EUROPEAN WORKSHOP ON ON-BOARD DATA PROCESING (OBDP2021), 14-17 JUNE 2021

DEPENDABLE MPSOC FRAMEWORK FOR MIXED CRITICALITY APPLICATIONS

Renato Costa Amorim[(1)], Rodolfo Martins [(1)], Prem Harikrishnan [(1)], Max Ghiglione [(2)] , Tim Helfers [(2)]

[(1)]Evoleo Technologies GmbH, Freibadstr. 30 D-81543 München – Germany, Email:infode@evoleotech.com
[(2)]Airbus Defence and Space GmbH, Robert-Koch-Straße, 85521 Taufkirchen - Germany, Email: max.ghiglione@airbus.com

ABSTRACT

System-on-Chip such as the Zynq UltraScale+ combine multi-core processors (PS), programmable logic (PL) and peripherals in a single device. For space applications, these devices offer the possibility for unprecedented functional integration and performance in a smaller form factor. Radiation effects introduce challenges to ensure reliability and availability of embedded applications.

EVOLEO and AIRBUS, in the frame of the ESA GSTP project CHICS, are developing an ADHA compatible radiation tolerant 3U onboard computer, based on the Zynq Ultrascale+ for mixed criticality space applications. The solution considers a clear separation between platform and payload functions within the MPSoC. It is oriented towards parallel but independent developments for platform and payload functions, which are often the responsibility of different entities.

This paper describes the challenges in such a system, proposed solutions along with the latest performance results. Plans for further developments and expected results are presented.

1. THE CHICS ONBOARD COMPUTER

The CHICS OBC (COTS-based Highly Integrated Computer System for Mini/Nano Satellites), is a joint development by EVOLEO Technologies GmbH and Airbus Defence and Space Ottobrunn. The goal is to provide an engineering model of a SAVOIR OBC based on COTS parts that can fulfil a large spectrum of the market/use cases needs in terms of functionality, performance, radiation tolerance, availability and commercial soundness.

The vision that drives the design is that current COTS components have reached such a level of maturity and radiation tolerance that it is possible to create highly performant, integrated Low Earth Orbit (LEO) avionics for a fraction of the cost of traditional rad-hard solutions without compromising on system availability.

The New Space Market opens the design possibilities to COTS components as lower orbits (LEO between 600 and 1200 km) have a less harsh radiation environment. Moreover, the faster pace of the market poses less strict requirements on availability and lifetime, enabling designs based on error detection and safe recovery, instead of complete avoidance.

It is therefore a balancing act to meet these goals by careful selection of COTS components with good radiation tolerance, robust implementation of critical functions and design oriented towards fault detection, isolation and recovery as opposed to fault avoidance.

If these orientations are followed, such avionics would provide unparallel solutions for small LEO satellites (30kg to 200kg), those that target high volume/ low cost (Earth observations constellations) or highly available in orbit edge computing (one-shot missions).

The CHICS OBC targets up to 4 years lifetime in LEO orbits and an availability better than 98,75%.

1.1. The rational for payload integration in OBCs

Current COTS multi-core System on Chip offer an unprecedented possibility to integrate intermediate sensor data processing and data fusing for AOCS, sensor acquisition and pre-processing for payloads and actuators data processing.

Current developments prove the possibility of employing SW-defined-GNSS, assisted by star tracker data fusion, visual based navigation including target pointing and tracking. Unprecedented possibilities in payload data fusion gathered from diverse sensors for physical measurements are made possible, but such applications pose high demands on processing power, memory capacity and bandwidth.

Current state of art space-borne processing platforms, like LEON-based SOC are not sufficient for such high-performance applications, even when employed in combination with current most performing FPGAs like the Microsemi RTG4. Moreover, the integration of such multiplatform systems is complex and prone to errors in design phase, increasing further the cost of the design.

The integration of AOCS and payload processing capabilities on a COTS platform will benefit all users due to the reduction in cost, mass and power consumption of the system. The main interest in payload integration lies in imaging services and visual navigation which could be supported by utilizing the GPU and Application processors of the Xilinx MPSoC families.

The applicability of the proposed OBC is being demonstrated via two distinct use-case situations: As an OBC with integrated AOCS capabilities (Star-tracker and GNSS) and as a generic processing board for edge Artificial Intelligence/Machine learning in Space towards Failure Prognostics and Detection in onboard equipment.

1.2. Advanced Data Handling Architecture

The goal of the ADHA activity supported by ESA, Advanced Data Handling Architecture, is to challenge the elements of the current data handling architectures, which may limit its growth for the coming years. It considers two viewpoints: technically, by aiming at reducing "mass, power and size" budgets, harness and AIT effort in the data handling system (DHS); industrially, by setting up new specifications and a roadmap that open up new possibilities in the supply chain of equipment modules to multiple Large Scale Integrators (LSI).

The main objective is to establish a versatile, compact, modular and scalable DHS architecture for application like mini and medium size Earth Observation (EO) satellites in LEO and constellation. Standardization of the DHS includes mechanical and electrical interfaces, with focus on interoperable and interchangeable modules based on space standards derived from commercial.

These objectives are highly correlated with the objectives of the CHICS OBC. Following ADHA specifications is an added value for both activities, as CHICS can become a versatile high processing unit for most types of ADHA solutions at a fraction of the cost of rad-hard solutions.

As an example, Sentinel type missions can resort to CHICS as a payload processing unit, capable of edge AI/ML for feature detection or other non-mission critical applications to increase the quality of the derived Earth observation products. Smaller missions can exploit the full potential of CHICS as a complete OBC with integrated AOCS.

1.3. Avionics Architecture

The baseline elements of the OBC are: 1) Zynq UltraScale +, support by DDR4 SDRAM 2)PolarFire FPGA supported by TMR'ed NAND Flash.
The Zynq MPSoC provides the high performance and integration of mixed criticality applications, whilst the PolarFire FPGA implements critical functions such as system reconfiguration and telemetry and telecommand encoding and decoding.

This functional allocation aims at potentiating the Zynq MPSoC capabilities as much as possible whilst relying on the PolarFire FPGA inherent radiation tolerance for safety critical functions. The European Space AVionics Open Interface aRchitecture (SAVOIR) specification is followed with little to no tailoring needed to implement all OBC functions predicted in the standard.

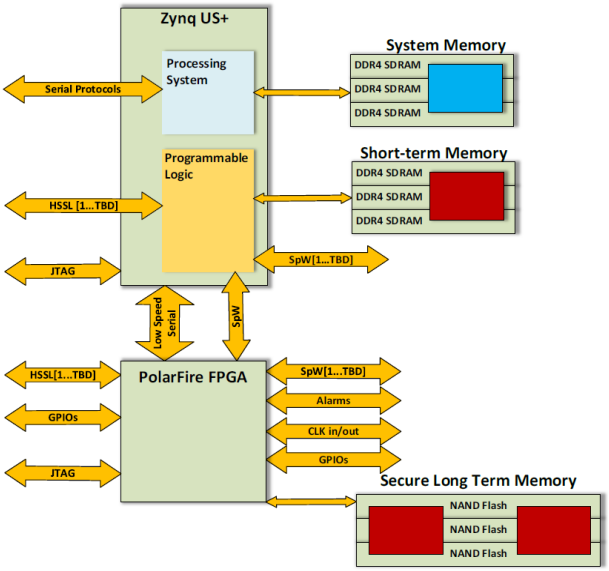

Figure 1. CHICS OBC high-level architecture

Fault detection, isolation and recovery are implemented through via a combination of hardware and software strategies and centralized in the PolarFire FPGA as the focal point for system reconfiguration.

All memories are protected via error correction codes, either through the embedded Zynq memory controllers or via bespoke VHDL controllers tailored to the known fault modes of each memory model. Most notably, the PolarFire FPGA interfaces with GB-sized commercial NAND Flash modules with TMR voting inside the FPGA. Focusing on radiation tolerance through IP cores/Software, as opposed to rad-tolerant components, accelerates the adoption of ever improving commercial EEE parts.

An array of analog, low speed and high-speed digital series interfaces are available to the front panel and

backplane. Both the Zynq MPSoC and PolarFire FPGA implement SpaceWire routers for a fault-tolerant SpaceWire network, always ensuring communications and control of the OBC.

2. MIXED CRITICALITY APPLICATIONS IN MPSOC

The criticality of applications such as AOCS, autonomous operations, payload data compression, payload data feature extraction and FDIR is distinct between themselves and across missions. Clearly, functions that ensure the control and operations of the satellite have a higher criticality and impact in mission safety compared to acquisition of payload data or extraction of knowledge from that data.

MPSoC technologies can combine these applications which have different criticality levels in the same chip. The challenge and innovation may lay in how to best implement these designs ensuring that higher criticality functions are not compromised by the lower criticality ones, whilst understanding that the hardware platform limits the maximum achievable reliability, in particular in radiation environments.

2.1. The challenges

The challenges associated with mixed criticality applications in MPSoC can be defined at a higher level without exhaustive understanding of the MPSoC architecture. In the scope of the CHICS OBC development, design challenges were subdivided into the following categories:

- Radiation effects: "Can the device survive the radiation environment?" and "How can one detect and mitigate non-destructive upsets?"
- Resource isolation: "How does one create isolated pools of resources inside the MPSoC?"
- Functional availability: "How can one optimize the individual availability for each application or function running on the MPSoC?"
- Fault propagation through data sharing: "Can two applications with different criticality levels exchange data without spreading faults?"
- Ease of adoption and tailoring: "How can one reuse this baseline HW/SW technology for multiple use-cases and combination of embedded applications?"
- Minimal baseline dependability: "Can one guarantee a high baseline level of availability, performance and FDIR features with little non-recurrent engineering effort?"

2.2. The approach

To solve these challenges, device agnostic or device-specific solutions can be employed. Considering the high internal complexity of the MPSoC and the features it provides, the team has followed a combined approach. The proposed approach is based on three design pillars: 1) Secure/non-secure side isolation; 2) Secure data exchange buffer and monitor; 3) External Supervisor for critical faults.
Each design pillar is described in more detail in the upcoming sections.

2.2.1. Secure/non-secure side isolation

The first presupposition is that resources and software shall be isolated as far as possible between the OBC, called secure side, and the payload applications called the non-secure side or "user" side in order to avoid fault propagation and deterioration of overall availability.
The secure side includes mission critical applications such as main OBSW, FDIR routines, AOCS control and TM/TC execution. These are implemented on the lockstep ARM-R5 real time cores (RPU).
The non-secure side is reserved to non-mission critical/ embedded payload applications mapped onto the ARM-A53 application cores (APU) as XEN hypervisor guests.
For the programmable logic, embedded memories and peripheral resources, they should also be reserved and isolated as far as possible.

2.2.2. Secure data exchange buffer and monitor

In order to ensure that faulty data does not propagate between isolation areas, and that data flow is managed by a secured element, the concept for the data exchange buffer and monitor was created.
This buffer is a shared data storage element that secure and non-secure sides can use to read and write data according to their own priorities and operational schedule. In parallel to this operation, a configurable monitoring block performs constant checks on the datasets to ensure these are fault free. These checks can range in complexity from thresholds checks on numerical data to more complex machine learning algorithms for fault detection in spacecraft equipment associated with the APU application.

2.2.3. External supervisor for critical faults

The design approach considers extensive fault detection, isolation and recovery features locally at the MPSoC secure side for lower criticality faults. Nevertheless, the know radiation sensitivity of the MPSoC require an external entity to recover from more critical faults which may impact the local FDIR functions.

Therefore, a PolarFire FPGA supervisor shall handle all faults that the secure-side cannot recover by itself, as a last resort. Besides these tasks, the Supervisor also monitors and control other PCB elements such as the power chain which increase the overall MPSoC reliability and availability.

Taking these assumptions, the native features provided by the Zynq MPSoC hardware, Xilinx-supported tools and open source software were researched and combined with other design ideas to achieve a complete solution.

2.3. The solution

Following the design pillars presented below a variety of design elements were design, implemented and configured towards the final solution. The following sections presents a description of these elements.

2.3.1. Secure/non-secure side isolation

The Xilinx MPSoC platform provides multiple hardware features and methods to isolate internal resources. These can mostly be summarized by the following:

- System Memory Management Unit (SMMU)
- Xilinx Memory Protection Unit (XMPU)
- Xilinx Peripheral Protection Unit (XPPU)
- TrustZone Isolation (TZ) [1]
- Hypervisors such as XEN for virtualization and Isolation.

These features and methods can be effectively used to isolate subsystems within the MPSoC, reduce interference, thereby providing a deterministic real time processing platform. The figure below shows an hardware isolation of the processing system running OBC on the secure side on top of FreeRTOS and AOCS applications running on the non-secure side as Hypervisor guests.

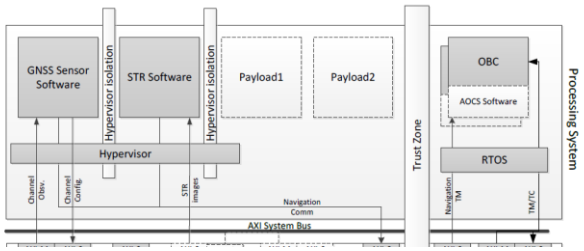

Figure 2. CHICS OBC software blocks for OBC-GNSS-Star Tracker use-case

2.3.1.1. Hardware Isolation

ARM TrustZone provides isolation across the MPSoC platform including the different processors, cores, memories, hardware accelerators, interrupt controllers, AXI transactions, other peripherals and execution states by storing them in dedicated physical registers with a single bit indicating whether the resource can be accessible by secure or non-secure side. TZ is enabled using the Xilinx Vivado toolsuite.
The MPSoC platform provides unique master IDs for each of the cores and peripherals in the Processing system, and AXI ports connecting to the Programmable Logic. Using these unique IDs, the MPSoC can be further isolated by filtering the master node AXI transactions using the XMPU, XPPU, AXI Timeout blocks (ATB) and AXI Isolation Blocks (AIB) by different cores, memories, AXI ports and peripherals. In the current OBC-AOCS use-case, a typical example of using hardware isolation is that the PL master of the GNSS IP block uses the AxPROT signal (TZ) on the FPD AXI ports connected to Processing System to write DMA bursts on the DDR4 memory (XMPU) via the SMMU Translational Buffer Units (TBU).

2.3.1.2. Virtualization of APU

The applications cores (APU) are implementing the Star Tracker and GNSS software elements of the AOCS applications processing chain. They are virtualized using the XEN hypervisor as Baremetal/FreeRTOS guests running on separate A53 cores. The Hypervisor uses the SMMU to provide address translation. The A53 cores share the L2 cache to access the DDR4 memory in a non-virtualized environment, which can cause interference and affect real time processing. The L2 cache will be further statically partitioned to use dedicated cache lines to access the DDR4 from each A53 cores using cache coloring, a feature provided by the XEN hypervisor [2]. This demonstrates Symmetric Multi Processing (SMP) with isolation and deterministic performance on the A53 cores.

2.3.2. Secure data exchange buffer and monitor

Two possibilities were identified to support a data exchange buffer: 1) the embedded inter-processor interrupts (IPI) and associated message buffers 2) a shared data region on the programmable logic. In the end, both functionalities were combined to create a bespoke secure data exchange buffer which is not only reliable but also expandable.

The Zynq MPSoC provides multiple inter-processor Interrupts (IPI) connecting the different cores on the Processing System, Programmable Logic and the Platform Management Unit. This interrupts also include dedicated physical message buffers protected by XPPU for each channel allowing up to 32 KB of data to be send with each interrupt. These interrupts, despite extremely useful for transmitting small amounts of data, are limited in size and configuration by the hardware itself. It is them not possible to adjust the buffer

configuration and size according to the mission use-case and associated set of critical and non-critical applications. Also, it would not be possible to snoop the data and thus have a third element monitoring errors in that data.

The solution is to instantiate dual port RAM blocks in the programmable logic, two for each APU A53 core (one area for the APU to RPU direction and another RAM block for the reverse direction). These are called the exchange buffers. Each APU core can only read or write to their assigned memory region whilst the RPU can read and write to all of them.

Updates on the written data are coordinated by use of the IPI. When a new dataset is put into the buffer, the core that has written to that buffer generates an IPI to the destination core. Since there is only one IPI reserved for the APU, this will be multiplexed with a time slice using the Hypervisor and with information on the message buffers for each target processor.

Besides this through-traffic data, all data being write to the buffer is duplicated to a monitoring buffer and made available to a fault detection block for snooping. This block can access the monitoring buffer via AXI interface and therefore can be customized to the use-case. These blocks can mimic PUS-type FDIR services such as parameter monitoring (limit check, expected value and delta checks) or perform more complex monitoring algorithms based on machine learning algorithms. This is particularly interest to detect faulty spacecraft equipment processed by the APU.

The Exchange Monitor, XMPU and XPPU will be set to generate an interrupt to the FDIR application running on the RPU in case of access violation, data corruption, synchronization loss or watchdog timeout. All BRAM blocks are protected from single event upsets via SECDEC codes.

The figure below shows an example of information flow of the exchange buffer along with the exchange monitor.

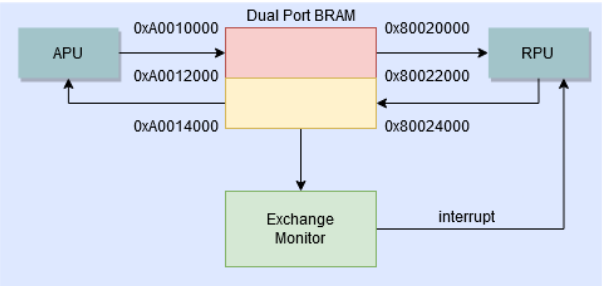

Figure 3: Exchange buffer and monitor basic architecture.

2.3.3. External Supervisor

The external supervision of the Zynq MPSoC is performed by a flash based, configuration upset immune PolarFire FPGA [3]. In the general context of the CHICS OBC, it provides centralized acquisition and evaluation of telemetries.

It communicates with the MPSoC via a SpaceWire link. Telemetry exchange is utilized to detect anomalies in the Zynq. A watchdog timeout in this communication signals a faulty Zynq. Additional GPIO lines and power chain monitoring elements provide an overarching understanding of the health status of the Zynq and other board elements.

The PolarFire can directly control latchup control limiter switches to reset individual power islands in the Zynq. Thus, a more granular fault recovery can be implemented, maintain part of the Zynq functionality operational while faulty resources are reset.

2.4. FDIR services

The Zynq implements a comprehensive set of FDIR services conforming to PUS standards centralized in the RPU cores. It provides onboard monitoring, housekeeping, event report generation and trigger recovery sequence. This FDIR application will monitor parameters, functions and generate housekeeping data in the MPSoC and trigger a recovery sequence based on the severity if a fault is detected.

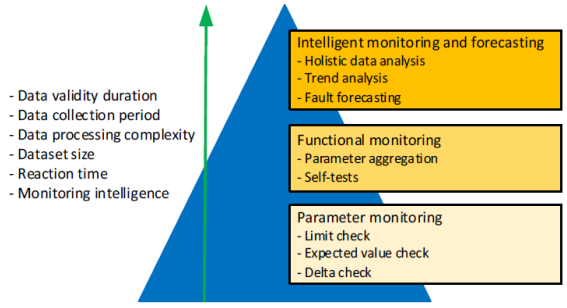

Figure 4: Conceptual FDIR intelligence pyramid

The source code is flexible enough to supporting increasing number of monitored parameters in the FDIR data pool. Therefore, as the design matures and use-cases are defined, more elements representing the system health can be added and monitored.

The FDIR services based on custom software are based on an escalating pyramid of monitoring complexity (figure 4). The initial design iterations will implement simple parameter and functional monitoring on dataset with low dimension and complexity. As the design matures and more knowledge on the system fault modes

is obtained, the processing capabilities of the ARM-R5 cores can be further exploited to increase the monitoring inteligence and isolation capabilities.

2.5. Increasing availability

The boot sequence of the A53 cores has a direct impact on the availability of the non-mission critical applications. The R5 cores run FreeRTOS which can typically boot in milliseconds while the A53 cores running XEN guests can take anywhere from 2 seconds up to 20 seconds depending on the XEN configuration.

XEN is a type 1 hypervisor meaning it can run directly on top of any hardware without any OS. This provides an important trade off on how the hypervisor guests can start. The traditional way is to start the guest after XEN boots the Domain-0 which is the most privileged domain running Linux kernel. In this configuration the LINUX kernel takes up to 20 seconds approximately to load and only then additional guests can be started. This configuration provides more control on how and when the guests can be started and reset. The figure 4 below shows the traditional boot sequence of XEN guests.

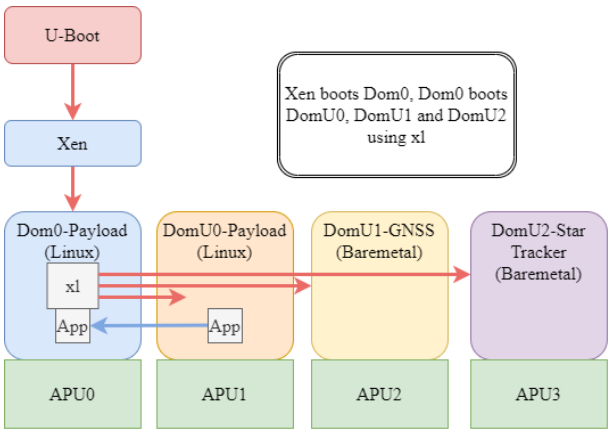

Figure 4: Dom0 boot sequence

A second possible configuration is called the Dom0less method [4], as seen in Figure 5. In this configuration the guests can be loaded from Uboot as soon as XEN loads. The total time for a guest to start is typically less than 2 seconds. It provides the benefits of faster boot time, complete isolation from Dom0, lower complexity with configuration and easier to certify system.

However, this configuration is not without its drawbacks. It also implies that there will be no monitoring of guest partitions by the Dom0. Thus, if a fault detected on one of the guests requires restarting, all other guests will also need to be restarted.

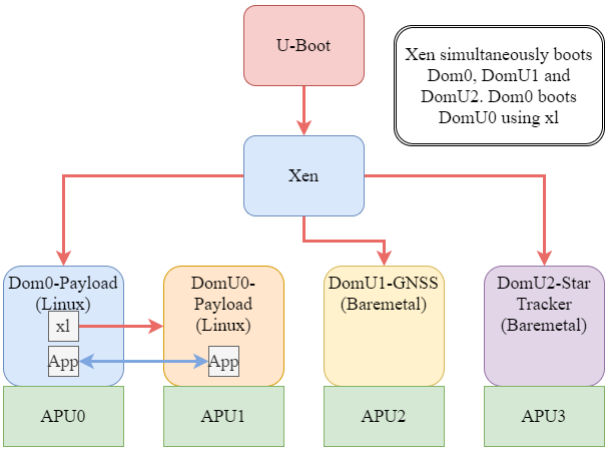

Figure 5. Dom0-less boot sequence

A Dom0 configuration is thus suggested for use-cases where there is a low number of XEN guests in an environment where guest restarting is uncommon.

2.6. Radiation effects mitigation

Several internal resources of the Zynq MPSoC have been characterized for radiation environments in terms of destructive and non-destructive effects [5]. Nevertheless, it is an extremely complex device which makes extensive characterization of all its fault modes extremely hard.

Destructive effects are arguably the most critical failure mode affecting a potential COTS device in space applications. Considering the design changes in the military grade XQ Zynq MPSoC, an XQ device in LVAUX mode has a relatively low probability of a destructive single event effect in low-Earth Orbit. Thus, latch up current limiters (LCL) can be effectively included in the design in order to protect the know sensitivity power inputs.

Non-destructive single event effects such as bit-flips, configuration upsets and functional interrupts (SEFI) are more common and widespread for SRAM-based devices. The CHICS OBC enables all embedded ECC on internal memories, including dedicated SECDEC IP blocks protecting programmable logic elements. Corruption of programmable logic configuration is handled by Xilinx Soft Error Mitigation IP Core. Counters for corrected errors and uncorrectable errors are all flagged to the central FDIR application.

The proposed Zynq architecture may be an effective approach to tolerate fault modes which have not been and may hardly be identified. By isolating as much as possible the chip resources, data flows and creating barriers in these dataflows, faults can be more easily detected and contained. Also, it provides a framework that can be used to observe these faults and propagation

paths in the effort to constantly harden the design as in-orbit experience is acquired.

3. EARLY TEST RESULTS

The CHICS OBC design has passed the preliminary design review and is moving towards a first engineering model and demonstration environment by the end of 2021. It is expected that TRL6 will be achieved in 2022 via validation of the engineering model in a ADHA rack. Furthermore, an exchange monitoring unit based on AI/ML algorithms will be implemented in 2022 for detection of faulty equipment running on the non-secure side as seen in figure 5.

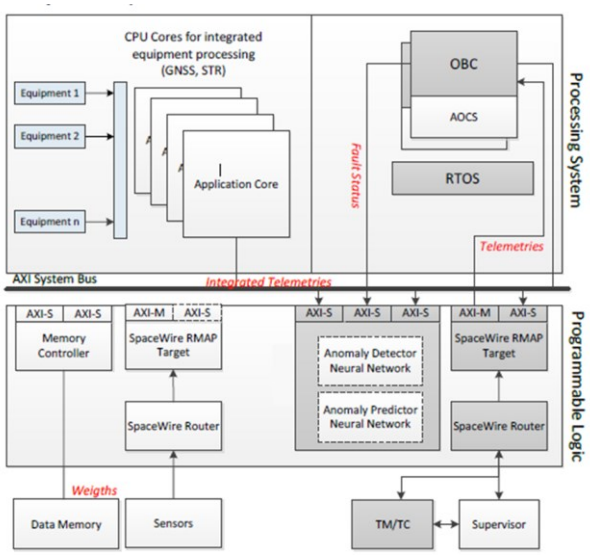

Figure 7: Exchange buffer with anomaly detector neural network

The current work on creating a demonstration environment has already been successful in demonstrating the proposed hardening approach. The team has programmed a ZCU102 development kit to demonstrate a prototype of the exchange buffer, isolation configuration and XEN hypervisors guests.

The XEN hypervisor on APU, at 1200MHz was configured with cache colouring enabled. It can run Baremetal, FreeRTOS and LINUX applications as guests also in Dom0less configuration. R5 cores (500MHz) were set to run on lockstep with FreeRTOS.
The exchange buffer was created with dual port BRAMs protected with SECDEC on the programmable logic to exchange information between APU /RPU using inter-processor interrupts to synchronize.

Custom VHDL elements, called demo blocks, were created to stimulate AXI DMA bursts with interrupts from PL to DDR targeted to be read by a XEN guest running on the A53 cores. Additionally, demo blocks sending AXI DMA bursts with interrupts from PL to OnChip Memory were created. These produced data to be read by R5 cores running on FreeRTOS.

3.1. Validated features

The Exchange Buffer was validated with FreeRTOS on R5 and FreeRTOS on A53 as XEN guest with inter-processor interrupts for synchronization as seen in figure 6.

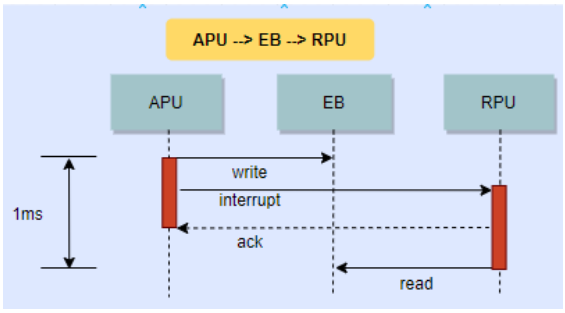

Figure 6: Exchange buffer APU write, RPU read sequence

These initial tests set read and write rates at 1KHz, which can be further increased. The size of the exchange buffer was set to 64 KB per core, which can also be increased further based on requirements.

The PL demo block generates periodic DMA bursts with interrupts for hypervisor guests on APU indicating the completion AXI DMA write in a specified memory location on the DDR. The DMA interval can be configured from anywhere between 10 microseconds to 1 second. The 10 microseconds intervals were validated on Xen as both baremetal and FreeRTOS guests with cache colouring enabled for a total of 50 samples.
The following figure shows the sequence of interactions between the APU, PL and the DDR.

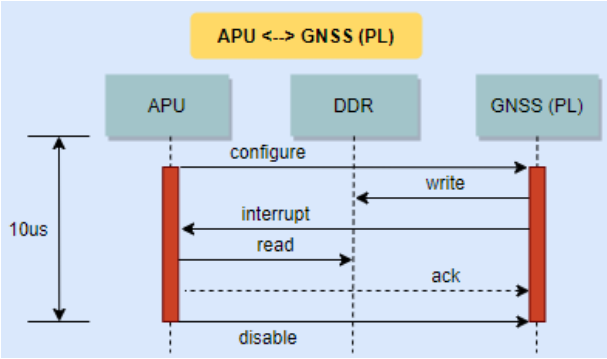

Figure 7: GNSS DMA sequence

Interrupt latency on the APU for baremetal guest was found to be 130 nanoseconds and FreeRTOS guest to be 230 nanoseconds without any interference on the L2 cache with null scheduler and static pinning of virtual cpus to each core, (XEN interrupt latency according to [2] can be up to 10 to 15 microseconds with interference

on the L2 cache which can be reduced up to 2 microseconds with cache colouring enabled).

As an example of the integration of navigation platform equipment, Star Tracker and GNSS sensor software have been designed on application cores isolated by a hypervisor. The advantage of the MPSoC is not only the multicore processing, but also the possibility of adding dedicated hardware in the programmable logic. Based on the AGGA4 GNSS receiver design, up to 24 channels can be integrated in the base configuration of the CHICS computer, enabling dual constellation dual frequency capabilities integrated in the OBC. Configurations with more channels, Star Tracker acceleration, or advanced anomaly detection algorithms employing neural networks can be envisioned exploiting the programmable logic.

4. Design reuse and further development

Through the design and implementation process, the team has identified several opportunities to reuse the CHICS OBC as a baseline hardware/software platform for technology development, as a basis for a line-up of ADHA avionics boards and as an initial design iteration of an evolving OBC composed the most recent system on chip.

It is currently being considered as a platform for the development of key enabling technologies and high performance applications in the field of COTS and MPSoC hardening for space applications.

4.1. Design reuse

The high-level architecture of the CHICS OBC presented in this paper has been proposed for additional processing boards of an ADHA avionics suite (onboard computer, payload interface and processing, mass memory, software defined radio).
The concept is to replicate for all boards the same basic HW/SW assets that provide power, processing functionality, interfaces and storage and augment them with the necessary items to tailor them to their particular functions. For example, further NAND Flash modules can be added to the underside of the PCB expanding mass memory capabilities. Another example is the replacement of front-panel interfaces according to mission needs with little changes to the main design elements.

By replicating across all avionics boards the same basic hardware and software components, they can all largely share the same fault mitigation techniques which only need to be developed once. The dependability issues arising for the use of COTS technologies, at all levels from architectural to power chain or software, can be solved for all boards simultaneously. Thus, multiple COTS-based boards can be designed in shorter time, lower cost, less risk and still retain dependability as a single board or as a full system.

In essence, there is a single reference design that is used as starting point to the design of other boards that implement similar or expanded functionality and interfaces.

4.2. Next generation ACAP OBC

Versal devices are claimed to be the first adaptive compute acceleration platforms (ACAP). These kinds of System on Chips combine adaptable processing and acceleration engines with programmable logic and configurable connectivity. This means that on a single system on chip FPGA, Processing and DSP capabilities are hosted.

Thanks to this integrated interconnect the Intelligent Engines for AI and DSP can be used directly by the processing system enabling a superior performance/watt that typical systems. This feature especially is very compelling in space applications due to the limited power budget and dissipation issues.

The Versal AI Core and Prime series are currently being considered for a follow-up design of the original CHICS OBC, guided by similar design principals. Such a design, aligned with cPCI Serial Space and ADHA would bring unprecedented processing capabilities to small satellites enabling key applications such payload processing based on AI/ML for non-mission critical applications. This includes optical payload use-cases such as object detection, segmentation, compression, autonomous calibration or autonomous definition of downlink priorities according to collected payload data.
This next generation can largely reuse the SW/VHDL assets developed in the scope of the CHICS OBC due to the similarities in the devices and reusage of Xilinx development tools.

5. Conclusions

In this paper, the CHICS OBC was presented with particular emphasis on its features enabling mixed criticality applications on the MPSoC. The OBC is in development under ESA contract and will achieve TRL6 during 2022.
Three design pillars were identified as essential to be able to run mission critical and non-critical applications within highly heterogenous MPSoC. This combination is thought to be key towards high functional integration for small satellites. The design exploits the resource isolation features of the MPSoC to create a secure and non-secure side. A dedicated data exchange buffer which is highly configurable and can be used to snoop on data exchanges between the two sides is used to

isolate faults. For faults that cannot be locally recovered by the FDIR services on the MPSoC, an external PolarFire FPGA performs recovery and reconfiguration functions.
Early test results suggest the feasibility of this design approach to meet a variety of use-cases such as embedded AOCS applications and AI/ML fault detection on equipment telemetries directly on the OBC.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the support and guidance provided by the ESA technical officers Gianluca Furano and Antonis Tavoularis under the scope of the CHICS OBC contract. Furthermore, the acknowledgement of the EVOLEO team members not mentioned as authors.

REFERENCESS

[1] Steven McNeil et al, October 26th, 2020, “Isolation Methods in Zynq Ultrascale+ MPSoCs” Xilinx XAPP1320 (v3.2)
[2] Stefano Stabellini, Embedded Linux Conference North America, 2020, “Xen Cache Colouring: Interference-Free Real-Time Systems”
[3] J. J. Wang et al, “Radiation characteristics of field programmable gate array using complementary-SONOS configuration cell”
[4] Stefano Stabellini, Embedded Linux Conference North America, August 2019, “Static Partitioning Made Simple”
[5] iRoC Technologies, April 2018, “Radiation Results of SER Test Xilinx XZU3EG”