DOI: 10.1109/ICMMTS62835.2025.10926015

Towards adaptive RISC-V based systems for non-terrestrial sub-THz communication

Hendrik Borchert*, Markus Ulbricht*, Marko Andjelkovic*, Diana Göhringer ‡, Milos Krstic*†
*IHP - Leibniz Institute for High Performance Microelectronics, Frankfurt (Oder), Germany
‡Technical University of Dresden, Dresden, Germany
†University of Potsdam, Potsdam, Germany
borchert|ulbricht|andjelkovic|krs@ihp-microelectronics.com, diana.goehringer@tu-dresden.de

***Abstract*—The upcoming 6G communication standard promises unrivaled bandwidth, connectivity, and coverage and will likely span from most remote places over densely populated areas into low earth orbit. The implementation of this vision, however, poses many considerable challenges to the underlying processing hardware with advanced solutions needed to meet these requirements – especially in space. These challenges include the need for significant technological advances, critical demands in terms of performance, reliability, and adaptability, and considerations in terms of the trustworthiness of devices, to name only a few of them. This paper presents our joint efforts to address these needs and enable open-source, adaptive, and fault-tolerant processing systems for 6G communication systems in low-earth orbit.**

***Index Terms*—RISC-V, wireless communications, non-terrestrial networks, 6G communications, baseband processor, MAC processor, multi-core processor, fault tolerance, resilience, adaptive system**

I. Introduction

The ongoing roll-out of fifth-generation (5G) mobile networks promises highly available, high-speed, ultra-low latency communication resources, enabling applications like smart healthcare, virtual reality, and edge computing. However, the continuous demand for even greater capabilities drives further research into the next iteration of mobile networks.

The upcoming sixth generation (6G) promises even higher networking capabilities, with features such as AI integration, increased sustainability, and advanced sensing, enhancing extended reality and next-generation robotics ([1]). Achieving these advancements requires academia and industry to explore new spaces - both technologically and physically.

Technological Space - Implementing communication devices that meet 6G's demanding requirements presents significant technological challenges, including data rates ranging from hundreds of Gbps to several Tbps, localization accuracy below 1 cm, latency below 100 ns, error rate less than 1e-9, and connectivity for millions of devices ([1], [2]). This will require groundbreaking advancements in radio transceiver hardware, efficient information processing, channeling, managing, and an optimized implementation of the protocol stack, with a focus on heterogeneous computing and digital acceleration ([3]).

Physical Space - To provide limitless connectivity everywhere—another ambitious vision of 6G—research and industry have begun focusing on solutions beyond Earth. Low Earth Orbit (LEO) satellite communication is expected to be a key component of 6G wireless communication ([4]). This was once unthinkable but is now possible due to significantly reduced launch costs, streamlined satellite designs enabling low-cost mass production, and a strong drive to connect underserved areas ([5]). Several high-tech private companies have already started deploying satellites into LEO, aiming to create mega-constellations with thousands or even tens of thousands of devices.

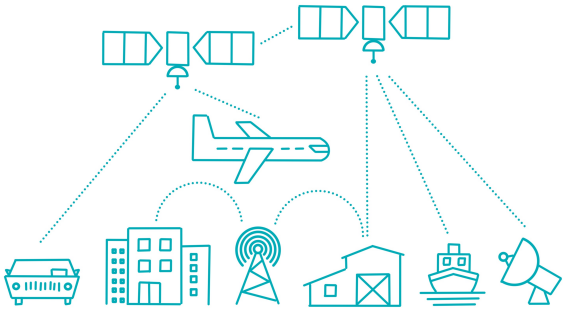

Fig. 1. Envisioned satellite-based ubiquitous 6G coverage (based on [5])

However, operating and maintaining advanced technology in space poses considerable challenges. Immense acceleration forces and strong vibrations during launch, ionizing radiation, high differences in temperature, solar particle events, and voltage variations in place put a lot of strain on the hardware. These conditions can cause temporary disruptions, as seen with Starlink ([6]) as a consequence of a solar particle event, or lead to complete device failures. Deploying systems in such harsh but also very dynamic environments requires high resilience, adaptability to ever-changing conditions, as well as high configurability and maintainability. In addition to being controllable from Earth, these systems must be capable of autonomously adapting to dynamic internal and external changes.

Design Space - Another crucial area to explore is the design space. Due to continuous efforts by space agencies or communication companies, a wide range of commercial processing platforms are already available that meet the high-performance and reliability demands of space-based communication systems. However, given the critical role of the next generation of mobile networks, these systems must also meet the highest standards of

trustworthiness and accessibility. One cornerstone of achieving this is formed by open-source soft- and hardware, offering transparency, low-threshold access, and high adaptability. Open source instruction set architectures like RISC-V have sparked significant interest in academia and industry alike, and chip solutions for space applications based on the open RSIC-V ISA have emerged, such as Trikarenos ([7]) and NOEL-V ([8]). Even ESA is considering using RISC-V as the basis for their processing systems ([9]). The primary challenge, however, remains: developing a trustworthy and adaptive base for safety-critical systems ranging from small IoT sensors to high-performance computing. Addressing this challenge will require combined efforts from both professional and academic partners.

The following paper summarizes our efforts to tackle the different challenges in all of these spaces as part of our shared vision to enable open-source, adaptive, and fault-tolerant processing systems for 6G communication in low earth orbit. As a first example, chapter II provides an overview of the TETRISC SoC - a fault-tolerant quad-core system based on the RISC-V ISA. Chapter III elaborates on efficient and scalable medium access control design for space applications, while chapter IV focuses on reconfigurable and fault-tolerant baseband processor design. Chapter V introduces a domain-specific computing architecture called Agiler that offers high parallelism and adaptivity. Chapter VI finally summarizes the presented approaches and presents an outlook on further steps on our path to the envisioned space-borne 6G systems.

II. The TETRISC SoC

Integrated circuits today face increasing challenges due to effects coming from radiation, aging, and temperature variations, which are particularly critical in sectors such as automotive and aerospace. Addressing these challenges requires a system that can adapt to varying conditions while maintaining high reliability. In order to address this, we have proposed an adaptive fault tolerant system, building on the foundation of our first-generation TETRISC SoC ([10]). The new system, based on RISC-V processor cores, combines cross-layer hardening techniques with real-time system status detection, enabling it to adjust to optimal operating modes based on current conditions.

The TETRISC SoC is a quad-core multiprocessor system that leverages on-chip monitors and the inherent redundancy of a multiprocessor system to achieve adaptive fault tolerance (see Fig. 2). The first-generation TETRISC featured SEU, aging, and temperature monitors, along with a reconfigurable voter network, which collectively ensured core reliability. The system could operate in various modes, such as single-core, DMR, TMR, and QMR, providing possibility for trade-off between power consumption, performance, and reliability. In the second-generation system, we have made several optimizations to enhance these features, ensuring better performance and adaptability.

Enhanced On-Chip Monitors - The new SoC incorporates significant upgrades to its reliability monitors. This includes the addition of dynamic scrubbing for the SEU monitor, precise aging measurements via shadow registers, and a shift to smaller, more efficient, ring oscillator-based temperature monitors.

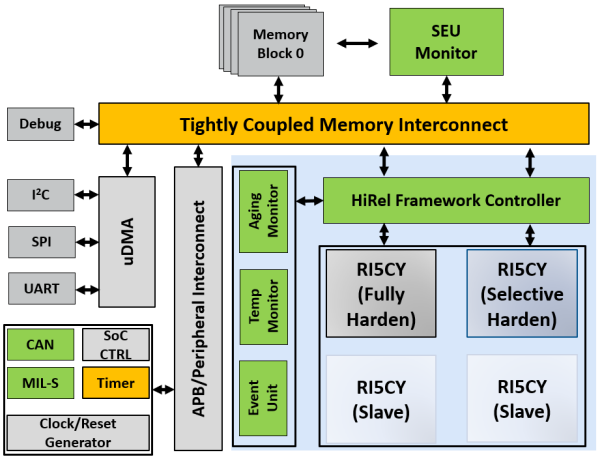

Fig. 2. The system architecture of the next-gen TETRISC SoC.

Adaptive Fault Tolerance Mechanism - The system leverages real-time data from the on-chip monitors to adapt its operational mode. It can operate in three modes: Performance Mode (independent core operation), Destress Mode (thermal management through clock gating), and Fault Tolerant Mode (N-modular redundancy for enhanced resilience).

Optimized System Architecture - The TETRISC SoC architecture has been refined to improve system performance while maintaining adaptability. These optimizations include enhancements to the HiRel Framework Controller (HFC) and the critical path design, ensuring the system can dynamically respond to reliability challenges without significant overhead.

III. Medium Access Control Design

Developing 6G capable communication systems for non-terrestrial deployment in space poses multiple hurdles. Meeting key performance indicator (KPI) goals regarding data throughput, connectivity, coverage, and latency while maintaining high energy efficiency already presents a substantial challenge but this is further aggravated by the effects of the low earth orbit on the reliable operation of electronics. One potential solution to these problems involves deploying heterogeneous computing platforms ([3]) that combine classical resources of computation, like CPUs, with highly efficient – but less flexible – domain-specific accelerators for the computationally expensive tasks of the wireless communication stack.

The medium access control (MAC), as part of layer 2 of the OSI model, plays a critical role in managing access to the shared wireless spectrum available to the communication system with potentially many devices contending for transmission time ([11]). Opposite to terrestrial networks, satellites cannot readily leverage cloud computation resources. Due to the high distance to Earth, it is imperative to tightly integrate the MAC into the edge to keep latency requirements and energy constraints in check. The same holds for baseband processing. To address this need for flexible yet powerful edge devices,

we adopt a control plane/data plane separation as depicted in Fig. 3, where the control plane deploys RISC-V CPUs, providing the necessary software-based reconfigurability. In contrast, the data plane primarily comprises memory-mapped hardware accelerators without requiring the same level of flexibility.

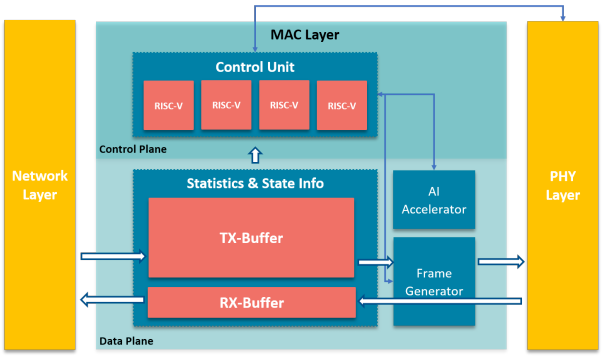

Fig. 3. High-level architectural view on the medium access control with the CRISPY platform serving as the control plane.

A. Data Plane

The data plane encompasses all components related to data flow with focus on packet/frame handling and processing. This includes, among other MAC functions, header generation, frame de/aggregation, channel prioritization and automatic repeat request (ARQ) ([11]). Building a MAC data plane consists of two major challenges: a) achieving efficient data processing through data parallelism while keeping critical path length as low as possible, and b) designing a scalable system that supports many nodes without compromising on throughput or latency. The latter problem becomes particularly pronounced in systems where frame aggregation is necessary to reduce frequent device switching and the associated latency, such as in line-of-sight communication in sub-THz bands.

a) Efficient Data Processing: A high degree of Bit-level parallelism and utilizing low-latency and high-throughput memory technologies, such as SRAM and High Bandwidth Memory (HBM), are advantageous for achieving ultra high data rates. Both FPGA and ASIC designs can effectively exploit bit-level parallelism. Incoming narrow data words are converted into words with high bit-width by a data aggregator and subsequently written into the TX buffer before further processing occurs. Notably, clock frequency tends to exhibit comparatively low degradation as bit-width increases, which is important given the limited possibilities for frequency scaling, especially on FPGAs. On tested platforms, we observed that the bottleneck in frequency scaling does not stem from the components of the MAC design itself but rather from the inherent constraints and limitations of the FPGA, which is an expected outcome.

b) Scalable System Design: The scalability of Medium Access Control (MAC) depends mainly on the algorithmic complexity of accessing the TX buffer, both in terms of space and time. Systems deploying dedicated buffer architectures lack good space complexity as the amount of memory needed scales with the number of network nodes supported by the hardware. Therefore, we adopt a shared buffer architecture, thus ensuring better utilization of the scarce memory resources. We opt for a list-like memory structure to allow for constant-time data retrieval and insertion, with data stored on a per-destination basis.

Our buffer architecture efficiently supports two key modalities: frame aggregation and Automatic Repeat reQuest (ARQ). It is essential that frame aggregation occurs *on-the-fly*, where frames are streamed, processed, and immediately forwarded from the buffer to the PHY layer without any intermediate buffering of the aggregated frame.

B. Control Plane

Control functions are typically centered around resource management and network coordination and involve tasks such as scheduling, random access procedures for UE registration, physical layer configuration, UE buffer status reporting, discontinuous reception (DRX), or uplink power control and are well-suited for CPU processing as they demand relatively modest computation power, yet require the high flexibility of general purpose computing. In particular, they profit significantly from efficient branching, a key strength of modern CPU architectures ([12]). Nevertheless, domain-specific accelerators could benefit certain control plane functionalities, like AI-based scheduling. To explore this potential, we are developing a RISC-V-based computing platform utilizing the RISC-V based PULP Cheshire ([13]) platform and the NVIDIA Deep Learning Accelerator (NVDLA) ([14]).

C. Reliable Medium Access Control

Achieving platform reliability includes the usage of shielding techniques, system self-monitoring of aging, redundant computation, and data processing. Consequently, we plan to transfer the TETRISC architecture to the CRISPY platform to ensure proper control functionality in space. Memory structures dominate within the data plane, for which adaptive error correction codes (ECC), combined with age monitoring, might be a suitable fit as adaptive ECCs ([15], [16]) allow for dynamic trading between reliability and efficiency.

IV. Reconfigurable Baseband Processing

A key component in a wireless communication system is the baseband processor, which performs final processing of data before transmission and preliminary processing of data received by the RF frontend. The baseband processor acts as an interface between the RF frontend and the MAC stage. In any application, the baseband processor should be designed to be tolerant to errors caused by noise in the communication channel. In addition, for space applications, the baseband processing stage must be also tolerant to errors resulting from ionizing radiation, particularly to soft errors.

In order to address the needs for high data rates and enhanced reliability in 6G communication systems, in the framework of the 6G-TakeOff project we aim to develop

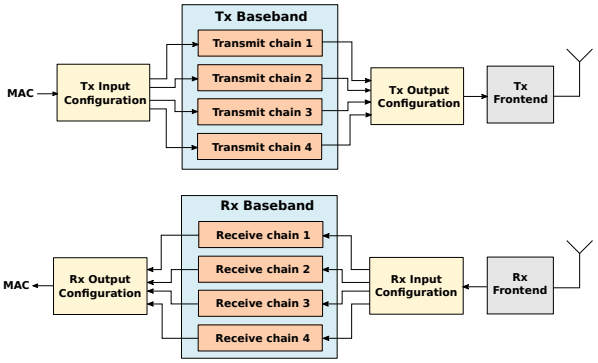

Fig. 4. General architecture of the parallelized baseband processing system.

a reconfigurable baseband processing unit for satellite-to-ground communications, supporting a single carrier signal and 8 modulation schemes: BPSK, QPSK, APSK16, APSK32, APSK64, QAM16, QAM32, and QAM64. The system will be able to select autonomously the optimal modulation scheme based on application requirements and channel state.

As the single baseband processing stage cannot support the expected data rates for 6G communication networks, the idea is to employ parallel transmission and reception stages. The system-level design of the baseband processor with four transmission and reception stages is shown in Fig. 4. One of the main design goals is to achieve the fault-tolerance of the baseband processor under harsh radiation conditions in space. To this end, we are considering a combination of local fault detection in each block and reconfigurable system-level fault-tolerance based on N-modular redundancy. The local fault detection will be based on duplication with comparison. When a fault in any block is detected, the respective control unit will initiate a re-execution of data processing only within that block. This approach reduces the overall latency, since the processing is repeated only in the faulty block and not in all blocks. At system level, the fault-tolerance of the whole baseband processor could be enhanced by employing a similar approach as in the case of TETRISC SoC, i. e., by configuring multiple transmit or receive stages into DMR or TMR settings. The system-level DMR or TMR modes could be suitable in the case of increased radiation intensity in space, which could be detected with special radiation sensors. However, it is important to mention that the data rate would be reduced during DMR and TRM modes because the number of parallel processing stages would be one (in the case of TMR) or two (in the case of DMR), instead of four (in non-fault-tolerant mode). The control logic for reconfiguration can be realized either as a custom-designed hardware accelerator or in the form of a general-purpose processor (e.g., RISC-V processor).

V. Self-Adaptive Domain-Specific Computing Architectures

To provide high-performance, energy-efficiency and reliability for 6G communication in low earth orbit, novel adaptive domain-specific hardware/software solutions are required. Data must be processed in parallel with a several of signal processing algorithms. In addition, the computing system must be able to detect faults and adapt quickly to changing situations at runtime, e.g., by exchanging a faulty compute unit using reconfiguration. This feature can be provided by self-adaptive domain-specific computer architectures using reconfigurable hardware, where each component (software, processor, accelerator, communication infrastructure, memory) can be adapted at design- and runtime according to application demands.

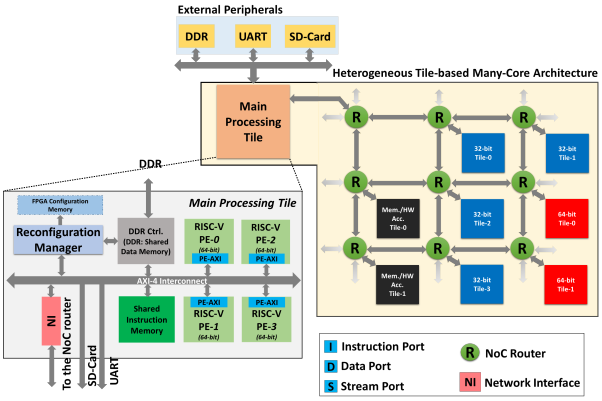

Fig. 5. Overview of the Agiler platform [17]

One example of such an computing architecture is Agiler [18] (see also Fig. 5), which offers a high degree of flexibility and scalability. Agiler is an adaptive tile-based many-core architecture. Different many-core configurations are possible based on the application requirements. A scalable 2D mesh topology is used to connect heterogeneous compute tiles. Each compute tile can be configured as a single-/multi-core architecture consisting of different RISC-V based processors, memory and dedicated hardware accelerators. The main processing tile is responsible for distributing the applications to the compute tiles as well as for reconfiguring the overall system using dynamic and partial reconfiguration via its reconfiguration manager [19]. Dynamic and partial reconfiguration offers a fine-grained reconfigurability on bit-level. This gives a high degree of freedom but results in a higher reconfiguration overhead in terms of longer reconfiguration times. If a faster adaptation is desired coarse-grained reconfigurable architectures (CGRAs) are a promising solution, as they offer reconfiguration on word- instead of bit-level. One example here is DA-CGRA [20], which consists of processing elements connected with each other over switch boxes. The processing elements are selected at design time based on the application requirements and are heterogeneous to achieve a high energy-efficiency. DA-CGRA requires a processor to control the reconfiguration and would therefore be a promising candidate to function as an accelerator compute tile within the Agiler platform.

VI. Summary and Outlook

This paper introduces our collective efforts to advance on open-source, adaptive, and fault-tolerant processing systems for 6G communication networks in low earth orbit. These

efforts share a significant overlap in the challenges they address, which we defined in the introduction of this paper. Technical challenges, such as providing sufficient processing power to meet the enormous demands of 6G, are tackled through adaptive and heterogeneous computing. The challenges posed by the physical operation of the devices in low Earth orbit, which increase the risk of premature system aging and early failures, are mitigated by continuous system health monitoring and fault tolerance techniques. Lastly, challenges related to accessibility and trustworthiness are addressed through the use of open-source solutions, such as processing systems based on the RISC-V ISA.

To advance our shared vision of a 6G and space-ready computing and communication platform, further improvements in the reliability and system development domain have to be made. We are envisioning a system both fully shielded against radiation by means of classic shielding techniques as well as redundancy not only on CPU level as shown in our TETRISC SoC but also in the data plane for components of both medium access control and baseband. Moreover, scalability and performance are two main concerns that need to be addressed in the context of reliability due to the added redundancy and higher complexity of the system and warrants investigation in future work.

VII. Acknowledgement

The work presented in this report has received funding from the German Federal Ministry of Education and Research (BMBF - Bundesministerium für Bildung und Forschung) under grant No. 16ME0134 (Scale4Edge) and grant No. 16KISK060 (6G-TakeOff), the European Union's Horizon Europe research and innovation program under grant No. 101092598 (COREnext) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 287022738 – TRR196 for the associated project S05. The authors further acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the programme of "Souverän. Digital. Vernetzt.". Joint project 6G-life, project identification number: 16KISK001K. The responsibility for the content of this publication lies with the authors.

During the preparation of this work, the authors used the large language model GPT-4 (Release October 2023) in order to improve readability and language. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

[1] M. A. Uusitalo, P. Rugeland, M. R. Boldi, E. C. Strinati, P. Demestichas, M. Ericson, G. P. Fettweis, M. C. Filippou, A. Gati, M.-H. Hamon, M. Hoffmann, M. Latva-Aho, A. Pärssinen, B. Richerzhagen, H. Schotten, T. Svensson, G. Wikström, H. Wymeersch, V. Ziegler, and Y. Zou, "6G Vision, Value, Use Cases and Technologies From European 6G Flagship Project Hexa-X," IEEE Access, vol. 9, pp. 160 004–160 020, 2021, conference Name: IEEE Access. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9625032

[2] H. Tataria, M. Shafi, M. Dohler, and S. Sun, "Six Critical Challenges for 6G Wireless Systems: A Summary and Some Solutions," IEEE Vehicular Technology Magazine, vol. PP, pp. 2–12, Jan. 2022.

[3] V. Ziegler, H. Viswanathan, H. Flinck, M. Hoffmann, V. Räisänen, and K. Hätönen, "6G Architecture to Connect the Worlds," IEEE Access, vol. 8, pp. 173 508–173 520, 2020, conference Name: IEEE Access. [Online]. Available: https://ieeexplore.ieee.org/document/9200631

[4] M. Giordani and M. Zorzi, "Satellite Communication at Millimeter Waves: a Key Enabler of the 6G Era." IEEE Computer Society, Feb. 2020, pp. 383–388, iSSN: 2325-2626. [Online]. Available: https://www.computer.org/csdl/proceedings-article/icnc/2020/09049651/1iERZ5WwiZO

[5] X. Lin, S. Cioni, G. Charbit, N. Chuberre, S. Hellsten, and J.-F. Boutillon, "On the Path to 6G: Embracing the Next Wave of Low Earth Orbit Satellite Access," Aug. 2021, arXiv:2104.10533 [cs]. [Online]. Available: http://arxiv.org/abs/2104.10533

[6] U. Shetti, "Musk's Starlink satellites disrupted by major solar storm," Reuters, May 2024. [Online]. Available: https://www.reuters.com/technology/space/starlink-experiencing-degraded-service-musk-says-satellites-under-pressure-2024-05-11/

[7] M. Rogenmoser and L. Benini, "Trikarenos: A Fault-Tolerant RISC-V-based Microcontroller for CubeSats in 28nm," Oct. 2023, arXiv:2310.02045 [cs]. [Online]. Available: http://arxiv.org/abs/2310.02045

[8] F. Gaisler, "NOEL-V (RISC-V)." [Online]. Available: https://www.gaisler.com/index.php/products/processors/noel-v

[9] R. Weigand, "RISC-V: a Rising Star in Space," Barcelona, Jul. 2023. [Online]. Available: http://microelectronics.esa.int/papers/PresentationSummit-EUR-RISC-V-RisingStarInSpace-2023-06-08.pdf

[10] M. Ulbricht, L. Lu, J. Chen, and M. Krstic, "The TETRISC SoC-A resilient quad-core system based on the ResiliCell approach," Microelectronics Reliability, vol. 148, p. 115173, Sep. 2023.

[11] 3rd Generation Partnership Project (3GPP), 5G; NR; Medium Access Control (MAC) protocol specification, ETSI Std. 38.321, 2023, release 18. [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/38_series/38.321/38321-g20.zip

[12] L. Vespa, A. Bauman, and J. Wells, "Algorithm Flattening: Complete branch elimination for GPU requires a paradigm shift from CPU thinking," in 2015 IEEE High Performance Extreme Computing Conference (HPEC), Sep. 2015, pp. 1–6. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7322477

[13] A. Ottaviano, T. Benz, P. Scheffler, and L. Benini, "Cheshire: A lightweight, linux-capable risc-v host platform for domain-specific accelerator plug-in," IEEE Transactions on Circuits and Systems II: Express Briefs, vol. 70, no. 10, pp. 3777–3781, 2023.

[14] N. Corporation, "Nvdla open source hardware, version 1.0," https://github.com/nvdla/hw, 2024, accessed: 2024-09-02.

[15] Y. S. Lee, G. Koo, Y.-H. Gong, and S. W. Chung, "Stealth ECC: A Data-Width Aware Adaptive ECC Scheme for DRAM Error Resilience," in 2022 Design, Automation & Test in Europe Conference & Exhibition (DATE), Mar. 2022. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9774775

[16] D. Kim and J. Kim, "Adaptive Granularity On-die ECC," in 2022 19th International SoC Design Conference (ISOCC), Oct. 2022. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10031324

[17] A. Kamaleldin and D. Göhringer, "An agile tile-based platform for adaptive heterogeneous many-core systems," in 2022 International Conference on Field-Programmable Technology (ICFPT), 2022, pp. 1–4.

[18] A. Kamaleldin and D. Göhringer, "Agiler: An adaptive heterogeneous tile-based many-core architecture for risc-v processors," IEEE Access, vol. 10, pp. 43 895–43 913, 2022.

[19] N. Charaf, A. Kamaleldin, M. Thümmler, and D. Göhringer, "Rv-cap: Enabling dynamic partial reconfiguration for fpga-based risc-v system-on-chip," in 2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), 2021, pp. 172–179.

[20] E. Aliagha, N. Charaf, N. Krishna Venkatesan, and D. Göhringer, "DA-CGRA: Domain-aware heterogeneous coarse-grained reconfigurable architecture for the edge," in Euromicro Conference on Digital System Design (DSD), 2024.