

Carfield: An Open-Research Platform for Safety, Resilient and Predictable Systems

Integrated Systems Laboratory (ETH Zürich)

University of Bologna

Angelo Garofalo angelo.garofalo@unibo.it

PULP Platform Open Source Hardware, the way it should be!



Outline



- Why an Open-Research Platform for Mixed-Criticality Applications (Automotive)?
- Goals of the Project
- Carfield: Heterogeneous Architecture Overview
 - Linux-Capable Host
 - Safety- and Timing-Critical Tasks
 - Secure Boot and Data Encryption/Decryption
 - Pluggable Accelerators
- The First Prototype in Intel16 FinFet Technology
- The Future Research Roadmap
- Conclusion





[Otani, Renesas, from Tutorial at ISSCC23]

ETH zürich

- ECUs are moving towards domain, zone architectures (more than simple MCUs)
 - Aggregation of multiple functions into single compute modules
- Computing systems require higher real-time processing performance, enanched safety & security features → Interesting challenges to be addressed (...not only limited to the Automotive world ^(©))



Automotive Trends: Platforms on the Market

- Microcontroller class of devices
 - Infineon AURIX Family MCUs
 - Control tasks, low-power sensor acquisition & data processing Features: lockstepped 32-b HP TriCore CPU , HW I/O monitor, dedicated accelerators
- Powerful real-time architectures
 - ST Stellar G Series (based on ARM Cortex-R cores)
 - Domain controllers and zone-oriented ECUs
 - Features: HW-based virtualization, Multi-core Cortex-R52 (+ NEON) cluster in split-lock, vast I/Os connectivity

Application class processors

- NXP i.MX 8 Family
- ADAS, Infotainment
- Features: Cortex-A53, Cortex-A72, HW Virtualization, GPUs





Limited possibility to explore and innovate at computing architectural level





RISC-V Solutions For Automotive





High-End RISC-V applications based on core IPs from SiFive: Infotainment, ADAS, cockpit...

Imagination

GPU linked by RISC-V CPU targeting ASIL-B

mobileye"

advanced driver assist systems chips, 12 RISC-V cores, 176 trillions OP/s

ETH zürich



ASIL-D processors



- The advantages of RISC-V open ISA •
 - Royalty-free
 - Extensible w/ Domain-Specific instructions •
 - Lightweight solution to boost efficiency of key kernels (DSP, DNNs, reduced precision INT/FP..)
 - Supported by a fast-growing ecosystem
 - Compilers, SW stack..
 - Its adoption can no longer be ignored (14.5% of the SoC market will be powered by RISC-V by 2025) *[Semico, 2020]

What about Open-Source Hardware?



5

ALMA MATER STUDIORUM Università di Bologna SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

Advantages of Open Source (Computing) Hardware



Hardware whose design is made publicly available so that anyone can study, modify, distribute, make, and sell the design or hardware based on that design (*source: <u>Open Source Hardware (OSHW) Statement of Principles 1.0</u>)*





Goals of The Project



- Develop a pre-competitive **RISC-V** powered Automotive SoC
 - Based on fully open-source HW and SW IPs
 - Scalable and configurable architectural template based on our PULP architectural ball-park
 - To adapt to different computing requirements
 - Full SW stack to address requirements of RISC-V based automotive applications
- Collaborative research roadmap for automotive-driven computing architectures
 - Functional safety, Hardware/Software Acceleration, Time-Predictability, Fast-Interrupts, HW-based Virtualization
- Benchmark RISC-V ecosystem and architectural solutions for automotive
 - Contributing to European interest to build an automotive reference platform around RISC-V
- Close the gap between RISC-V and ARM-powered solutions



We Started Joining Forces



- Project's Leaders
- Digital Systems Design, PULP, Open-Source, RISC-V
- Processors/Ips/Interconnects/Interrupts/HW Acceleration
- SW stack, compilers, runtime and optimized routines
- ➤ Real-Time (RT) Systems and On/Off-Chip RT Communication
- Safe/Secure Cyber-Physical Systems
- Virtualization-assisted systems, OS, Hypervisors, RISC-V
- Intel16 FinFet technology (for the first prototype)
- o ASIC design support and packaging
- Supporters: STMicroelectronics, BOSCH

ETHZÜRICH () OKMAGENESTER ST KELOSUM SSH-Soc Workshop, DAC, San Francisco, 9 July 2023







Accelerators Domain

Carfield: Architectural Template Based on Fully Open IPs

Main Computing and I/O System

L2 MULTI-BANK **FP VECTOR CLUSTER (SPATZ)** SECURITY ISLAND SAFETY ISLAND SPM SPI JTAG L1 MULTI-BANKED SPM BK BK DATA SPM INSN SPM BK BK SRAM wECC BOOT ROM LOCAL INTERCO BK BK IBEX **RV-PLIC** CV32E4 CV32E4 CV32E4 BK BK CTRL FPU • • • FPU FPU • • • FPU SHA2 **KEYs** TRI-LOCKSTEP TRNG CC ECC PE0 CC0 VRF PE1 CC0 VRF **AES128** K-H-MAC OTPs Safe Hart DMA WATCHDOG MAILBOXes CLIC INTC MULTI-PORT (AXI) I\$ PREDICTABLE AXI INTERCONNECT I/Os AND MMU MMU **iDMA** L1 MULTI-BANKED SPM LAST LEVEL PERIPHERALS CACHE (LLC) D\$/I\$ wECC D\$/I\$ wECC TLB LOCAL INTERCONNECT UART QSPI CVA6 CVA6 Serial wH-EXT wH-EXT CAN RV Link DMA RV0 RV1 **HYPERBUS** ... Tensor (N-1) FPU FPU **GPIOs** ETH Core MEMORY CONTROLLER CLIC INTC CLIC INTC HMR WATCH I2C DOG I\$ HOST SUBSYSTEM HMR ACCELERATION CLUSTER TIMERS SPI

ALMA MATER STUDIORUM SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

ETH zürich



9



Host-Domain for Low-Criticality Linux-Based Applications



P D

Host-Domain for Low-Criticality Linux-Based Applications



HyperRAM

- Coherent Cluster of 2 CVA6 cores based on a light-weight self-invalidation protocol
- CVA6: Application mid-end processor
 - Linux-capable processor, 1.7GHz, 1.65 DMIPS/MHz
 - RV64GC ISA, six-stage, in-order pipeline
 - Supports 48-bit virtual memory MMU (Sv48)
 - M, HS (Hypervisor-extended Supervisor) and U privilege modes
 - Tightly integrated D\$ and I\$
- Shared Last Level [Data] Cache (LLC)
- HyperBUS Ctrl for off-chip HyperRAM access



11

[Zaruba et al., IEEE Trans. VLSI, 2019]

ETH zürich () AMMARATER ST VOICENM SSH-SoC Workshop, DAC, San Francisco, 9 July 2023



How Do We Handle Safety-Critical and Real-Time Tasks?

SAFETY ISLAND		SECURITY ISLAND	L2 MULTI-BANK SPM	FP VECTOR CLUSTER (SPATZ)						
DATA SPM	INSN SPM	SPI JTAG	ВК ВК	L1 MULTI-BANKED SPM						
CV32E4 CV TRI-L CLIC INTC	/32E4 CV32E4 OCKSTEP Safe Hart	SRAM wECCBOOT ROMIBEXRV-PLICSHA2KEYSTRNGAES128OTPsK-H-MACWATCHDOGMAILBOXes	BKBKBKBKBKBKECCMULTI-PORT (AXI)	CTRL CC DMA	LOCAL INTERCO • • • FPU FPU CCO VRF PE I\$	J • • • FPU 1 CC0 VRF				
PREDICTABLE AXI INTERCONNECT										
I/Os AND PERIPHERALS		MMU MMU	iDMA	L1 MULTI-BANKED SPM						
UART QS	PI • Protectio	on against transient faults	(safety)							
Serial CAI	Serial CAN • Predictable On-Chip Communication (RT)									
GPIOs ETI	• Reduced	Reduced contentions to access critical shared memory resources (RT)								
WATCH DOG		HOST SUBSYS	STEM		I\$					
SPI TIME	RS			ŀ	IMR ACCELERAT	ION CLUSTER				





12

The Safety Island





- Safety-critical applications running on top of a RTOS
- Three CV32E40 cores physically isolated operating in lockstep (single HART) and fast HW/SW recovery from faults
- ECC protected scratchpad memories for instructions and data
- Fast and Flexible Interrupts Handling through RISC-V compliant CLIC controller
- AXI-4 port for in/out communication

13



Predictable On-Chip Communication (AXI RT)



- AXI4 inherently unpredictable
- Minimal Intrusive Solution
 - No huge buffering, limited additional logic
 - Solution verified in systematic worst-case realtime analysis
- AXI Burst Splitter
 - Equalizes length of transactions to avoid unfair BW distribution in round-robin scheme
- AXI Cut & Forward
 - Configurable **chunking unit** to avoid long transaction delays influencing access time to the XBAR

AXI Bandwidth Reservation Unit

- Predictably enforces a given max nr of transactions per time period (to each master)
- Per-address-range credit-based mechanism



14

• Periodically **refreshed** (or by user)

Contention-Free Shared L2 Scratchpad Memory

1. Dual-AXI-Port L2 Mem Subsystem

Multi-banked L2 SPM accessible from two different AXI ports



4. We determine in SW which port and which mode to use

• By using different address space!



2. Two Address Mapping Modes





Non-interleaved

Interleaved

3. Dynamic Address Mapping by Address spaces, eg:



Point to the same L2 physical Mem space





What About Security and Data Encryption/Decryption?





16

ALMA MATER STUDIORUM SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

ETH zürich

The Security Island

Derived from the OpenTitan project by lowRISC



- Root of Trust
 - Stores Cryptographic Secrets



- Early Secure Boot Stages
 - Verify Cryptographic Signatures and Measurements before unlocking next boot stages
- Cryptographic Services
 - Available to the whole System through Mailboxes
 - Wide Set of Crypto-Accelerators
 - SHA2, AES128, etc.
- Secure Monitor for the platform
 ip opentitan





The I/O Communication

ETH zürich



ALMA MATER STUDIORUM SSH-SoC Workshop, DAC, San Francisco, 9 July 2023



18



The Spatz Acceleration Cluster

ETH zürich





19

ALMA MATER STUDIORUM SSH-SoC Workshop, DAC, San Francisco, 9 July 2023



- Multi-precision FPU support
 - FP64, FP32, FP16, FP8, SDOTP operations supported
- Physically-driven implementation: small footprint, high operating frequency, high scalability



20

[Cavalcante et al. , IEEE/ACM ICCAD 2022]

Memory bandwidth [B/cycle] 8

ETHzürich (C) GMMARMATER ST KOLOSYM SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

16



The HMR Acceleration Cluster

ETH zürich





21

ALMA MATER ST BOLORUM SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

The HMR Cluster for DNN-Oriented INT/FP Workloads





[Rogenmoser et al., arXiv, 2023][Tortorella et al., arXiv, 2023]**ETH** zürichCompared bit VoltorialSSH-Soc Workshop, DAC, San Francisco, 9 July 2023

- 12x 32-bit RISC-V cores with support for DSP/QNN ISA Extensions
- Single-Cycle Multi-Banked Tightly-Coupled Data Memory (Scratchpad)
- Hardware Synchronizer
- DMA Controller for Explicit Memory Management
- L1-coupled TensorCore (RedMule)
- Runtime-configurable Dual/Triple core redundancy mode + hw/swbased quick recovery mechanism



Hybrid Modular Redundancy (HMR): Reconfigurable



23

Independent Mode: high performance, no reliability



Hybrid Modular Redundancy (HMR): Reconfigurable



24

DMR Mode: good performance, good reliability, slow recovery



Hybrid Modular Redundancy (HMR): Reconfigurable



25

TMR Mode: low performance, high reliability, quick recovery



Rapid Recovery: shared hardware extension



- Cycle-by-cycle backup of the cores state in ECCprotected Status Registers
- Quick recovery procedure (24 cycles!)
- Shared logic between TMR and DMR modes





HMR, yes... but at which cost?





DI II D Cluster	Area [mm ²]	Overhead
		Overneau
Baseline	0.604	
DMR	0.605	0.3%
TMR	0.608	0.7%
HMR	0.612	1.3%
With I	Rapid Recove	ery
DMR	0.654	8.4%
TMR	0.657	8.8%
HMR	0.660	9.4%

	DMR	TMR	DMR Rapid Recovery	TMR Rapid Recovery
Recovery Latency [cycles]	Application dependant	363	24	24
Mode Switching [cycles]	703	598	603	515

[Rogenmoser et al., arXiv, 2023]





Carfield In Intel FinFet Technology

- Project started in March 2023 •
 - After one month of specifications and discussions with partners and supporters
- Tape-Out Date: 11 November 2023 •
 - 4x4mm² Chip in Intel16 FinFet technology
 - Advanced BGA Flip-Chip Packaging
- First prototype available by end of Q1 2024 ٠
- Today we have the full platform in place •
 - FPGA Emulation available

ETH zürich

- Started functional verification effort
- Started SW ecosystem development



Software Stack (Under Development)





ALMA MATER STUDIORUM Università di Bologna SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

ETH zürich



29

The (Near) Future Research Roadmap

Some of the topics we are addressing:

- SW Stack Development
- Design Optimizations
- Concurrent OS Support for RTOS and GPOS
- Virtualization Assisted Processors
- Real-Time analysis of I/O communication
- Area-Optimized Safety solutions for RISC-V processors
- Reconfigurable Architectures for Image Processing
- Radiation tests









Overcoming the Safety-Island Concept





31

Mixed-Criticality CVA6 Host Subsystem



- Concurrent execution of GPOS and one or more RTOSs (through Hypervisor) Safety
- Safety and Reliability solutions ٠
 - ECC protected memories
 - Redundancy modes for critical hardware
- HW Cache partitioning (LLC) ٠
 - Add predictability properties to cache
 - Predictable loads/stores for RTOSs
 - To prevent cache lines to be used by different quests (OS)
- Fast and predictable Interrupt • Handling (Virtualization-assisted)







RI

Fast Virtual Interrupts (part of RISC-V AIA Spec)

Without Virtualization (current)





- High-Latency to claim the Interrupts
 - Need intervention of the Hypervisor (Trap & Emulate)



- Fast, predictable, low-latency interrupts
 - Vritualization of CLIC controller
 - Interrupts Virtualization
- Direct IRQ routing/forwarding (to guests)
 - No Trap & Emulate Routines
- Reduce time-penalty to handle interrupts
 - Critical for real-time systems (E.g. timer interrupts



33

ETH zürich () OLMVERSITER ST VOICENA SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

ETH zürich SSH-SoC Workshop, DAC, San Francisco, 9 July 2023

github.com/pulp-platform/carfield

Conclusion

- Carfield: Open-Source Research Platform for Safety, Predictable and • Secure Systems
 - Hardware Architecture based on fully open-source lps
 - Complete Software stack (open-source)
- Collaborative Research opportunities among universities and industry
- First prototype soon to silicon-prove initial architecture and get feedback for next generation platform
- Looking forwards for feedbacks/contributions from industry and academia







Thank you!



ETHzürich 🛞 Manager angeland



Parallel Ultra Low Power

Luca Benini, Alessandro Capotondi, Alessandro Ottaviano, Alessio Burrello, Alfio Di Mauro, Andrea Borghesi, Andrea Cossettini, Andreas Kurth, Angelo Garofalo, Antonio Pullini, Arpan Prasad, Bjoern Forsberg, Corrado Bonfanti, Cristian Cioflan, Daniele Palossi, Davide Rossi, Fabio Montagna, Florian Glaser, Florian Zaruba, Francesco Conti, Georg Rutishauser, Germain Haugou, Gianna Paulin, Giuseppe Tagliavini, Hanna Müller, Luca Bertaccini, Luca Valente, Manuel Eggimann, Manuele Rusci, Marco Guermandi, Matheus Cavalcante, Matteo Perotti, Matteo Spallanzani, Michael Rogenmoser, Moritz Scherer, Moritz Schneider, Nazareno Bruschi, Nils Wistoff, Pasquale Davide Schiavone, Paul Scheffler, Philipp Mayer, Robert Balas, Samuel Riedel, Segio Mazzola, Sergei Vostrikov, Simone Benatti, Stefan Mach, Thomas Benz, Thorir Ingolfsson, Tim Fischer, Victor Javier Kartsch Morinigo, Vlad Niculescu, Xiaying Wang, Yichao Zhang, Frank K. Gürkaynak, all our past collaborators and many more that we forgot to mention





http://pulp-platform.org

@pulp_platform

References



- [Restuccia et al. DAC 2020]: Restuccia, F., Biondi, A., Marinoni, M., Cicero, G., & Buttazzo, G. (2020, July). Axi hyperconnect: A predictable, hypervisor-level interconnect for hardware accelerators in fpga soc. In 2020 57th ACM/IEEE Design Automation Conference (DAC) (pp. 1-6). IEEE.
- [Pagani et al. ECRTS 2019] Pagani, M., Rossi, E., Biondi, A., Marinoni, M., Lipari, G., & Buttazzo, G. (2019, July). A bandwidth reservation mechanism for AXI-based hardware accelerators on FPGAs. In 31st Euromicro Conference on Real-Time Systems (ECRTS 2019).
- [Modica et al., IEEE ICIT 2018] Modica, P., Biondi, A., Buttazzo, G., & Patel, A. (2018, February). Supporting temporal and spatial isolation in a hypervisor for arm multicore platforms. In 2018 IEEE International Conference on Industrial Technology (ICIT) (pp. 1651-1657). IEEE.
- [Sá et al., IEEE TC 2021] Sá, B., Martins, J., & Pinto, S. (2021). A first look at RISC-V virtualization from an embedded systems perspective. IEEE Transactions on Computers, 71(9), 2177-2190.
- [Semico, 2020] https://semico.com/content/risc-v-cores-approach-115-cagr-2020-2025-says-semico-research
- [Cavalcante et al., IEEE/ACM ICCAD 2022]: Cavalcante, M., Wüthrich, D., Perotti, M., Riedel, S., & Benini, L. (2022, October). Spatz: A Compact Vector Processing Unit for High-Performance and Energy-Efficient Shared-L1 Clusters. In Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design (pp. 1-9).
 The zürich SSH-Soc Workshop, DAC, San Francisco, 9 July 2023



References - 2



- [Rogenmoser et al., arXiv, 2023]: Rogenmoser, M., Tortorella, Y., Rossi, D., Conti, F., & Benini, L. (2023). Hybrid Modular Redundancy: Exploring Modular Redundancy Approaches in RISC-V Multi-Core Computing Clusters for Reliable Processing in Space. arXiv preprint arXiv:2303.08706.
- [Tortorella et al., arXiv, 2023]: Tortorella, Y., Bertaccini, L., Benini, L., Rossi, D., & Conti, F. (2023). RedMule: A Mixed-Precision Matrix-Matrix Operation Engine for Flexible and Energy-Efficient On-Chip Linear Algebra and TinyML Training Acceleration. arXiv preprint arXiv:2301.03904.
- [Zaruba et al., IEEE Trans. VLSI, 2019]: Zaruba, F., & Benini, L. (2019). The cost of application-class processing: Energy and performance analysis of a Linux-ready 1.7-GHz 64-bit RISC-V core in 22-nm FDSOI technology. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 27(11), 2629-2640



