A Guide to Advanced Data Processing and AI for Satellite Missions

Edwin Faier edwin.faier@xiphos.com

About Xiphos

Background

- Xiphos began operations in 1996
- Created with objective of migrating terrestrial computing and network communication techniques into harsh environments
- Xiphos uses industrial-grade components in a fault-tolerant architecture, providing robust performance costing a fraction of traditional space-grade solutions.
- AS9100-D & ISO 9001:2015 Certified

Processor Products

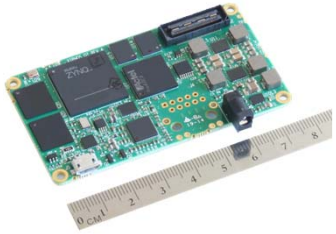

Processor Boards

Subsystems

Target Markets

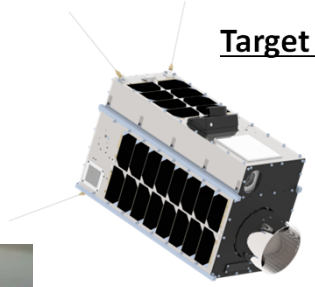

Satellites

Unmanned Vehicles

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Q7S Processor

High Performance, Low Power & Small Form Factor

- Most prolific member the Xiphos Q-Card family of low-cost embedded nodes for control, processing and interface applications with more than 250 units built
- Hybrid environment with powerful CPU and dense programmable logic, providing consistent and reliable high performance at extremely low power

CHARACTERISTICS

- Zynq-7020 SoC
  - Dual-core ARM Cortex-A9 @ 766 MHz
  - 106k flip-flops, 53k look-up tables and 220 DSP slices
- ProASIC3-based supervisor
- 1x 512 MB + 1x256 MB LPDDR2 RAM (with ECC)
- 2x 128 MB QSPI Flash (NOR)
- 2x 32 GB MicroSD
- 6-28V input
- Multiple interfaces (90 I/O, 24 LVDS pairs)
- 78 mm x 43 mm x 9 mm, 24g
- 1.5W for typical applications
- Radiation effects mitigation and 25krad TID lifetime

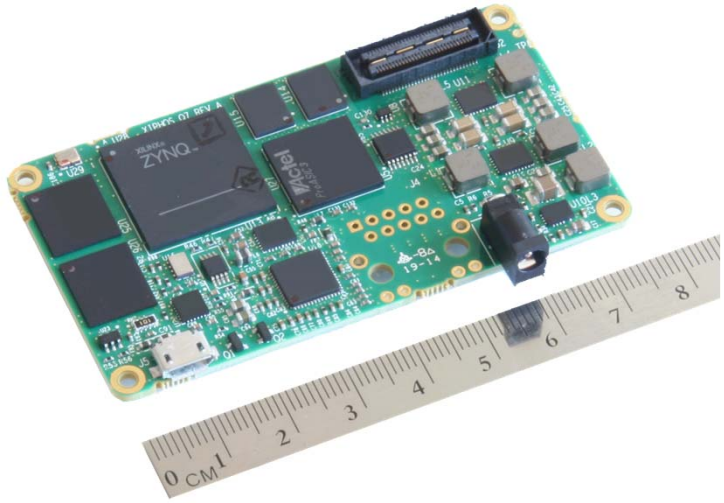

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Q8S Processor

Highest Performance Q-Card Processor

- Latest addition and highest performance member of the Xiphos Q-Card family, featuring a Multi-Processor System-on-Chip (MPSoC)
- Hybrid environment, including multi-core CPUs supported by massive programmable logic resources and a wide array of hardware interfaces at extremely low power
- Ideally suited for onboard synthetic aperture radar (SAR) processing, hyper/multispectral compression, stereo and monocular visual odometry, image registration and alignment, convolutional neural networks, advanced software defined radios (SDR)

CHARACTERISTICS

- Xilinx Zynq UltraScale+ XCZU7EG MPSoC
  - Quad-core ARM Cortex-A53 @ 1.2 GHz
  - Dual-core ARM Cortex-R5 @ 500 MHz
  - ARM Mali-400 GPU @ 600 MHz
  - 504k system logic cells, 461k flip-flops (FF), 274k look-up tables (LUT) and 1,728 DSP slices
- 4 GB LPDDR4 DRAM (with EDAC)
- 2x 256 MB QSPI Flash (NOR)
- 2x 128 GB eMMC storage
- 6-16 V input; 3.5-25 W, scalable
- 80 mm x 80 mm x 11.2 mm, 56 g
- Multiple interfaces (>130 I/O, 64 LVDS pairs, 3 Gbps transceivers)
- Radiation effects mitigation and 30krad TID lifetime

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Q8JS Processor

Highest Performance Q-Card Processor with Additional Gbps Interfaces

- Extends the capability of Q8S processor, adding support for JESD204B interfaces as well as access to external DDR3 or DDR4 memory. The Q8JS is ideal for advanced SDR applications.
- Hybrid environment, including multi-core CPUs supported by massive programmable logic resources and a wide array of hardware interfaces, including Gigabit per second interfaces supporting JESD204B and PCIe Gen 4
- Ideally suited for very wideband software defined radios (SDR)

CHARACTERISTICS

- Xilinx Zynq UltraScale+ XCZU7EG MPSoC
  - Quad-core ARM Cortex-A53 @ 1.2 GHz
  - Dual-core ARM Cortex-R5 @ 500 MHz
  - ARM Mali-400 GPU @ 600 MHz
  - 504k system logic cells, 461k flip-flops (FF), 274k look-up tables (LUT) and 1,728 DSP slices
- 4 GB LPDDR4 DRAM (with EDAC)
  - Access to external DDR3/DDR4 memory
- 2x 256 MB QSPI Flash (NOR)
- 2x 128 GB eMMC storage
- 6-16 V input; 3.5-25 W, scalable
- 80 mm x 80 mm x 11.2 mm, 58 g
- Multiple interfaces (147 I/O, 50 LVDS pairs and 16 Gbps transceivers supporting JESD204B and PCIe Gen 3)
- Radiation effects mitigation and 30krad TID lifetime

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Avionics Systems

Typical Project

Processor Boards

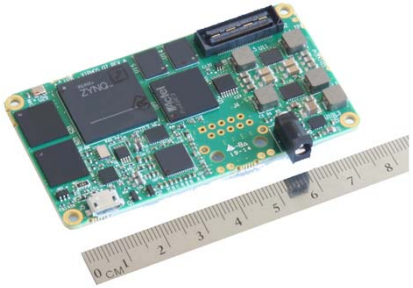

Q7 and Q8 are the successors to three previous generations of processors, proven in spaceflight since 2002

+

Custom Daughterboards

Q7 (shown with piggyback Product Integration Module) + custom daughterboard (2x Camera Link interfaces, IMU)

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Xilinx Zynq MPSoC FPGAs

Hybrid Processing and Logic Environment

- Combination of programmable logic and processing system delivers superior parallel processing power, real-time performance, fast computational speeds, and connectivity versatility

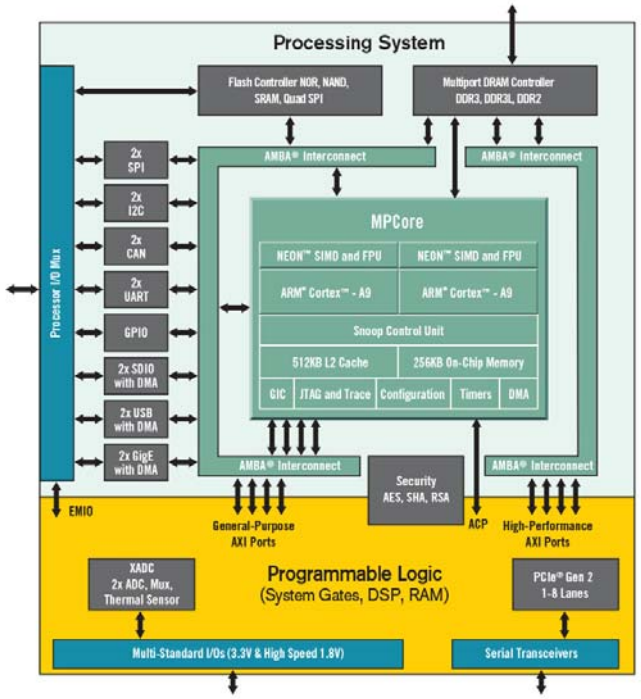

Xilinx Zynq 7020 (Q7)

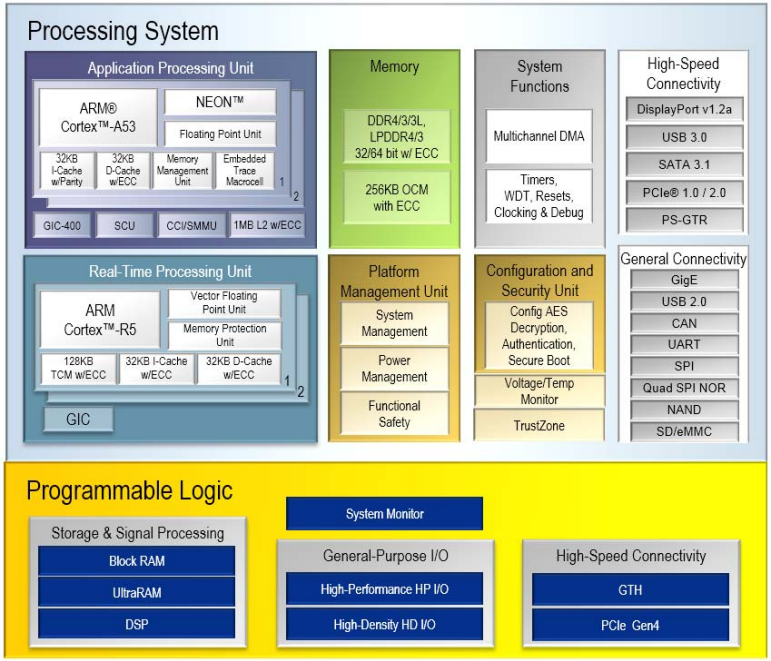

Xilinx Zynq UltraScale+ (Q8/Q8J)

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Advanced Data Processing

Why is it needed?

- Complex algorithms are increasingly required on smaller payloads, lunar vehicles, and spacecraft with constrained size, mass and power
- High resolution sensors with high data throughputs generally require real-time processing by logic prior to data being transferred to memory or CPU
- Modern sensors generate large amounts of data requiring onboard processing due to constrained downlinks and the need to maximize precious downlink usage over ground stations
- While not data processing per se, Software Defined Radios require very high-speed Digital Signal Processing (DSP) blocks to processed the digitized RF


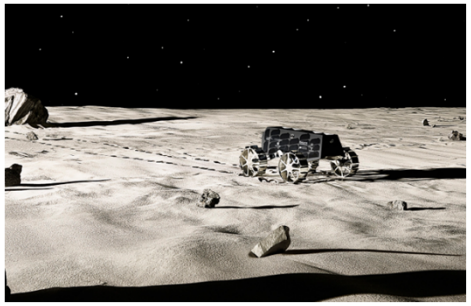
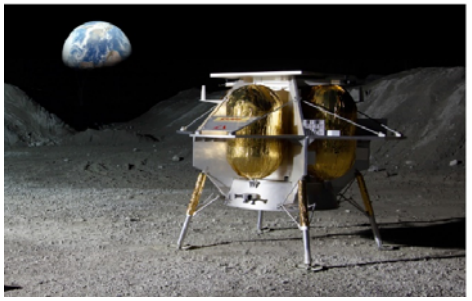


©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Advanced Data Processing

Leveraging Logic to process High-Speed Data

- High-speed interfaces available from the MPSoC FPGA allows the data to be extracted from the sensor (e.g, CameraLink, SpaceWire, LVDS, Gbps transceivers)
- Real-time data processing is performed in the Programmable Logic (PL) of the MPSoC FPGA before data is provided to the CPU (PS)
  - Allows use of standard Linux OS, simplifying and reducing cost of application development
- For high-speed imaging data processing:
  - Digital gain and offset adjustment
  - Lens distortion and image correction
  - Bad pixel correction
  - Binning
  - Coadding (adding multiple image frames to improve Signal to Noise Ratio) or Time Delay Integration (TDI)
  - Centroiding (30k frames/sec)
  - Feature detection, extraction & matching
  - Compression
- For Software Defined Radio (SDR) signal processing:
  - Various DSP cores including Fast Fourier Transforms (FFT) & Inverse FFT (IFFT), Finite Impulse Response (FIR) & Infinite Impulse Response (IIR) filters, Digital Down Converter (DDC), etc

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Hybridization of Algorithms

Hybrid Processing - sharing computation between CPUs and programmable logic

- Hybridization leverages the tight coupling of processors and logic in an MPSoC FPGA
- Logic excels for computations involving high volumes of data undergoing similar calculations, amongst other things
  - e.g., image processing, compression
- CPUs excel at other things
- A Q-Card allows both types of resources to be optimally exploited, simultaneously in a single application
- Xiphos has developed a methodology that starts with conventional C, C++ code
  - Profiling identifies where the CPU effort is being applied
  - Processing-intensive areas are evaluated to determine whether they are amenable to acceleration in logic
  - Software tools predict the resulting performance if those areas can be converted to logic
- Selected software is converted to logic, always considering return on effort
- Software is updated to utilize converted logic, as opposed to software routines
- Key algorithms implemented in logic can substantially reduce the power and mass required for space missions; this factor may be a key mission-enabler

The target performance is always known in advance

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Hybridization of Algorithms - Why it works so well

Parallels and Pipelines

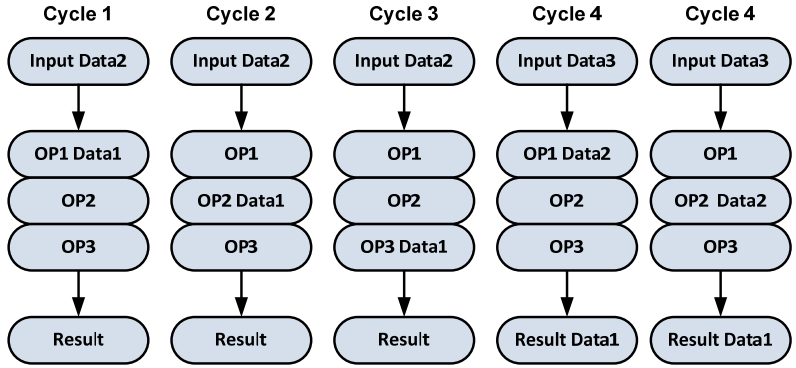

A conventional CPU does only one operation at a time

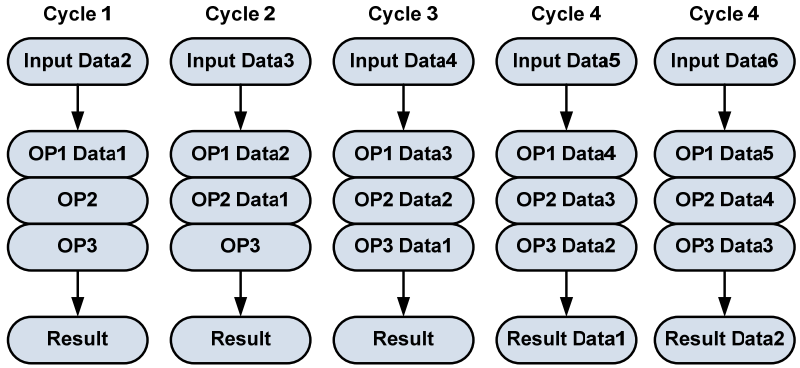

Logic can perform all operations simultaneously.
A "pipeline" is a purpose-built computing engine that never rests

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Why it works so well

Parallels and Pipelines

If one pipeline works well…

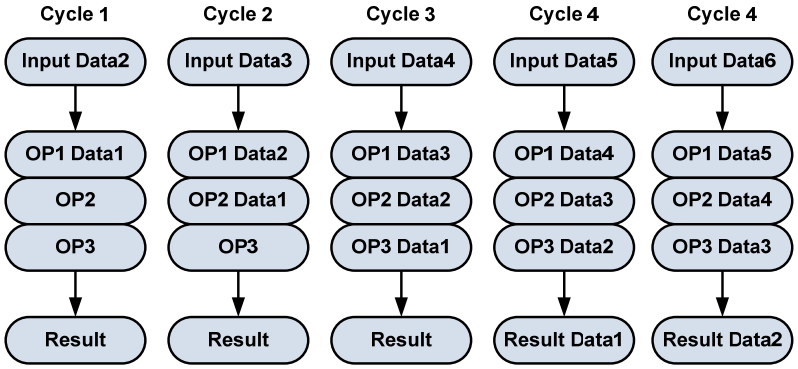

Create more!

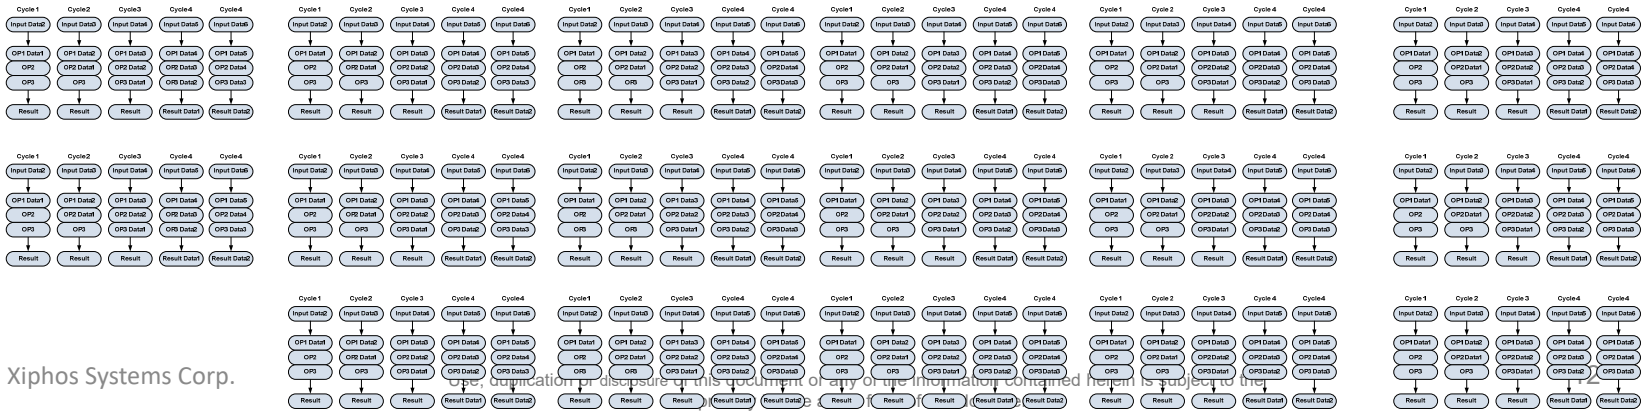

©2021 Xiphos Systems Corp. Use, duplication or disclosure of this document or any of the information contained herein is subject to the [illegible]

Hybridization of Algorithms

Sample Implementations

- *Hybridized Algorithm Examples:*
  - Location & Orientation (ICP)
  - Feature tracking (ORB)
  - Hyperspectral image compression (CCSDS 123)
  - Quantum Key Distribution

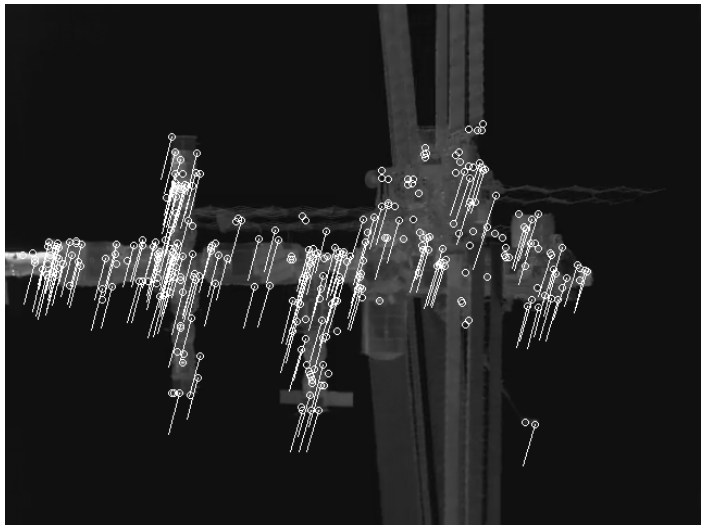

Performance Comparisons

	Performance		Power/Energy Savings	
Algorithm	**Q7 Hybrid vs PC[1]**	**Q7 Hybrid vs ARM Cortex A9[2]**	**Q7 Hybrid vs PC[1]**	**Q7 Hybrid vs ARM Cortex A9[2]**
Location & Orientation (ICP)	1X	12.9X	20.6X	9.3X
Feature Tracking (ORB)	2.3X	11.4X	49.6X	8.9X
Hyperspectral Image Compression	0.5X	9.6X	11X	7.5X
Quantum Key Privacy Amplification	4.4X	54.6X	95.6X	42.8X

1] PC based on i7-860 CPU @ 3.46 GHz

2] Algorithm run strictly in software on Q7 CPU, ARM Cortex A9 @ 750 MHz

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Advanced Data Processing (Hybridized) Application

Embedded Visual Odometry (EVO)

- EVO consists of a Q7, Camera Board and two cameras (stereo)
- Embedded algorithms include:
  - Visual Odometry
  - Hazard Detection & Avoidance
  - Disparity Map & 3D Point Cloud
- Highest performance, smallest size and lowest power using space-qualified processor

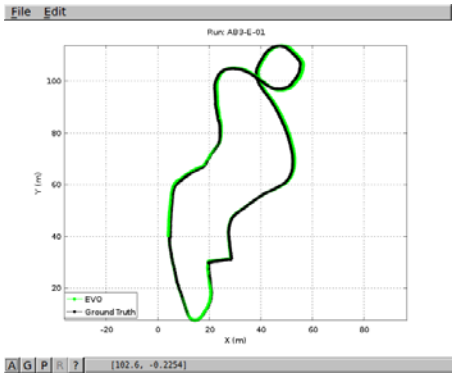

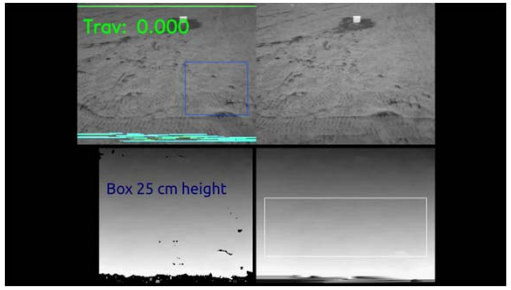

State-of-the art results, using hybrid algorithm running on Q7 processor

- Best result: average error 0.3% and peak error of 0.7%
- Typical result: average error 1.1% and peak error of 2% over 295 m traverse
- Algorithm running at 11 Hz using < 6 W (including stereo cameras)
- Rover operating at up to 6 km/h

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Artificial Intelligence (AI) Toolset

Xilinx Adaptable and Real-Time AI Inference Acceleration

- Xilinx's development platform for AI inference on Xilinx hardware platforms
- Supports mainstream frameworks and the latest models capable of diverse deep learning tasks
- Open source quantizer that supports pruned and unpruned model quantization, calibration, and fine tuning
- Vitis Unified Software Platform provides the tools to optimize, compress and compile trained AI models
- Xilinx Vitis AI has been ported to the Q8

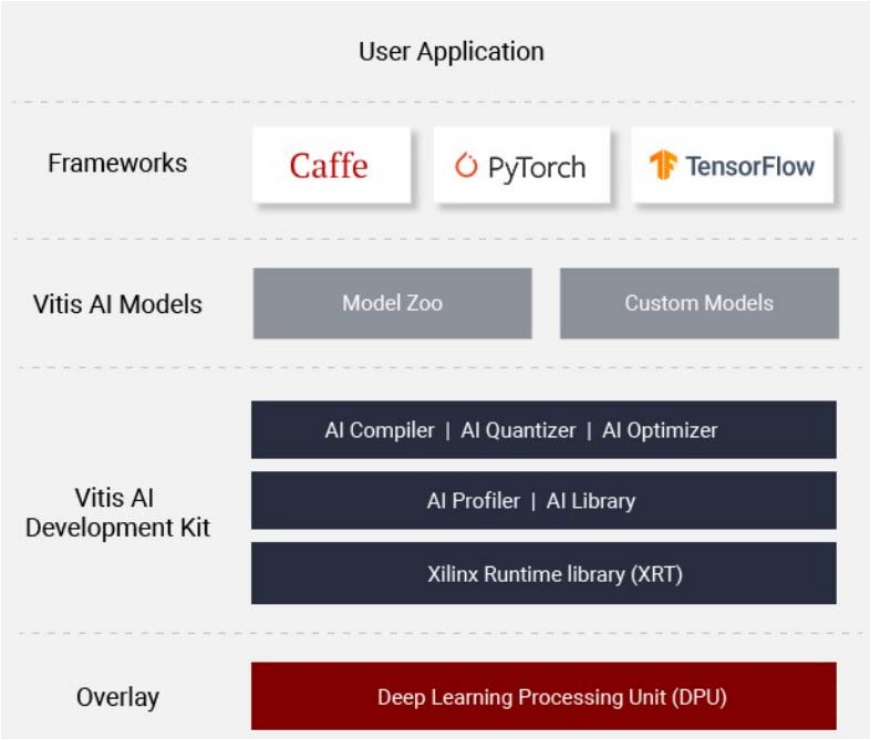

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Artificial Intelligence (AI) Toolset

Deep Learning Processing Unit (DPU) Core

- AI acceleration coprocessor
- Executes compiled application
- Direct access to RAM
- Highly configurable

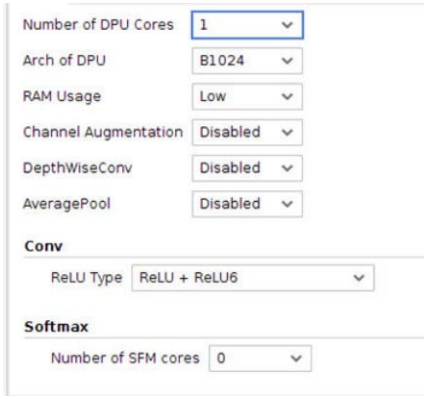

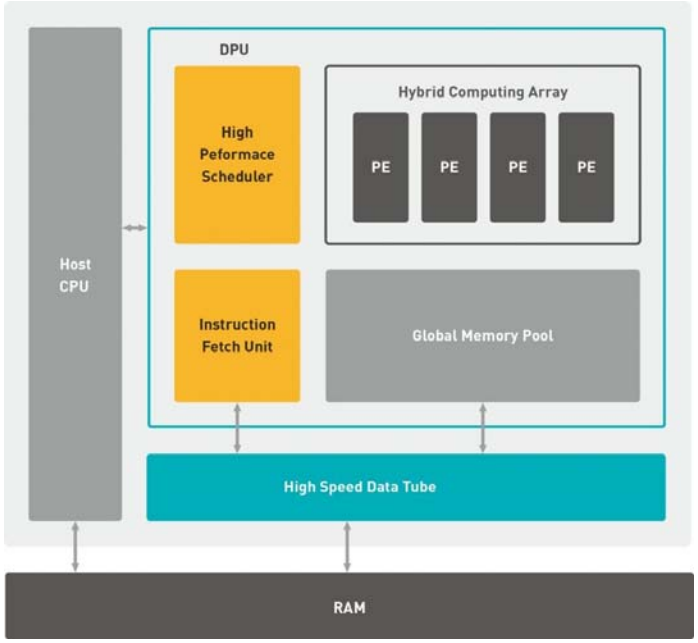

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

AI applications: Automating rover navigation and science operations

- Automated geologic scene characterization through terrain classification and novelty detection
  - Input from rover navigation cameras
  - Supports rover traversability analysis of geometric & non-geometric hazards
  - Automated science target identification and science payload operations (feature targeting)
  - Data aggregation, integration into GIS tools
- Algorithms and performance
  - State-of-the-art deep learning models
    - Semantic segmentation architecture for terrain classification
    - Using convolutional autoencoders for novelty detection
    - Automated sample identification using classification maps
  - Estimated inference performance for DeepLabV3 is 10Hz and MobileNetV2 is 15Hz inference performance on UltraScale+ using Mission Control's AI toolchain
- Upcoming use-cases:
  - Real-time inference of terrain classification and novelty detection on Q8 using analogue (in Iceland) and indoor test data (Moonyard)
  - Flight demonstration on Emirates Lunar Rover mission on a Q7, launching in 2022

Rover testing in Iceland

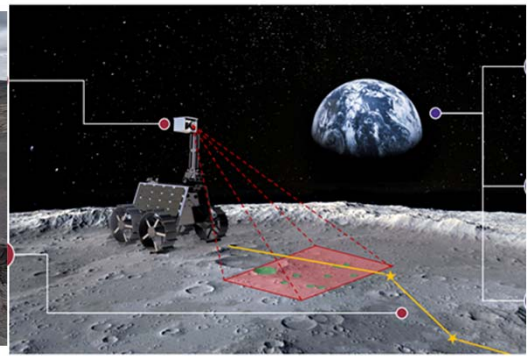

Upcoming ELR mission

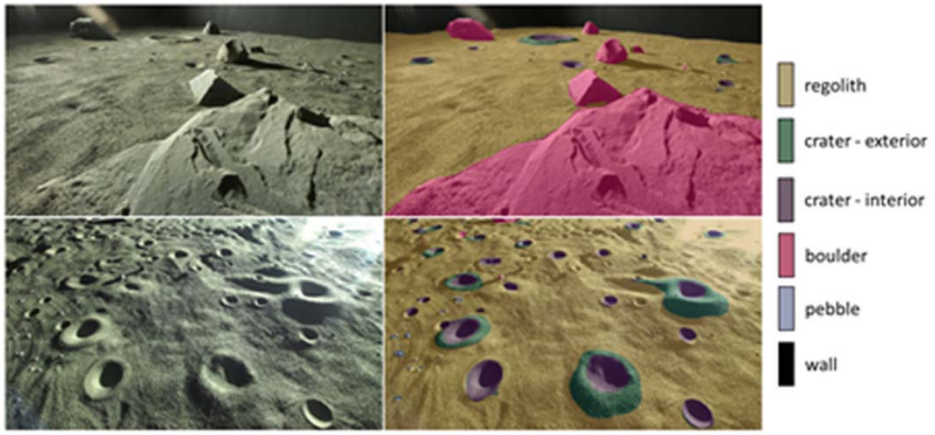

Mission Control's indoor Moonyard

©2021 Mission Control Space Services Inc. RESTRICTION ON USE, PUBLICATION OR DISCLOSURE OF PROPRIETARY INFORMATION

Mission Control's AI Toolchain

- Custom AI toolchain development to circumvent limitations of Vitis AI toolchain:
  - Extensible platform to accelerate CNN inference across multiple platforms: Xilinx SoCs, CPUs, GPUs, other low powered SoCs
  - Semantics independent of ML framework, so any input framework can be used
- Compiles neural networks from the standard NNEF format to a bytecode which can be interpreted on device
  - NNEF allows the toolchain to import networks from more frameworks than supported by Vitis AI and include a wider range of supported operations
  - Compiler can be retargeted for other platforms (i.e. GPUs)
- Bytecode interpreter runs on CPU and uses JIT-compilation to dispatch low-level instructions to a custom FPGA overlay
  - Two-level bytecode allows the interpreter to seamlessly switch between CPU- and FPGA-based computation based on the current context
  - Models/weights can be swapped on-the-fly since no changes are required to the FPGA's bitstream
  - Tunable FPGA overlay depending on targeted hardware platform
- Bytecode design allows operations to execute efficiently and minimizes the logic resources required to run large networks
- Currently being implemented on Zynq UltraScale+ and 7000 SoCs and will be used in flight mission in 2022
- Contact michele@missioncontrolspaceservices.com for more information

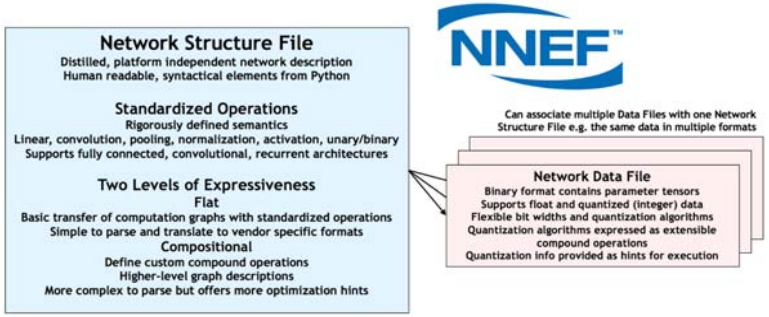

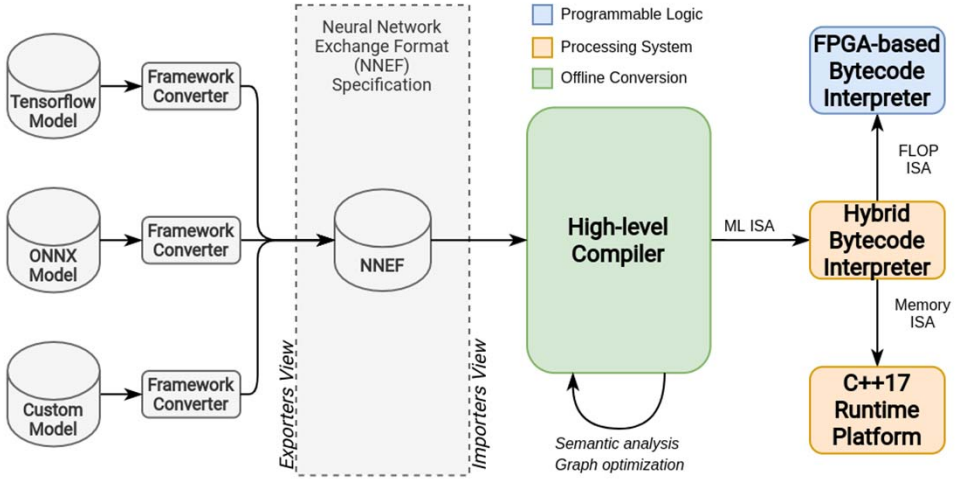

©2021 Mission Control Space Services Inc. RESTRICTION ON USE, PUBLICATION OR DISCLOSURE OF PROPRIETARY INFORMATION

Conclusion

- Hybrid processing and logic environment of MPSoC FPGA's can be leveraged to perform advanced data processing
  - Xilinx Zynq 7020 on Q7, Zynq UltraScale+ on Q8 and Q8J
- Real-time data processing is performed in the logic of the MPSoC FPGA before data is provided to the CPU
  - Allows use of standard Linux OS (non-RT OS)
- Hybridization allows the execution of complex algorithms in real time using low-power space processors
  - Can enable a mission
- Inference can be done very quickly and inexpensively on an FPGA, enabling advanced AI data processing
  - Xilinx Vitis toolchain available for UltraScale+ FPGAs (Q8/Q8J)
  - Custom toolchains can also be developed or are available through third parties

©2021 Xiphos Systems Corp.

Use, duplication or disclosure of this document or any of the information contained herein is subject to the Proprietary Notice at the front of this document.

Contact

- **For more information, please contact:**

Edwin Faier
President, Xiphos Systems Corporation
Email: edwin.faier@xiphos.com

- **Or visit www.xiphos.com**

©2021 Xiphos Systems Corp.

contained herein is subject to the
nent.