

Microprocessors and FPGAs for Intensive, In-Orbit, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



Dr. Rajan Bedi is the CEO and founder of Spacechips, which designs and builds a range of advanced, L to K-band, ultra high-throughput on-board processors, transponders and OBCs for telecommunication, Earth-Observation, navigation, internet and M2M/IoT satellites. The company also offers Space-Electronics Design-Consultancy, Avionics Testing, Technical-Marketing, Business-Intelligence and Training Services. (www.spacechips.co.uk). Rajan can also be contacted on Twitter to discuss your space-electronics' needs: <https://twitter.com/DrRajanBedi>

Spacechips' Design-Consultancy Services develop bespoke satellite and spacecraft sub-systems, as well as advising customers how to use and select the right components, how to design, test, assemble and manufacture space electronics. We teach semiconductor memories on our FPGAs for Space Applications Training Course.

CONTEXT & INTRODUCTION

Satellite operators are increasingly acquiring **more and more data in-orbit**, and **would prefer to process this on-board the payload** to extract value-added insights **rather than downlink huge amounts of information** to a cloud for post-processing on the ground. Limitations with existing space-grade semiconductor technology and/or RF bandwidth constraints have hindered the amount of data that can be processed in real-time. I know of several customers who have had to descope their mission aspirations because of both reasons as their downlink needs would have violated ITU regulations.

In contrast, **localised processing** as close to the originating data source as possible, i.e. at **the Edge**, is based on **real-time computation of large amounts of information from multiple sensors**, acquired using low-latency, deterministic interfaces, in a small, low-power form factor with unique thermal and reliability requirements. **Extracting the analytics in-orbit significantly reduces delay and the RF downlink bandwidth** – we are effectively moving the data centre to the origin of the raw data!

In this post, I want to **discuss and compare microprocessors and FPGAs for intensive on-board processing at the Edge**. Some applications ingest huge amounts of data from multiple sensors with different bandwidths, e.g. RF, LIDAR, imaging and GNSS, and **require critical decisions to be made in real-time**, e.g. recognition and classification of objects for spacecraft situational awareness, i.e. identification of friend or foe, space-debris collision avoidance, high-definition video Earth observation and space-exploration in-situ, resource utilisation. There is also an **increasing trend for autonomous on-board processing using machine-learning** techniques to extract analytics in-orbit.

Microprocessors and FPGAs for Intensive, In-Orbit, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



EXISTING SOLUTIONS AND LIMITATIONS

Current on-board processing is based on microprocessors or FPGAs, neither of which are optimised for AI in-orbit characterisation of objects. The former are good for control, complex decision making and OS support, while the latter can process a diverse set of computationally demanding algorithms, excelling in data movement, custom acceleration, bit-oriented functions and interfacing. However, existing solutions cannot process linear algebra, matrices or vector processing efficiently, nor exploit parallelism at low power for autonomous machine-learning, AI inference and the implementation of neural networks for feature detection and classification.

In the commercial world, GPUs, originally developed for gamers, are being used to accelerate diverse computational tasks including encryption, financial modelling, networking and AI. GPUs use multiple cores and parallel processing to execute thousands of threads simultaneously, operating significantly faster and more cost effectively than microprocessors, allowing the calculation of data-intensive analytics from multiple sensors in milliseconds as opposed to seconds, minutes or hours. GPUs are optimised for performing the same operations over and over very quickly on large amounts of stored information, whereas CPUs tend to jump all over the place.

While there are almost thirty space-grade microcontrollers, microprocessors, FPGAs and dedicated DSP engines, only a small subset of these can be considered for in-orbit Edge-based applications. Many existing devices don't have the computational horsepower or low-latency memory/I/O interfaces. Some consume too much power requiring large and expensive thermal-management solutions: [previously](#) I described how to keep your space-qualified semiconductors cool to ensure their safe operation and maximise reliability. Table 1 lists the legacy standard processing products which I considered. For the FPGAs listed below, the specified performance is the theoretical peak based on the number of resources and clock frequency. The V5QV does not contain microprocessor IP as standard.

Supplier and Part Number	Performance (DMIPS)	Memory (MB)	Power Consumption (W)	Size (mm)
AAC Kryten-M3	62.5	8	0.4 to 1	96 x 90
BAE RAD750 3U	435	1,000	11 to 12.2	100 x 160
BAE RAD750 6U	260	48	11 to 14	233 x 160
BAE RAD5545	2,600	4,000	35	TBC
CAES DS4350272	44 to 95	256 SRAM + 8 NV	< 5	100 x 177
DDC SCS750	1,800	64,000	7 to 30	6U x 160
SEAKR Medusa	2,450	1,000	TBC	TBC
SEAKR Athena	2,450	1,000 + 4,000 NV	16	TBC
Spaceteq GR712	160	256 + 512 NV	< 5	233 x 166
ZIN SBCFA1000	30	2 + 0.5 NV	TBC	127 x 101.6

Table 1 : Off-the-shelf, space-grade, on-board processing solutions.

Microprocessors and FPGAs for Intensive, In-Orbit, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



With the amount of on-board data predicted to increase exponentially, what type of processor should you use for intensive, Edge-based, on-board computation? Is an MPU or an FPGA better? ESA's recent [workshop](#) on On-Board Data Processing highlighted current concerns, trends and future needs.

The fundamental technical limitations which have prevented in-orbit Edge processing are:

1. The lack of high-capacity, low-latency, low-power, space-grade memory. Current, fast space-grade storage is limited to volatile DDR3/DDR4 SDRAM. [Previously](#), I explained that to realise 1 Tb of on-board storage would require 64, 16 Gb chips, consuming a total of 17 W power, requiring 152.3 cm³ of physical space and a financial cost of £468,060. This is simply not a viable implementation at any level and space-qualified non-volatile memory is very slow.
2. The lack of power-efficient microprocessors or FPGAs for space applications offering the required processing capability. Over the last decade, 65 and 20 nm SRAM-based FPGAs have provided payload processing consuming 20 W, while 28 nm flash-based devices have offered a lower-power solution. Ultra-deep-submicron performance, logic densities and resources have resulted in an increase in consumption. Space-qualified MPUs with the required raw performance dissipate over 30 W.
3. The inability of existing space-grade, microprocessors or FPGAs to efficiently fuse and process inputs from multiple sensors. Moving large amounts of information to and from processors creates performance bottlenecks for data-intensive computation.
4. The inability of existing space-grade microprocessors or FPGAs to efficiently implement deep-learning algorithms for object identification and classification.

NEW SOLUTIONS FOR EDGE-BASED PROCESSING

To realise those applications which require in-orbit, Edge-based, on-board processing, the latest FPGAs and microprocessors are addressing the above limitations:

1. The availability of fast (up to 2,400 MT/s), 4 GB, space-grade DDR4 memory in a small form factor which I introduced in a [previous](#) The hardware design-in of this SDRAM has also been [discussed](#).
2. The availability of low-power, 28 nm flash FPGAs has reduced power consumption and more energy-efficient microprocessors have increased the GFlops / W metric.
3. Since 2020, Teledyne e2v's Radiation-Tolerant, QLS1046-4GB, compute-intensive, microprocessor includes a data-path acceleration architecture (DPAA) to increase packet parsing, queue management, hardware buffer management and cryptography, as well as supporting the IEEE 1588 precision time protocol. Also since 2020, Xilinx's [XQRKU060](#) has improved information flow and throughput, with the data path, I/O and memory interfaces optimised for low latency.
4. The next generation of 7 nm FPGAs contain dedicated AI tiles optimised for processing linear algebra to accelerate the performance of deep-learning algorithms. The QLS1046-4GB's four cores each contain native vector co-processors, e. NEON.

Microprocessors and FPGAs for Intensive, In-Orbit, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



Table 2 includes the latest, space-grade FPGAs and microprocessors: the former combine re-configurable logic, MPUs and the next generation of parts will contain AI tiles for efficient vector processing. For the FPGAs/MPSoCs listed in green, the specified performance is the theoretical peak based on the number of resources and clock frequency. Actual levels of computation will be lower depending on how these are used, implemented, memory and I/O usage, but Table 2 provides a useful comparison including the soft-core, RISC CPU. The highly parallel nature of the KU060 and Versal devices are reflected in their large TOPS values.

Supplier and Part Number	Performance (DMIPS) TOPS/DMIPS	Memory (MB)	Power Consumption (W)	Size (mm)
AAC Kryten-M3	62.5	8	0.4 to 1	96 x 90
BAE RAD750 3U	435	1,000	11 to 12.2	100 x 160
BAE RAD750 6U	260	48	11 to 14	233 x 160
BAE RAD5545	2,600	4,000	35	TBC
CAES DS4350272	44 to 95	256 SRAM + 8 NV	< 5	100 x 177
DDC SCS750	1,800	64,000	7 to 30	6U x 160
SEAKR Medusa	2,450	1,000	TBC	TBC
SEAKR Athena	2,450	1,000 + 4,000 NV	16	TBC
Spaceteq GR712	160	256 + 512 NV	< 5	233 x 166
ZIN SBCFA1000	30	2 + 0.5 NV	TBC	127 x 101.6
Teledyne e2v QLS1046-4GB	30,000	4,000 (DDR4)	6.5 to 19.4	44 x 26
Xilinx KU060	5.7 / 310	4	20	40 x 40
Microchip PolarFire	1.33 / 130	4.1	4 to 10	40 x 40
NanoXplore NG-ULTRA	1.34 / TBC	4	TBC	45 x 45
Xilinx Versal	157/19,516 & 2,505 38/19,516 & 2,505	6 to 24	40 to 120	45 x 45

Table 2 : Comparison of space-grade, on-board processing solutions.

With the amount of on-board data increasing significantly, there is an increasing trend for autonomous payload processing using AI and machine-learning techniques to extract analytics in-orbit for both timing critical and non-real-time applications. For example, a space-debris retrieval spacecraft outside of its ground-station coverage would not be able to receive a late command to initiate a collision-avoidance manoeuvre. On-board situational awareness acquired from multiple sensors followed by object detection and classification, would allow this timing-critical decision to be taken in real-time, independent from human intervention. Similarly, high-definition SAR imagery generates huge amounts of Earth-observation data and rather than clog precious RF downlinks, in-orbit AI inference and the implementation of neural networks would allow for feature identification, scene segmentation and characterisation.

Microprocessors and FPGAs for Intensive, *In-Orbit*, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



Teledyne e2v offers its Radiation-Tolerant, [Qormino](#) QLS1046-4GB, quad-core processor combining four ARM® Cortex® A72 cores operating up to 1.8 GHz with 4 GB of fast DDR4 SDRAM in a tiny, 44 x 26 mm form factor as shown below. Integrating off-chip memory with multiple CPUs onto a single substrate removes the need to design this complex, timing-critical interface, delivering significant size, weight and power (SWaP) advantages to enable in-orbit, Edge processing. The part delivers a computing performance of 30,000 DMIPS or greater than 45,000 CoreMarks.

The four MPUs execute the ARMv8-A architecture each with their own L1 32KB data and 48 KB instruction caches, as well as sharing a common 2 MB L2, as shown in Figure 2. With a core frequency of 1.2 GHz, a supply voltage of 1 V and a DDR rate of 1.6 GT/s, the total power consumption of the QLS1046-4GB ranges from 6.5 to 12 W (excluding peripherals) depending on the maximum permitted junction temperature. Likewise, at 1.8 GHz, a supply of 1 V and a DDR4 rate of 2.1 GT/s, the device dissipates from 9.3 to 19.4 W. Its raw computing performance together with memory bandwidth to avoid I/O bottlenecks and a small form factor differentiates the QLS1046-4GB from the solutions listed in Table 1.



Figure 1 : Qormino® QLS1046A-4GB processor and memory [Teledyne e2v].

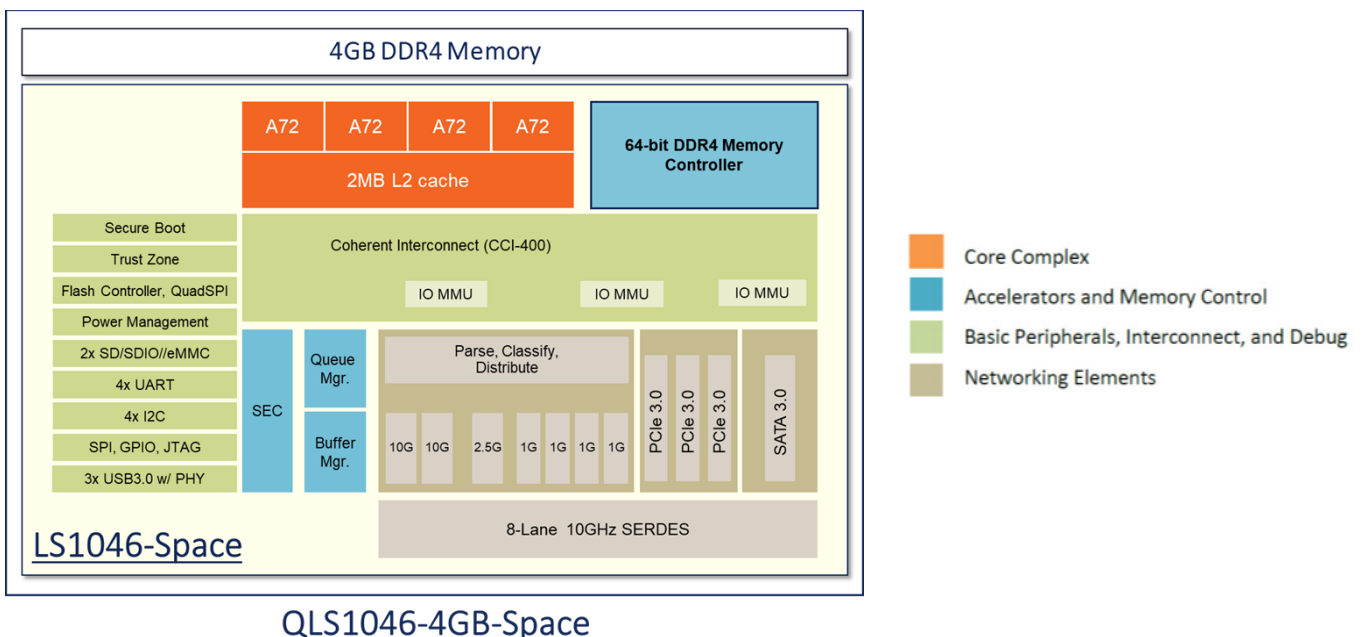


Figure 2: Block diagram of Qormino® QLS1046A-4GB.

Microprocessors and FPGAs for Intensive, In-Orbit, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



In September, Xilinx announced it will be releasing a rad-tolerant version of its Versal ACAP (Adaptive Compute Acceleration Platform). This device contains an array of AI engines comprising VLIW SIMD high-performance cores containing vector processors for both fixed and floating-point operations, a scalar processor, dedicated program and data memory, dedicated AXI channels and support for DMA and locks.

The AI tiles provide up to 6-way instruction parallelism, including two/three scalar operations, two vector reads and one write, and one fixed or floating-point vector operation every clock cycle. Data level parallelism is achieved via vector-level operations where multiple sets of data can be operated on a per-clock-cycle basis. Compared to the latest FPGAs and microprocessors, AI engines improve the performance of machine learning algorithms by 20X and 100X respectively, consuming only 50% of the power. Compared to the off-the-shelf processing solutions listed in Table 1, the AI tiles are a key distinguishing feature enabling intelligent, autonomous, in-orbit Edge processing.

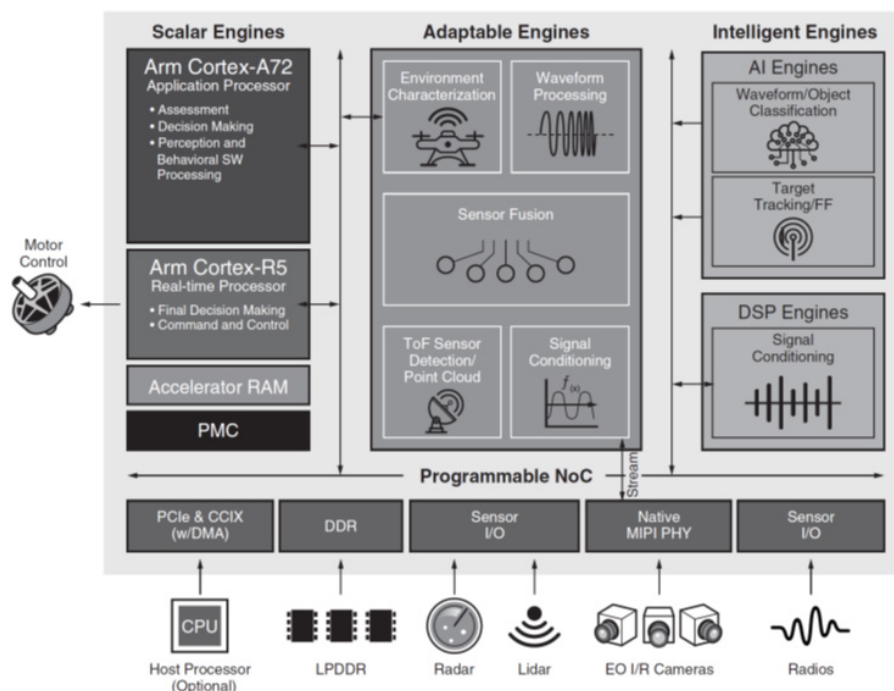


Figure 3 : Block diagram of Xilinx's Versal ACAP [Xilinx].

CONCLUSIONS

For the applications which I'm currently developing for, which type of on-board processor is better? FPGA, microprocessor or ACAP? A lot depends on how algorithms are implemented, e.g., the use of on-chip caching, the number and frequency of external memory accesses, pipelining, parallelisation and buffering. The latest space-grade devices can out-perform commercial GPUs, while also achieving higher power and price efficiency.

Microprocessors and FPGAs for Intensive, *In-Orbit*, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



For high-definition SAR video, the raw computing performance of the QLS1046-4GB together with its fast, memory interface and small form-factor makes it suitable for extracting real-time insights from Earth-Observation imaging data. DDR4 rates up to 2.1 GHz avoid traditional I/O bottlenecks.

For situational awareness, e.g., for identification of friend or foe, or for space-debris collision avoidance, the latest FPGAs such as the KU060 are able to ingest and process Tbps of data from multiple sensors with low latency in real-time to deliver ASIC-class, system-level performance. Likewise for space-exploration in-situ, resource utilisation. FPGAs process a diverse set of computationally demanding algorithms, excel in data movement, custom acceleration, bit-oriented functions and interfacing.

For object classification, AI inference and autonomous decision making to enable feature identification for late commanding of debris-retrieval spacecraft or re-configurable, cognitive transponders based on real-time traffic needs, Xilinx's ACAP would result in the most efficient Edge-based, vector compute solution. The implementation of neural networks requires TeraOPS of performance and domain-specific parallelism offered by Versal. These 7 nm devices can be power consuming so please check the early power-predictor spreadsheets to ensure they meet your allotted budget. The QLS1046-4GB may deliver deep learning at lower power dissipation and less financial cost.

Space-grade microprocessors, FPGAs and ACAPs are complementary on-board processing technologies each offering unique strengths. In-orbit, Edge-based processing requires real-time computation of large amounts of information acquired from multiple sensors right at the data source, necessitating low-latency, deterministic interfaces in a small, low-power form factor with unique thermal and reliability requirements.

When deciding on the most suitable on-board processor for intensive in-orbit, Edge-based computing, there are also time-to-market, implementation and procurement considerations, e.g., FPGAs typically need more power rails than microprocessors which means more regulators and, therefore, larger PCBs to accommodate these. FPGAs also have a reputation of being more difficult to design-in. For some projects, the time-to-orbit can be very short and OEMs will stick with existing devices from familiar suppliers to expedite hardware design. Some manufacturers don't have the skills or the time to learn new development tools or a different programming language. The six-figure, price of the latest, ultra-deep-submicron, space-qualified FPGAs is also a barrier for many OEMs, particularly those targeting lower-cost, NewSpace applications.

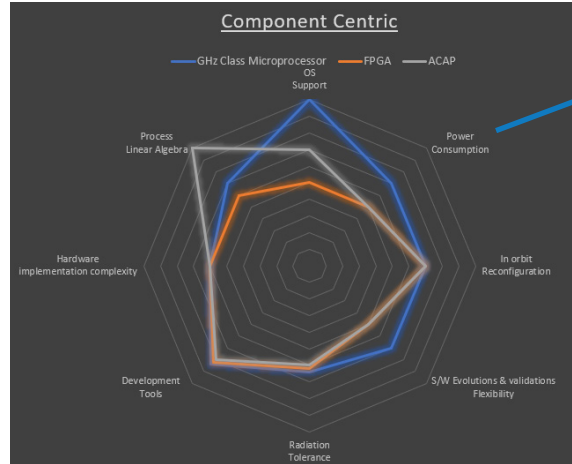
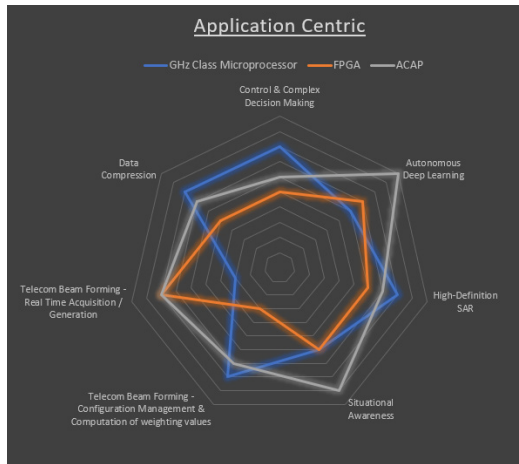
The next generation of in-orbit Edge processing will combine microprocessor, FPGA and intelligent computation to form a tightly-integrated heterogeneous platform. Multiple engine types are required because no single one is capable of optimally performing all the tasks required for an application. Scalar microprocessors are ideal for control, complex decision making and OS support, re-configurable FPGAs add flexibility to handle a diverse set of demanding algorithms, while intelligent engines optimize the calculation of linear algebra and vector arithmetic for machine learning and AI inference.

The following radar charts (Figure 4) compare the QLS1046-4GB, the latest, ultra-deep-submicron, space-grade FPGAs and ACAPs for in-orbit, EDGE-based processing:

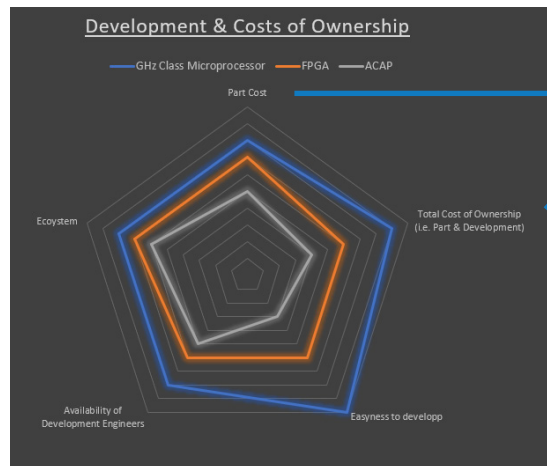
Microprocessors and FPGAs for Intensive, In-Orbit, Edge-Based Computation

Jan 2022

By Dr. Rajan Bedi



Note: Higher mark means lower consumption



Note: Higher mark means lower cost

Figure 4: Comparison of on-board processing solutions.



For further information, please contact:

Thomas Guillemain,
Marketing & Business Development,
Data Processing Solutions
thomas.guillemain@teledyne.com



For further information, please contact:

Thomas Porchez,
Application Support,
Data Processing Solutions
thomas.porchez@teledyne.com



For further information, please contact:

Jane Rohou,
MarCom Manager.
jane.rohou@teledyne.com

