Commercial Off-the-Shelf GPU Qualification for Space Applications

Wesley Powell, Michael Campola, Teresa Sheets - NASA/GSFC
Abigail Davidson, Sebastian Welsh - Cubic Aerospace

National Aeronautics and Space Administration

Space Technology Mission Directorate

Game Changing Development Program

Relevance to NASA Needs

- With increased sensor data rates, and limited downlink capability, NASA missions require increased onboard processing throughput for applications ranging from synthetic aperture radar (SAR) data reduction to hyperspectral image processing and recognition.
- Further into the future, onboard autonomy and artificial intelligence for robotic and crewed missions will present even greater demands for onboard processing.
- Graphics Processor Units (GPUs) offer an attractive processing architecture to meet these increased demands due to their massive parallelism.
- As no radiation hardened GPU devices currently exist, any near term GPU-based onboard processors must use commercially available devices.

Project Objectives

- Characterize the capability of GPUs to meet the demands of a candidate onboard processing application, thereby demonstrating their ability to improve mission performance, reduce spacecraft SWaP, and potentially enable new missions.
- Evaluate the radiation tolerance of capable COTS GPU devices to determine their suitability for spaceflight applications and understand any mitigations that are needed.

Candidate GPU Application – Onboard Data Processing for the Advanced Energetic Pair Telescope (AdEPT)

- AdEPT is a high resolution medium-energy (5-200 MeV) gamma-ray polarimeter telescope consisting of multiple Three-Dimensional Track Imager (3-DTI) modules. Data must be processed to remove cosmic rays and other noise, leaving only the medium-energy gamma-ray science data, in the time it takes to accumulate new data in the detector, ~ 20-50ms.
- The raw data stream is produced at a rate of 0.2 Gb/s up to 2.5 Gb/s depending on the size and number of 3-DTI modules. The full detector will consist of eight 3-DTI Modules producing eight streams of for a total of 20 Gb/s.
- One of the algorithms being tested will first eliminate the low-energy gamma-ray interactions (blips) by removing low amplitude values, and then eliminate the Cosmic Rays (CRs) by checking for any straight tracks that cross the detector's boundary. This processing must be done for both the xz and yz planes, when corresponding tracks are found, that data will be removed. The resulting data set will contain a refinement of data that will aid in the detection of medium energy gamma-ray tracks.
- The figure to the right is a graphical representation of one-half of the raw data(xz plane), that on-board software processing must reduce to science data.

Results Thus Far

- The software for the AdEPT algorithm has been ported to the GPU processing system, using CUDA (Compute Unified Device Architecture) routines to interface with the GPU, and was re-structured to take full advantage of GPU core processing.
- A similar algorithm was run on a Tilera multi-core processing system using pthreads, demonstrating real-time processing of input streaming data at maximum of 0.8 Gb/s per 4 cores. Scaled by 2 for 1.7 Gb/s, scaled by 4 for 3.2 Gb/s, scaled by 8 for **5.5 Gb/s**.
- For the NVIDIA Jetson TK1, the GPU version of the algorithm, processing the same amount of data, is currently running at **13.97 Gb/s**.
- While further optimization is needed, these interim results demonstrate the utility of GPUs to perform high bandwidth onboard instrument data processing.

Radiation Testing Plans

- Radiation tests will be performed on the Nvidia Tegra Jetson-TX2 and Tegra Xavier GPUs using commercial development boards.
- As the Xavier boards are just now becoming available, this has required radiation testing to be deferred to FY19.
- An initial Total Ionizing Dose (TID) test will be performed at a Co-60 source.
- Devices that successfully pass TID testing will undergo Single Event Effects (SEE) testing at the Texas A&M University (TAMU) Cyclotron Institute Radiation Effects Facility.
- SEE tests will indicate if the devices are susceptible to destructive events that would prohibit their use for space applications.
- These tests will also provide insight into the "signatures" of non-destructive events and how they could be mitigated on a spaceflight GPU processor board.

Future Plans

- Prepare COTS development boards and conduct radiation tests.
- Explore options for further accelerating AdEPT algorithm (CUDA streams, etc.) on the Jetson TK1 board.
- Port and benchmark the AdEPT algorithms on the faster Jetson TX2 and Xavier boards.
- Investigate additional spaceflight GPU applications for prototyping.

- ***Pave the way for the development of a spaceflight GPU processor board for future NASA missions.***

GPU Performance Metrics

Device	Cost ($k)	Power (W)	Floating Point Calculations (MFLOPS)	Performance Figure of Merit (MFLOPS/W)	Cost Figure of Merit (MFLOPS/$k)	Mass Figure of Merit (MFLOPS/kg)
BAE RAD750 (general purpose rad hard SBC)	200	5	266	50	1	133
MPC8548E (high performance rad tolerant processor)	20	33	8,000	250	13	4,000
Virtex 5QV (SIRF) (high performance rad hard FPGA)	75	15	80,000	5,000	1,000	40,000
Tegra K1 (tablet GPU with up-screening)	**10**	**10**	**326,000**	**32,600**	**32,600**	**163,000**
Tegra X2 (tablet GPU with up-screening)	**20**	**15**	**750,000**	**50,000**	**37,500**	**375,000**
Tegra Xavier (tablet GPU with up-screening)	**25**	**20**	**1,300,000**	**65,000**	**52,000**	**650,000**

Notional GPU Processor Board

AdEPT Data Processing Steps

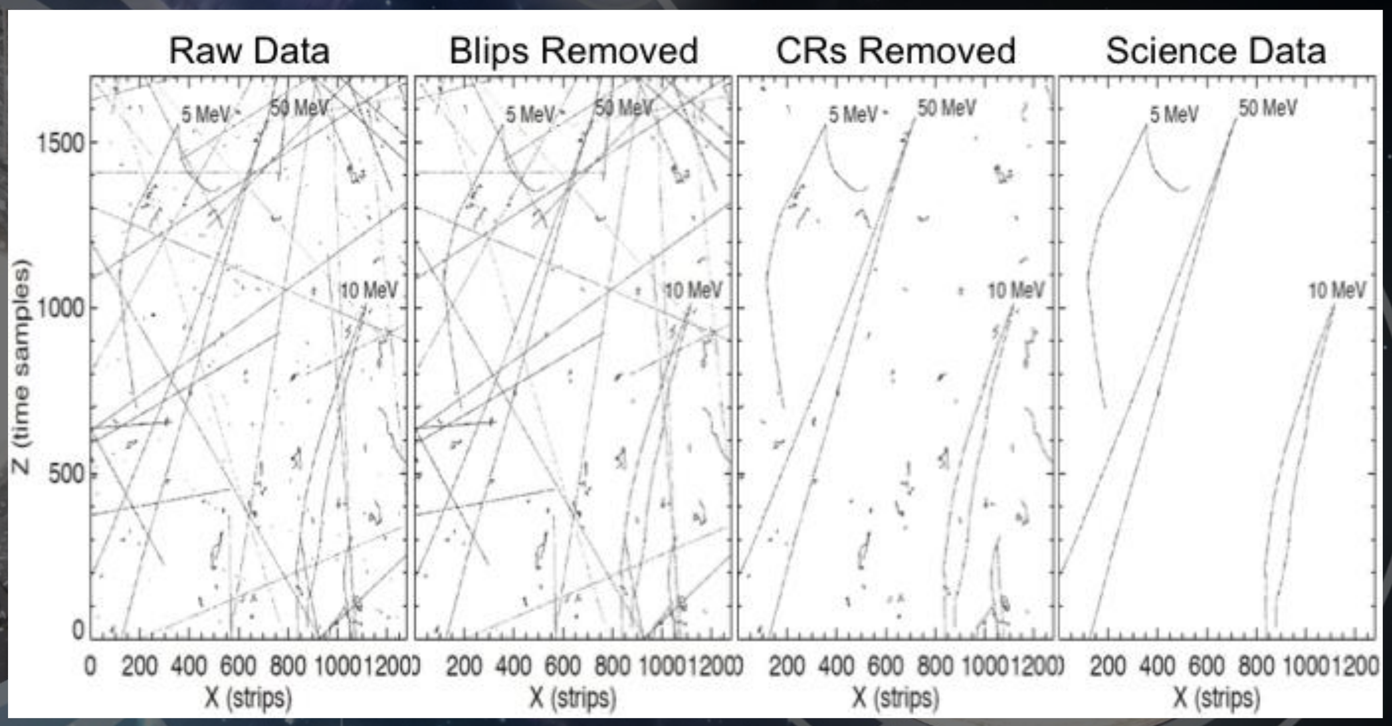

TAMU Heavy Ion Beam

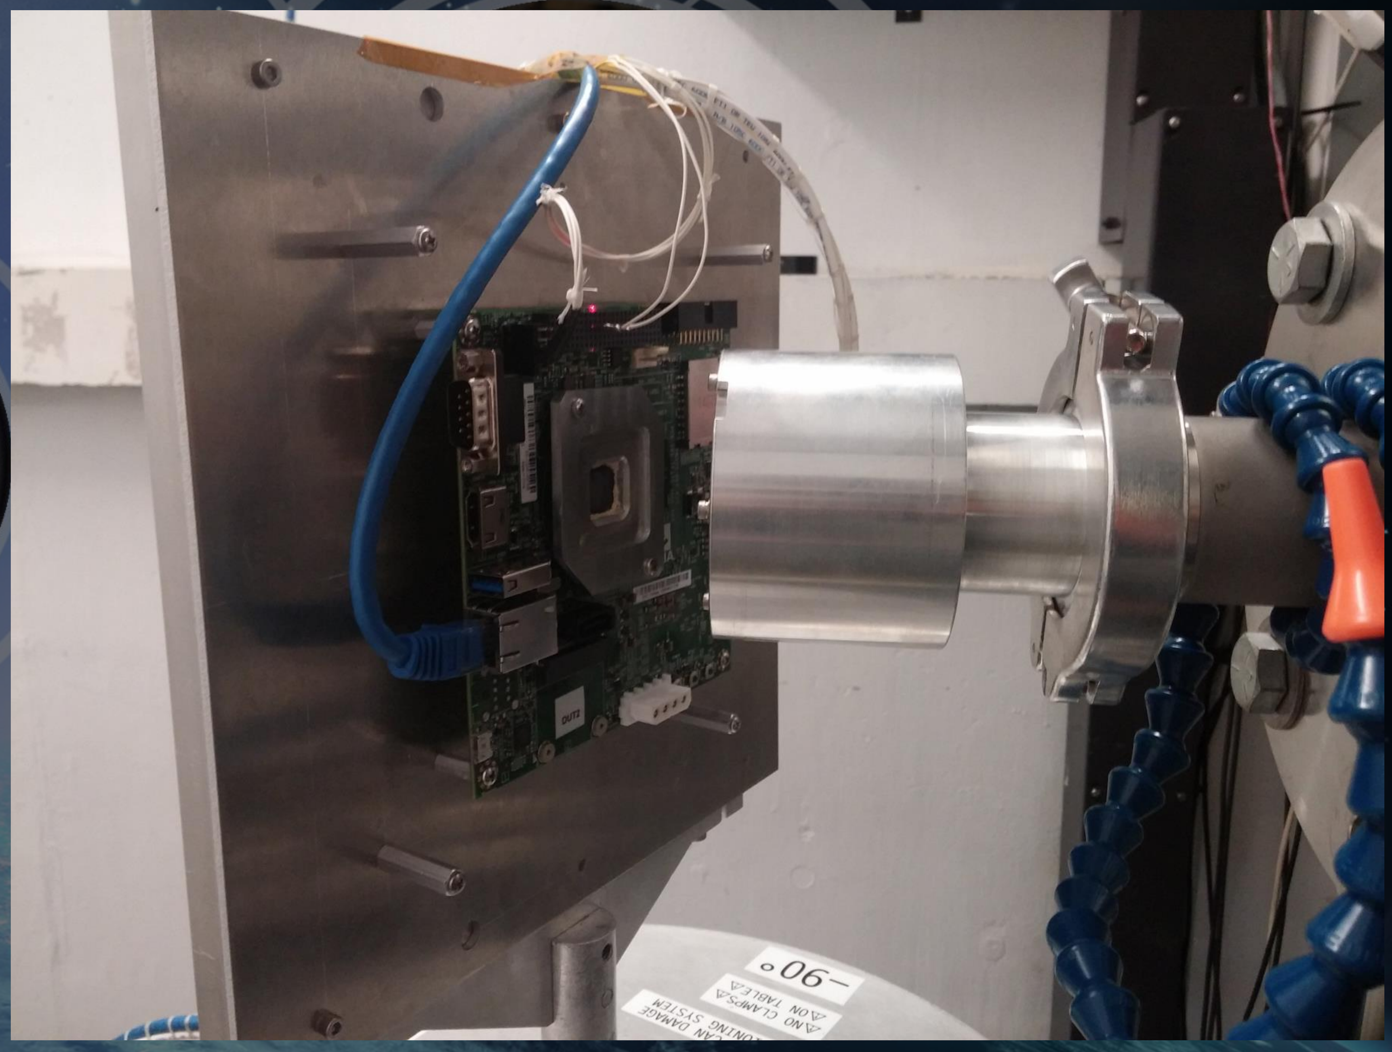

To be presented by Wes Powell at Space Technology Mission Directorate (STMD) Annual Program Review (APR), Cleveland OH, September 26, 2018.