free(ST))) \\
 &c.slot(SU) := \#\{j | opc(j) = store \wedge LSU \sqsubseteq_S c.stg(j) \sqsubseteq_S post\} < sq_size \vee \exists j'. c(j') = (ST, 0) \\
 &c.slot(LSU) := \#\{j | c.stg(j) = LSU\} < mq_size
 \end{aligned}$$

Unpredictable Software

(1/2)

- ▶ Predictable systems require predictable hardware **and** predictable software.
- ▶ **Threads** are the dominant model for concurrent programming.
 - Threads are also notoriously nondeterministic [8].
 - Nicer models are available.


- ▶ We have focused on RTIC [9], [10].
 - Lightweight, minimal runtime overhead.
 - Static guarantees provided by Rust, RTIC itself.
 - Solid formal foundation based on the stack resource policy (SRP) [11].

Our Research (So Far)

Predictable System Design

(1/2)

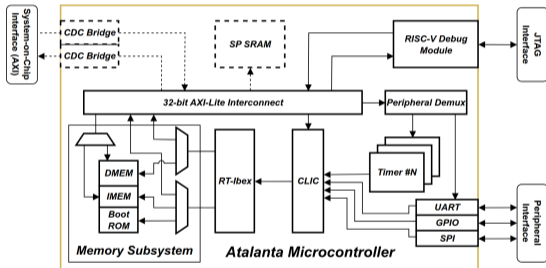
- ▶ Atalanta: our initial work to build a predictable RISC-V microcontroller system.
- ▶ Custom Ibex fork for advanced interrupt controller support, Arm-style ISR fetch.
- ▶ Integrated to (then) brand-new CLIC [12], other open components.



Predictable System Design

(2/2)

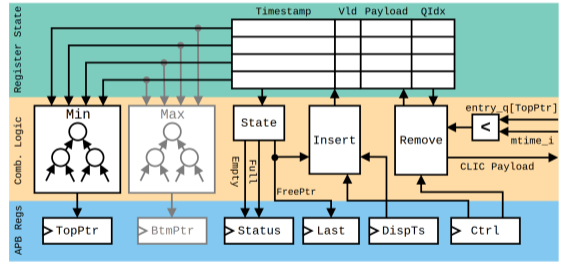
- ▶ Results:
 - Competitive (~21) cycle interrupt latency.
 - Solid platform for subsequent research.
- ▶ Published in [13], poster presented at 2024 European RISC-V Summit in Munich



Hardware Priority Queues

(1/2)


- ▶ Priority queues (PQ) are fundamental to implementing task schedulers.
- ▶ Software-based PQs have a logarithmic time-complexity for at least one of its access operations.
 - A hardware PQ can provide constant access times.



Hardware Priority Queues

(2/2)

- ▶ Initial implementation in [14], improved follow-up in [15].
- ▶ Results:
 - Proof of HW feasibility.
 - Latency & jitter comparison against SW data structures.



Fast & Heterogeneous Interrupts

(1/3)

- ▶ Reactive real-time systems need to do a lot of context switching (CS).
 - Highly repetitive low-level routine.
 - Hardware acceleration has the potential to save a lot of instructions.

```
trampoline:
  addi sp, sp, -(4 * 9)
  sw ra, 0(sp)
  sw t0, 4(sp)
  sw a0, 8(sp)
  sw a1, 12(sp)
  sw a2, 16(sp)
  sw a3, 20(sp)
  sw t1, 24(sp)
  csrr t0, mepc
  csrr t1, mcause
  sw t0, 28(sp)
  sw t1, 32(sp)
  csrsi mstatus, 8 // enable global irqs
  jal functional_isr
  csrwi mstatus, 8 // disable global irqs
  lw t0, 28(sp)
  lw t1, 32(sp)
  csrwi mcause, t0
  csrwi mepc, t0
  lw ra, 0(sp)
  lw t0, 4(sp)
  lw a0, 8(sp)
  lw a1, 12(sp)
  lw a2, 16(sp)
  lw a3, 20(sp)
  lw t1, 24(sp)
  addi sp, sp, (4 * 9)
  mret
```

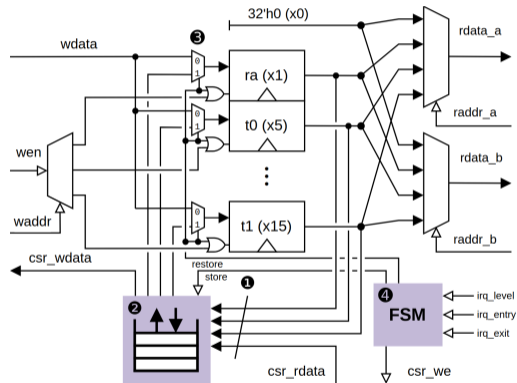
Entry, 12 Instructions

Exit, 13 Instructions

Fast & Heterogeneous Interrupts

(2/3)

- ▶ Our initial work presented a stacked register file to eliminate the need for CS [16].
- ▶ Extension in [17] presented the heterogeneous interrupt (HETI) scheme.
 - Combines accelerated and conventional interrupts.

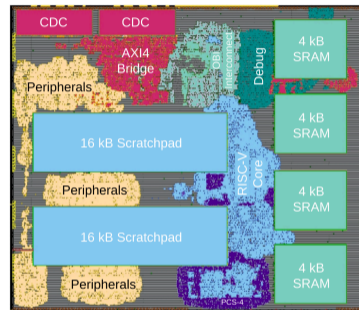


Fast & Heterogeneous Interrupts

(3/3)

► Results:

- Representative ASIC implementation on TSMC 22-nm.
- ~25 % instruction & clock cycle reduction in CS-heavy benchmark.

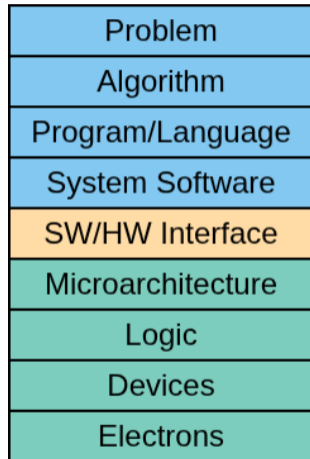


- ▶ Potentially a new RISC-V CPU implementation for better integrating our custom features.
- ▶ More work on hardening and evaluating the presented solutions in ASIC physical design.
 - Hopefully with tapeouts & demonstrators in future projects.

- ▶ New publication on hardware-based earliest-deadline first (EDF) scheduling accepted for publication in 2026 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS).
 - Privileged information for now :)

Interested in Contributing?

- ▶ Broad research field that relates to most layers of the transformation hierarchy.
- ▶ Got inspiration for a related MSc. thesis?
 - Let's have chat!
- ▶ Interested in paid work @TAU?
 - No RA openings **right now**, but hopefully we can hire again very soon.



References

(1/9)

- [1] G. C. Buttazzo, **Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications**, 3rd ed., vol. 24. in Real-Time Systems Series, vol. 24. New York, NY: Springer Nature, 2011.
- [2] F. Reghenzani, G. Massari, and W. Fornaciari, "The Real-Time Linux Kernel: A Survey on PREEMPT_RT," **ACM computing surveys**, vol. 52, no. 1, pp. 1–36, 2019.
- [3] S. A. Edwards and E. A. Lee, "The Case for the Precision Timed (PRET) Machine," in **Proceedings of the 44th Annual Design Automation Conference**, in DAC '07. Asso-

References

(2/9)

- ciation for Computing Machinery, 2007, pp. 264–265. doi: [10.1145/1278480.1278545](https://doi.org/10.1145/1278480.1278545).
- [4] D. Oliveira, W. Chen, S. Pinto, and R. Mancuso, “Shared Resource Contention in MCUs: A Reality Check and the Quest for Timeliness,” in **36th Euromicro Conference on Real-Time Systems (ECRTS 2024)**, R. Pellizzoni, Ed., in Leibniz International Proceedings in Informatics (LIPIcs), vol. 298. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, pp. 5:1–5:25. doi: [10.4230/LIPIcs.ECRTS.2024.5](https://doi.org/10.4230/LIPIcs.ECRTS.2024.5).

References

(3/9)

- [5] A. Seznec and P. Michaud, "A Case for (Partially) Tagged Geometric History Length Branch Prediction," **Journal of Instruction-Level Parallelism**, vol. 8, p. 23, 2006.
- [6] D. A. Jiménez and C. Lin, "Neural methods for dynamic branch prediction," **ACM Trans. Comput. Syst.**, vol. 20, no. 4, pp. 369–397, Nov. 2002, doi: [10.1145/571637.571639](https://doi.org/10.1145/571637.571639).
- [7] A. Gruin, T. Carle, C. Rochange, H. Cassé, and P. Sainrat, "MINOTAuR: A Timing Predictable RISC-V Core Featuring Speculative Execution," **IEEE Transactions on Comput-**

References

(4/9)

- ers**, vol. 72, no. 1, pp. 183–195, 2023, doi: [10.1109/TC.2022.3200000](https://doi.org/10.1109/TC.2022.3200000).
- [8] E. Lee, “The Problem with Threads,” **Computer**, vol. 39, no. 5, pp. 33–42, 2006.
- [9] P. Lindgren, P. Dzialo, and H. Lunnikivi, “Hardware support for Static-Priority Stack Resource Policy based scheduling,” in **2023 IEEE 32nd International Symposium on Industrial Electronics (ISIE)**, 2023, pp. 1–5. doi: [10.1109/ISIE51358.2023.10228088](https://doi.org/10.1109/ISIE51358.2023.10228088).

References

(5/9)

- [10] H. Lunnikivi, Z. Madaoui, P. Dzialo, and P. Lindgren, "Modular RTIC: Lightweight Real Time for Customized Architectures," **IEEE Transactions on Very Large Scale Integration (VLSI) Systems**, vol. 33, no. 11, pp. 2952–2960, 2025, doi: [10.1109/TVLSI.2025.3595712](https://doi.org/10.1109/TVLSI.2025.3595712).
- [11] T. Baker, "A stack-based resource allocation policy for real-time processes," in **[1990] Proceedings 11th Real-Time Systems Symposium**, 1990, pp. 191–200. doi: [10.1109/REAL.1990.128747](https://doi.org/10.1109/REAL.1990.128747).

References

(6/9)

- [12] R. Balas, A. Ottaviano, and L. Benini, "CV32RT: Enabling Fast Interrupt and Context Switching for RISC-V Microcontrollers," **IEEE Transactions on Very Large Scale Integration (VLSI) Systems**, vol. 32, no. 6, pp. 1032–1044, 2024, doi: [10.1109/TVLSI.2024.3377130](https://doi.org/10.1109/TVLSI.2024.3377130).
- [13] A. Nurmi, P. Lindgren, A. Kalache, H. Lunnikivi, and T. D. Hämäläinen, "Atalanta: Open-Source RISC-V Microcontroller for Rust-Based Hard Real-Time Systems," in **Architecture of Computing Systems**, D. Fey, B. Stabernack, S. Lankes, M.

References

(7/9)

- Pacher, and T. Pionteck, Eds., Cham: Springer Nature Switzerland, 2024, pp. 316–330.
- [14] A. Nurmi, P. Lindgren, T. Szymkowiak, and T. D. Hämäläinen, “AnTiQ: A Hardware-Accelerated Priority Queue Design with Constant Time Arbitrary Element Removal,” in **2023 26th Euromicro Conference on Digital System Design (DSD)**, 2023, pp. 462–469. doi: [10.1109/DSD60849.2023.00070](https://doi.org/10.1109/DSD60849.2023.00070).
- [15] A. Nurmi, H. Lunnikivi, P. Lindgren, and T. D. Hämäläinen, “Towards Predictable Ultra-Low Latency End-Nodes with Hardware-Accelerated Abstract Timers,” in **2025 IEEE Nordic**

References

(8/9)

- Circuits and Systems Conference (NorCAS)**, 2025, pp. 1–6. doi: [10.1109/NorCAS66540.2025.11231291](https://doi.org/10.1109/NorCAS66540.2025.11231291).
- [16] A. Nurmi, A. Kalache, and T. D. Hämmäläinen, “Hardware Solutions for Eliminating Context Switching Latency in Processor-Based Hard Real-Time Systems,” in **2024 IEEE Nordic Circuits and Systems Conference (NorCAS)**, 2024, pp. 1–6. doi: [10.1109/NorCAS64408.2024.10752471](https://doi.org/10.1109/NorCAS64408.2024.10752471).
- [17] A. Nurmi, A. Kalache, H. Lunnikivi, P. Lindgren, and T. D. Hämmäläinen, “Efficient and Predictable Context Switching for Mixed-Criticality and Real-Time Systems,” **IEEE Transactions**

References

(9/9)

on Very Large Scale Integration (VLSI) Systems, vol. 33, no. 11, pp. 2907–2915, 2025, doi: [10.1109/TVLSI.2025.3612433](https://doi.org/10.1109/TVLSI.2025.3612433).

Thank you!
Questions?