Towards an Integrated GPU Accelerated SoC as a Flight Computer for Small Satellites

Caleb Adams
Department of Computer Science
University of Georgia
1510 Cedar St.
Athens, GA 30602
770-314-8422
pieninja@uga.edu

Allen Spain
Department of Electrical and Computer Engineering
University of Georgia
314 Barnett Shoals Rd.
Athens, GA 30605
770-833-1784
avs81684@uga.edu

Jackson Parker
Department of Electrical Engineering
University of Georgia
1510 Cedar St.
Athens, GA 30602
706-542-2374
jop80923@uga.edu

Matthew Hevert
Department of Mechanical Engineering
University of Georgia
1510 Cedar St.
Athens, GA 30602
303-250-2594
Matthew.Hevert@uga.edu

James Roach
Department of Computer Science
University of Georgia
1510 Cedar St.
Athens, GA 30602
770-845-4787
jhr08134@uga.edu

Dr. David Cotten
Department of Geography
University of Georgia
210 Field St.
Athens, GA 30602
706-542-2374
dcotte1@uga.edu

***Abstract*—Many small satellites are designed to utilize cutting edge technology with the goal of rapidly advancing space based capabilities. As a result, many components take advantage of developments from the miniaturization of smartphone technology. Within the past 2 years, the UGA Small Satellite Research Laboratory has extended this concept into embedded GPUs for high-performance processing in LEO. Here we showcase advances in our research of high-performance space-based computation by integrating a traditional flight computer with existing miniaturized GPU/SoC systems. Such a system paves the way for many of NASA's goals that require space based AI, neural networks, computer vision, and high performance computing. Our system fits a standard CubeSat PC/104+ form factor, and implements many standard protocols such as I2C, SPI, UART, and RS422. The system also has several GPIO pins, 2 USB-C ports, a micro USB port for flashing, an Ethernet port, and a micro SD card slot for development. Additionally, the system is designed to be modular, so that GPU accelerated SoCs can be stacked to form a distributed system. For our primary computer, which handles I/O and initializes processes on the SoC, we choose to use the radiation tolerant Smart Fusion 2 SoC with an ARM Cortex-M3 processor and a FPGA. In addition to this primary computer, we use the Nvidia Tegra X2/X2i as the GPU/SoC workhorse. The primary computer and the TX2i are designed to share memory space with peripherals mounted onto the board, so that no significant file transfer is required between the subsystems. Additionally, Nvidia's Pascal architecture enables GPU-CPU or GPU-GPU communication without PCIe, enabling dense interconnected networks for monitoring and computation. To address thermal concerns, we cap the TX2i's power draw at 7.5 Watts, provide recommendations for thermal interface materials, and ensure that the primary computer only enables the GPU/SoC when parallel computation is specifically requested. Furthermore, radiation mitigation techniques are explored with ECC, software mitigation techniques, and aluminized kapton sheets. In conclusion, this system is a step towards a miniaturized high-performance flight computer well suited for future computational demands.**

TABLE OF CONTENTS

978-1-5386-6854-2/19/$31.00 ©2019 IEEE

1. INTRODUCTION

Overview

The University of Georgia's Small Satellite Research Laboratory (SSRL) is utilizing a Nvidia TX2i Graphics Processing Unit (GPU) within the Multi-view Onboard Computational Imager (MOCI) Satellite. Although the system utilizes a GPU, an additional On Board Computer (OBC) is still required for control and communication with core avionics. The UGA SSRL has developed a board, the Core GPU Interface (CORGI), that is capable of interfacing the Nvidia TX1, Nvidia TX2, or Nvidia TX2i into a PC/104+ compliant CubeSat [1][2][3]. In this paper we discuss the previously designed CORGI board and elaborate on lessons learned, including how the CORGI is being used to combine an OBC and an Nvidia TX2i into one board. We call this new board, still undergoing development, the Accelerated Flight Computer (AFC).

The GPU (Nvidia TX2/TX2i) being used is a complete System on a Chip (SoC), capable of running GNU/Linux on an ARMv8, with a 256 core Pascal GPU. For an exhaustive list of TX2/TX2i capabilities see the user guide [4] Currently the TX2/TX2i utilizes CUDA 9.1 [4].

The Microsemi Smart Fusion 2 (SF2) SoC serves as the central control node for the AFC. The SF2 SoC contains a 166MHz ARM Cortex-M3 Processor with embedded flash and an FPGA[5]. The SoC has as static power draw of 7mW and has Single Event Upset (SEU) protected and tolerant eSRAM and DDR bridges [6]. The SF2's FPGA has demonstrated that it operates well in heavy ion and proton radiation environments, though some mitigation is still required [7]. Additionally, past high-performance space processors have

used the SF2 SoC with success [8]. The architecture of the SF2 SoC makes it ideal for space environments and its wealth of documentation makes us confident that it will work well with our design.

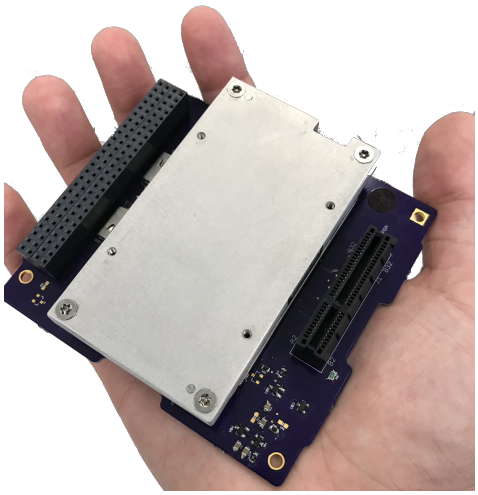

Figure 1. A previous iteration of the AFC concept, known as the CORGI (Core GPU Interface).

Additionally, we seek to adhere to the IEEE Std 1156.4-1997 standard for Spaceborne Computer Modules. This standard provides requirement levels for thermal performance, pressure, shock, vibration, and radiation [9].

2. Form Factor

The CubeSat form factor is one of the primary considerations in the design of the AFC. The form factor complies with available deployers such as the Poly Picosatellite Orbital deployers (P-POD), and the JEM Small Satellite Orbital Deployer (J-SSOD), and is modular to accommodate the payload and support I/O both spatially and functionally [10][11]. The CORGI (see Figure 1) board was a design benchmark which satisfied requirements, such as mission and development requirements. However the design still required optimization and additional features to provide a comprehensive accelerated heterogeneous computing platform[8]. CORGI's successor, the AFC, has useful features such as bi-directional logic shifters to convert the 1.8V (CMOS) logic from the onboard Nvidia Jetson TX2i [4] to the Smart Fusion 2 3.3v (LVTTL) logic. Shifters are required for serial data transfer, GPIO flipping, and IC enabling. The AFC will maintain a PC/104+ form factor (90 x 96 mm)[3] with an ergonomic cutout to accommodate the high speed LVDS expansion header (Samtec LSHM-120-04.0-L-DV-A-N-K-TR) mounted beneath the TX2i module. Tight integration of both CPU and GPU allows support of all aspects of satellite performance, such as power, command and data handling, attitude determination, payload etc. The Smart Fusion 2 ARM Cortex-M3 SoC (M2S150T-1FCG1152) is a low power and low profile FCBGA (Flip Chip Ball grid Array) combining program accelerators: FPGA, and CPU [8]. The SF2, primary computer contains an ARM Cortex SoC and is equipped with a flash based FPGA. The Nvidia TX2/TX2i's (50 mm x 87 mm) compact design provides high density compute, and parallelized Pascal microarchitecture. A 400-pin compatible board-to-board connector will interface with the TX2i module via a 8x50 connector (SEAM-50-02.0-S-08-2-A-K-TR). The Jetson TX2i was added to the AFC due to its form factor and recent usage in high performance computing (HPC) in industrial environments, long life, and ECC memory [4]. The AFC form factor maximizes compute, while eliminating general purpose parts sold with COTS On Board Computers (OBC). This has the added benefit of simplicity of development, and reduced quiescent current draw.

Table 1. Processors

Smart Fusion2: ARM Cortex-M3 SoC (M2S150T-1FCG1152)
NVIDIA Jetson TX2: -2x Denver ARM Cortex-A57 -256-core Pascal GPU

Table 2. Memory

SF2: 4x256 DDR3 Memory Bank
TX2: 8GB LPDDR4 and 32 GB ECC support
Both: -2 Gb (Cypress CYRS16B256) -1 Gb SPI Flash on SPI0 via FPGA MSS -1Gb Flash on SPI1 via FPGA fabric

Table 3. IO Interfaces

PC/104+, 12C, SPI, GPIO
Expansion header (QSPI)
2x RJ-45
2x USB 3.0 Type C
Micro USB (FTDI)

3. Hardware Design

Hardware

As shown in Figure 3, H1 and H2 are the main PC/104+ CubeSat headers. These route peripheral components onto the bus, as well as routing the PC/104+. The H1 and H2 headers on the AFC also act as pass-throughs, so there is no break in communication or power. The Smart Fusion 2 shares all primary communication lines and acts as the control node for the command and data handling (CDH) subsystem of the satellite. In addition to two RJ45 Ethernet receptacles the AFC board also makes a QSPI expansion header available for added support and 2x USB Type C for debugging the TX2i.

Prototyping

The SmartFusion 2 advanced development kit [12] was used for the prototyping of the integrated CPU-GPU board. The SmartFusion 2 advanced development kit contains a wide range of headers (SPI, I2C, GPIO, UART, etc.) [12]. The versatility of this development kit makes it ideal for prototyping.

The Nvidia TX1/TX2 development kit includes a developer board that is compatible with the Nvidia TX2i [4]. The developer board, connected to the Nvidia TX2/TX2i, can be interfaced with the Smart Fusion 2 advanced development kit via 10/100/1000 PHY Ethernet connection, SPI, LVTTL UART, I2C, PCIe, or GPIO[4][12].

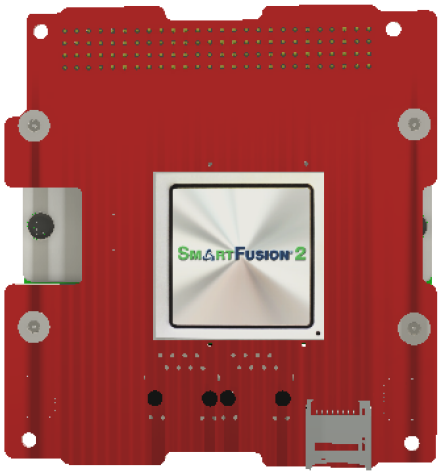

Figure 2. AFC 3D design back

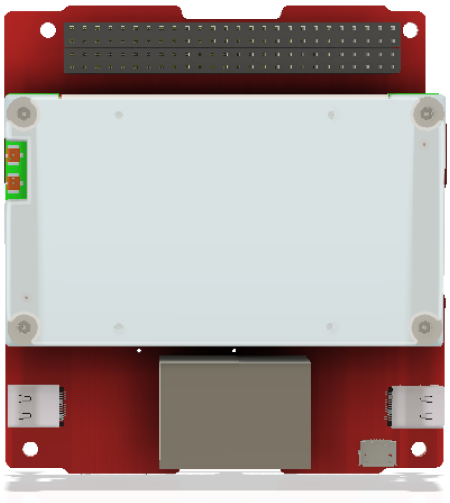

Figure 3. AFC 3D design front

4. Architecture

Hardware Architecture

With ubiquitous computing and 10/100/1000 Ethernet Switching, Ethernet provides wired communication for the satellite bulk data transfer between the SF2 and the one or more TX2i. This design contains carrier level Ethernet whose data link and physical layer protocol are described by the IEEE 802.3. [13] This design utilizes a node switch which acts as a physical medium.The SF2, being the primary distributer of the workload, design contains SEU Protected/Tolerant Memories i.e: eSRAMs, DDR3. Its robust nature played a considerable roll in its selection as the primary processor on the accelerated flight computer. In the event of one or all of the TX2is destruction, the SF2 will be able to perform I/O monitoring and control duties without a co-compute device. This option, although functional, will diminish payload capabilities. In response, the system-on-chip (SoC) field programmable gate array (FPGA) microcontroller subsystem (MSS) includes a watchdog which provides a system which needs no operating system to detect and respond to lockups caused by heavy ion bombardment [12]. The SF2 will utilize 4 x 256 DDR3 (MT41K256M8DA-125 IT:K) external memory banks providing volatile memory. In addition, the TX2i modules and SF2 share RAD NOR SPI flash memory (Cypress CYRS16B256). The SF2 will allow mediated read/write access with the purpose of sharing mutually relevant files. This functionality is defined in hardware via MUX whose selection is determine by GPIO pins from SF2 SoC, active high.

Software Architecture

This section details how software will make use of the hardware system presented above in order to provide a modular fault tolerant accelerated computing platform. The Smart Fusion 2 will act as the main control node for the whole system and will make use of four main interfaces to interact with the TX2is. These interfaces will be a Gb Ethernet local network, GPIO, SPI, and shared flash memory that the Smart Fusion 2 and TX2is will access via multiplexed SPI connections. The ethernet line will be used by the Smart Fusion 2 to distribute computation tasks and data to the TX2i's. These ethernet communications will take the form of multicast messages, as each TX2i will need to receive and direct messages for communication to individual TX2is. The computation tasks to be distributed to each TX2i will take the form of CUDA code. The TX2is will make use of the ethernet line or GPIO to kick the watchdogs on the Smart Fusion 2 FPGA and to log their progress or telemetry. The Smart Fusion 2 will use the SPI accessed shared flash memory as its primary persistent memory where it will store tasks and data to distribute to the TX2is. The TX2is will make use of the shared flash memory by writing computational results to it.

5. Radiation and Thermal Mitigation

System Hardware Mitigation and Shielding

For radiation and thermal mitigation in Low Earth Orbit (LEO) we recommend utilizing the Dunmore Aerospace "Satkit" which contains the standard STARcrest MLI materials cut into manageable sizes for small satellite developers. This allows one to develop their own thermal protection blankets according to various mission requirements. The kit includes an outer layer material, inner layer material, first surface tape, and clear polyimide tape[14]. This kit is optimal for small satellite systems as it is small, cheap, and easily customizable. More specifications of the materials are listed below (See Figure 4) and values were calculated from the Solar Radiation, Earth's IR radiation, and Albedo Radiation equations [15].

$$q = G_s \alpha_s \cos\phi \tag{1}$$

$$q = \sigma T_e^4 \alpha_I R F e \tag{2}$$

$$q = G_s (AF) \alpha_s F e \cos\theta \tag{3}$$

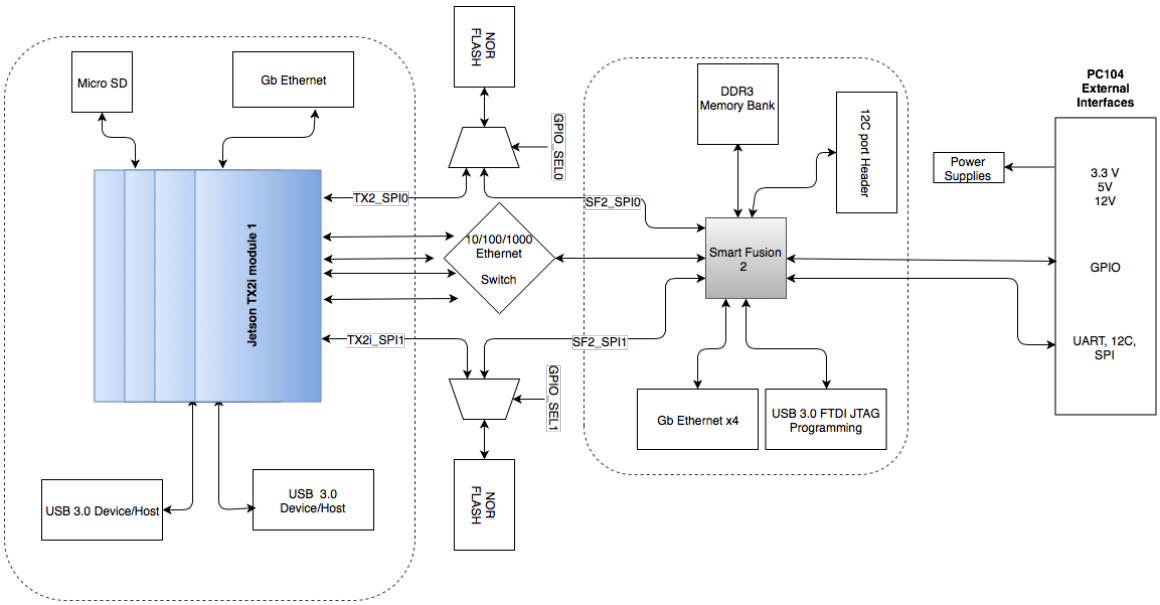

Figure 4. **The Hardware Design of the Cubesat GPU/SoC-CPU system**

The heat flux due to the LEO sources using equations 1, 2, and 3 was calculated with the help of data from literature [15]. Values have an expected error of ± 0.4. Values for the solar radiation were determined using the mean values provided by data from the World Radiation Center in Davos Switzerland [16].

The necessary thickness of the radiation shield was calculated in part using values provided by Dunmore. The primary source of heat flux is solar, these values vary on an annual basis due to the Sun's elliptical orbit. Meaning that the maximum and minimum amount of flux expected from the Sun would range between 1322 and $1414 w/m^2$. The target ambient temperature inside the CubeSat is 293.15 K (20 degrees C). This is calculated with the following formula:

$$Q = \frac{\Delta T K A}{L} \tag{4}$$

Where Q is the heat flux into the system, K is the thermal conductivity of the material, A is the area in meters squared, and L is the thickness of the radiation shielding. Assuming a thermal conductivity of 0.014 and a ΔT of $101K$. The expected required thickness of the radiation shielding is $1.03438 \cdot 10^{-7} m^2$. Radiation from Free Molecular Heating (FMH) was determined to be negligible for the stage in which the CubeSat would be launched from the ISS. FMH is almost exclusively encountered during launch ascent just after the booster's payload fairing is ejected.

Radiation Mitigation

The radiation shielding thickness (See Equation 4) is also driven by the 1997 IEEE Standard for Environmental Specifications for Spaceborne Computer Modules [9]. We design for level I radiation in preperation for the LEO environment. An additional benefit of the Dunmore Aerospace "Satkit" is that it meets this standard while accruing minimal mass gains.

Proper shielding does not have to incur heavy mass gains, and in fact shielding that is too thick can increase the effects of some kinds of radiation events. This is due to the higher levels of secondary particles created when a high-energy GCR particle impacts a thick shield [17]. If necessary, a

Table 4. **Recommended Dunmore Aerospace Satkit Part**

VDA / 200 GA Polyimide	
Tensile Strength	24000 psi
Elongation	50%
Thickness	50.8 micron
Density	1.42 g/cc
Yield	13.8 $m \cdot m/kg$
Weight/Area	72 g/m^2
Operating Temp	-250 - 400 C
Metalization	99.99 % pure aluminium
VDA / 200 GA Polyimide	
Tensile Strength	26000 psi
Elongation	110 %
Thickness	6.35 micron
Density	1.39 g/cc
Yield	124.9 $m \cdot m/kg$
Weight/Area	8 g/m^2
Operating Temp	-250 - 150 C
Metalization	99.99 % pure aluminium
VDA / 25 GA PET / VDA, Embossed & Perforated	
Thickness	76.2 microns
Yield	10.4 $m \cdot m/kg$
Weight/Area	96 g/m^2
Operating Temp	-40 - 220 C
Metalization	99.99 % pure aluminium
100 GA Polyimide / 966 PSA	
Thickness	76.2 microns
Yield	10.4 $m \cdot m/kg$
Weight/Area	96 g/m^2
Operating Temp	-40 - 220 C
Metalization	99.99 % pure aluminium

properly designed shield may act as a heat sink due without exceeding mass limitations. This reduction of a device's operating temperature can greatly reduce the risks posed to that device by radiation [17].

When working with Commercial Off The Shelf (COTS) electronic components for use in the space environment, some of the problems that must be addressed are data integrity performance and accuracy in high radiation environments. In our design, the TX2i, one of the Tegra SoCs, is one of the more vulnerable parts with respect to this, due to its highly dense hardware design[18]. As stated above, much can be done at a hardware level but often times additional software mitigation is needed.

One of the major problems radiation presents for software is its effect on memory. Overtime as bit flips [15] [9] aggregate, they will corrupt information and make some systems unusable. To address these issues, in the AFC we plan on modifying the file system and bootloader. We will triplicate the file system on the TX2i so that any part of it which is corrupted will be able to use the principle of triple modular redundancy (TMR) to automatically correct damage. In addition to this, we will use the U-Boot bootloader software to modify the bootloading process with TMR [19]. The bootloader will store 3 copies of the operating system (OS) image for the TX2i and a hash for each image. At boot time the bootloader will recalculate the hash for each OS image it attempts to load and compare the calculated hash against the stored hash to determine if the OS has been corrupted. If corruption is detected then boot loader will attempt to load the next OS image. If all OS images are determined to be corrupted, the bootloader shall attempt to construct an uncorrupted image by bit voting between the corrupted images.

Our design uses the SF2 FPGA as a trusted control node due test results showing high levels of radiation tolerance [7]. Thus, we plan to have the Smart Fusion 2 FPGA operate as a watchdog for each TX2/TX2i. If the Smart Fusion 2 detects that one of the TX2/TX2i's has anomalous power draw levels (over 7.5 Watts), it will send a command over ethernet (via the SF2 SoC) commanding the TX2i to reduce GPU usage until the power level stabilizes.

Thermal Mitigation

Simulations have shown, given a large enough mounting structure, that the Nvidia TX2 is capable of dissipating heat effectively. To achieve this goal it is imperative that the TX2's Thermal Transfer Plate (TTP) is adequately interfaced into the heat dissipating mass[1]. To interface the TTP with the heat dissipating mass we highly recommend the use of a low thickness and high conductance thermal interface material (TIM) that adheres to the NASA standards for collected volatile condensible materials ($\leq$ 0.1% CVCM) and total mass loss ($\leq$ 1% TML) [20]. Previous findings have suggested that the Carbice Space TIM is ideal for these purposes due to its low thickness ($0.065mm$) and high conductance ($13,330 Wm^{-2}K^{-1}$) [1].

It is important to note that all thermal simulations are run with a heat load according to subsystem power draws and assume 0% efficiency for a worst case scenario. The steady state thermal simulation clearly shows that the GPU's core temperature has been lowered with the addition of the TIM. The overall maximum has been brought down from around 160°C, to under 50°C. While this model did not include the entire spacecraft structure as a conduction medium, this would only serve to decrease the temperature further, as conduction is a much faster heat transfer mechanism than radiation.

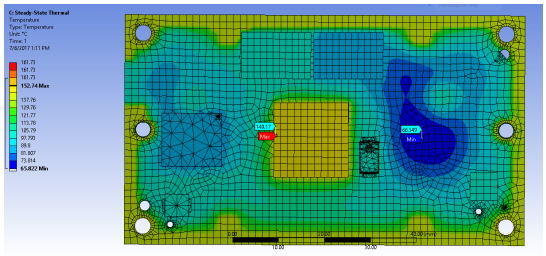

Figure 5. The Nvidia TX2 daughter-board with the Parker Series SoC Exposed. No TIM is used in this simulation.

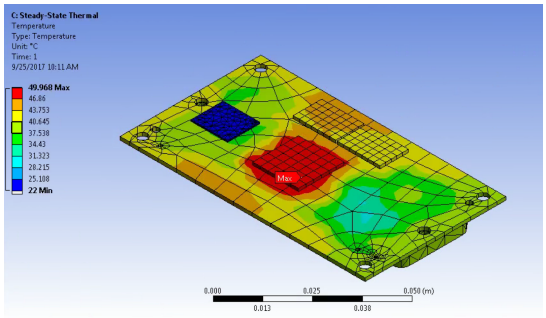

Figure 6. The Nvidia TX2 daughter-board with the Parker Series SoC Exposed. TIM is used in this simulation

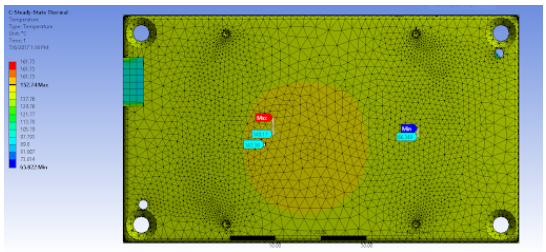

Figure 7. The Nvidia TX2 with TTP integrated and TIM used.

However, some caution must be used with these results. The thermal environment of the satellite is inherently transient, so the ambient temperature assigned to the model here is likely inaccurate. This also means a "steady state" analysis might not necessarily be appropriate. However, the purpose of this analysis is not to provide conclusive thermal information about the AFC, rather to serve as a data point for design, and to show that under extremely approximate conditions, the AFC will be able to operate in orbit.

6. Conclusion and Future Work

The UGA SSRL's CORGI design, the starting point for designing the AFC, has made the usage of the Nvidia TX2/TX2i more feasible in LEO. The recommendations provided in the

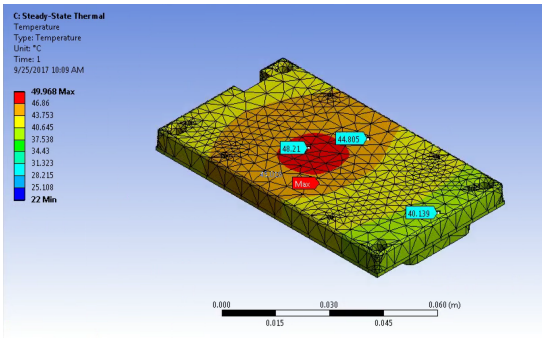

Figure 8. The Nvidia TX2 with TTP integrated, no TIM is used.

paper regarding radiation and thermal mitigation provide a starting point for those who wish to utilize the Nvidia TX1, TX2, or TX2i in LEO. Additionally, when designing a fully integrated system with the Nvidia TX2/TX2i, the use of a SmartFusion2 SoC FPGA is highly recommended. It is well understood, well documented, radiation tolerant, and operates on low power. The combination of the SmartFusion2 with the Nvidia TX2/TX2i has the potential to meet many of NASA's goals that require space based AI, neural networks, computer vision, and high performance computing[21].

Future work will entail testing of the aforementioned components. The system shall be able to halt and restart actions on the subsystem and utilize Paralleled Thread Execution assembly to checkpoint GPU operations. Additional methods for thermal dissipation will be simulated and those mentioned above will be tested in our thermal vacuum chamber. Radiation testing (heavy ion and proton) will be performed during heavy computational load. The system will be used to form a modular computational cluster and distributed computation will be further tested in our vacuum chamber while the system is under heavy load.

APPENDICES

A. CODE AND RESOURCES IN THIS PAPER

The resources used in this paper can be requested at any point by contacting the authors of this paper.

ACKNOWLEDGMENTS

The authors would like to thank the Georgia Space Grant Consortium for funding these GPU research projects and the Air Force Research Laboratory's University Nanosat Program for giving us tremendous opportunities and for funding the projects that led us to this point. Additionally the authors would like to thank Erick Gavilanes for making the initial design for the CORGI in 2017 and getting us to a point where we could perform iterations and follow up research. Thanks to Casper Versteeg for providing significant knowledge and results relating thermal mitigation. Nir Patel, Nicholas Heavner, and Michael Buzzy did great work by contributing to the thermal models. A big thanks Justin Heimerl, for supplying CAD and working on electrical design and to Sydney Whilden and Adam King for doing initial radiation research and laying the foundations for this paper in that respect. A special thanks to Dr. Deepak Mishra for helping support the University of Georgia's Small Satellite Research Laboratory from the beginning. Another thank you to Roger Hunter.

REFERENCES

[1] C. Versteeg, "Thermal management and design of high heat small satellite payloads," *32nd Annual AIAA/USU Conference on Small Satellites*, 2018.

[2] C. Adams, "A near real time space based computer vision system for accurate terrain mapping," *32nd Annual AIAA/USU Conference on Small Satellites*, 2018.

[3] (2003) Pc/104-plus specification. [Online]. Available: www.winsystems.com/wp-content/uploads/specs/PC104PlusSpec.pdf

[4] (2017) Jetson tx2 developer kit user guide. [Online]. Available: www.developer.nvidia.com/embedded/jetpack

[5] (2015) Arm cortex-m3 processor technical reference manual. [Online]. Available: www.infocenter.arm.com/help/topic/com.arm.doc.ddi0337e/DDI0337E_cortex_m3_r1p1_trm.pdf

[6] (2018) Product brief smartfusion2 soc fpga. [Online]. Available: www.microsemi.com/document-portal/doc_view/134215-ug0557-smartfusion2-soc-fpga-advanced-development-kit-user-guide

[7] (2014) Igloo2 and smartfusion2 65nm commercial flash fpgas interim summary of radiation test results. [Online]. Available: www.microsemi.com/document-portal/doc_view/134103-igloo2-and-smartfusion2-fpgas-interim-radiation-report

[8] F. Bruhn, "Introducing radiation tolerant heterogeneous computers for small satellites," *32nd Annual AIAA/USU Conference on Small Satellites*, 2015.

[9] "Ieee standard for environmental specifications for spaceborne computer modules," *IEEE Std 1156.4-1997*, 1997. [Online]. Available: www.ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=603627&tag=1

[10] (2016) 6u cubesat design specification. [Online]. Available: http://org.ntnu.no/studsat/docs/proposal_1/A8%20-%20Cubesat%20Design%20Specification.pdf

[11] "Nanoracks cubesat deployer (nrcsd) interface definition document (idd). nr-nrcsd-s0003," 2018. [Online]. Available: www.nanoracks.com

[12] (2016) Smartfusion2 soc fpga advanced development kit user guide. [Online]. Available: www.microsemi.com/document-portal/

[13] (2008) Ethernet theory of operation. [Online]. Available: www.microchip.com/downloads/en/appnotes/01120a.pdf

[14] Dunmore aerospace satkit. [Online]. Available: www.dunmore.com/products/satkit-mli.html

[15] G. D. Badhwar, "The radiation environment in low-earth orbit," *Radiation Research*, vol. 148, no. 5, pp. S3–S10, 1997. [Online]. Available: www.jstor.org/stable/3579710

[16] "World radiation data centre," 2018. [Online]. Available: www.re3data.org/repository/r3d100010177

[17] R. H. Maurer, M. E. Fraeman, M. N. Martin, and D. R. Roth, "Harsh environments: Space radiation en-

vironment, effects, and mitigation," *Johns Hopkins APL Technical Digest*, vol. 28, no. 1, pp. 17–29, 2008.

[18] E. Wyrwas, "Body of knowledge for graphics processing units (gpus) nepp-bok-2018," 2018. [Online]. Available: www.nasa.gov/sites/default/files/atoms/files/2017-8-1_stip_final-508ed.pdf

[19] "The denx u-boot and linux guide (dulg) for canyonlands," 2018. [Online]. Available: www.denx.de/wiki/DULG/Manual

[20] R. R. J. Roe, "Standard materials and processes requirements for spacecraft," *Technical Report NASA-STD-6016A*, November 2016.

[21] (2017) Nasa strategic technology investment plan. [Online]. Available: www.nasa.gov/sites/default/files/atoms/files/2017-8-1_stip_final-508ed.pdf

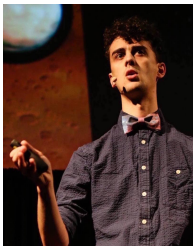

Caleb Adams *Caleb Adams is a Graduate student in Computer Science and co-founder of the UGA Small Satellite Research Laboratory, which now has 2 satellite missions. Caleb's research mostly focuses on high performance computation systems in space environments for AI and Computer Vision. He has also worked on audio telemetry systems for Orion at NASA, was also a google glass beta tester, and a TEDxUGA speaker.*

Allen Spain *Allen Spain is currently a M.S. Student in the School of Electrical and Computer Engineering at the University of Georgia. He specializes in electronic circuit design, and is the primary designer of Electronics Interface Boards for the UGA Small Satellite Laboratory. He is currently a Research Assistant at the University of Georgia focusing on Biomimetic Photonic algorithms.*

Jackson Parker *Jackson Parker received his B.S. in biological engineering in spring of 2018 from the University of Georgia, where he is currently a Masters student in Electrical and Computer Engineering. Current research activities and interests include accelerated computation, machine learning, algorithmic development, and computer vision. He plans on furthering his knowledge in those areas of interest while working on his M.S. and then pursue a PhD in computer science.*

Matthew Hevert *Matthew Hevert is currently pursuing a Bachelors of Science in Mechanical Engineering at the University of Georgia with an anticipated graduation in 2020. He has been working at the Small Satellite Research Laboratory performing work in optics, thermodynamics, and mechanical radiation mitigation. He plans on continuing his education after graduation in optical sciences and engineering where he one days hopes to help design the next generation of space-based telescopes.*

James Roach *James Roach is an undergraduate computer science student graduating from the University of Georgia in May 2019. He has been working at the small satellite research laboratory since fall 2015. He has also worked on small satellites for Space Dynamics Research Laboratory. Current research interests include accelerated computation, the effect of radiation on electronics, and fault tolerance in cubesat platforms.*

David Cotten *David L. Cotten, Ph.D., received a B.S in Physics from Louisiana State University in 2005. He currently serves as the manager and Co-PI of the UGA's Small Satellite Research Laboratory and is an Assistant Research Scientist at the Center for Geospatial Research in the Geography Department at UGA. He graduated from UGA in 2011 with his doctorate in Physics and Astronomy. His research as a Post-Doctoral Associate focused on surface-atmosphere exchange, and he is currently using remote sensing (multispectral/hyper spectral sensors) and micrometeorology techniques to quantify carbon storage in wetland regions at both the local and regional scales. Dr. Cotten is also using unmanned aerial vehicles, air photos, and satellite imagery to build 3D models of terrestrial objects using photogrammetric Structure from Motion.*