

Guidance on Standardizing GPU Test Approaches

Edward Wyrwas, Member, IEEE, Kenneth A LaBel, Michael Campola and Martha O'Bryan

Abstract-- A standardized test method has been created to characterize and stress graphics processing units (GPU) during radiation effects testing.

I. INTRODUCTION

When it comes to defensible scientific results, repeatability is always key to prove out experimental hypotheses. The logistics behind repeatable testing are often guided by a standard or test method which defines the environmental or electrical stress conditions to which the device under test (DUT) is subjected. While the stress conditions are defined, seldom does a test method lay out the complexities with respect to setting up the physical test. Often, however, the vagueness in the methods permits a wide variation in how a test is performed and subsequently how the results are collected. This opens up even greater uncertainty in the comparability of results across different DUT within a technology family, especially for complex devices such as Graphics Processing Units (GPUs).

Semiconductor reliability is already a challenge for terrestrial applications of high performance computing (HPC), aerospace, defense and automotive electronics. Typically, these devices can be evaluated for performance marginality and long-term risks associated with radiation effects using simulation, first and second order approximations, or environmental testing. While these methodologies have historically worked for monolithic devices, such as diodes or MOSFETs, it is increasingly more difficult to assess modern microprocessors as they contain processing elements, memories, control structures, passive chip components and an interconnect structure. Tests need to traverse multiple test vectors to accomplish a decent level of confidence in the test results. Further, new generations of microprocessor devices tend to have fine-tuned performance and other capabilities when compared to previous generations in the device family. Standardized tests should disregard hardware and driver optimizations as newer products will

almost always perform better (smaller size, lower thermal design power (TDP) per transistor, more FLOPS, more memory) than their predecessors.

Radiation testing mirrors the logistics and electrical monitoring associated with reliability and qualification testing. Ideally, a component or component on a minimalist daughter board should be used with a pin and socket interconnect to a carrier board. This is often the case with discrete components (i.e. diodes) when undergoing radiation tests. Practical repeatability is often overlooked in test creation due to resource constraints and haste. The monitoring should take place from the carrier board for consistency and mitigation against handling damage. Power supplies should be controlled and monitored with software so that timings or intervals between operational steps are consistent between each DUT and each test run. Electrical control using network-based software controlled relays permits rapid creation of test benches without intensive development. These are a few broad examples of system-level control that facilitates autonomy. Within this paper, we discuss a standardized approach to radiation testing of GPU devices which facilitates an apples to apples comparison between device generations and device types across various vendors. Table 1 compares three devices.

TABLE I
COMPARISON OF GPUS

| Part Model | GTX 1050 | APQ8096 | Jetson TX1 |
|-----------------|----------------------|-----------------------|----------------------|
| Manufacturer | nVidia | Qualcomm | nVidia |
| Technology | 16nm FinFET | 14nm FinFET | 20nm CMOS |
| REAG ID | GSFC 17-039 | JPL | GSFC 16-038 |
| Board Model | EVGA 02G-P4-6152-KR | Intrinsyc Open-Q 820 | 699-82180-1000-100 U |
| Packaging | Flip Chip, BGA, PCB | Flip Chip, BGA, SOM | Flip Chip, BGA, SOM |
| Memory Capacity | 2GB GDDR5, >8GB DDR4 | 3GB LPDDR4 | 4GB LPDDR4 |
| Performance | 1.86 TFLOPs | 0.50 TFLOPs | 1.00 TFLOP |
| Test Bench OS | Windows 2016 | Android 6 Marshmallow | Linux for Tegra |

II. HEAT SINKS AND COOLING

An advanced microprocessor testing strategy was started using nVidia's GTX 1050 GPU. The nVidia GTX 1050 GPU is a graphic coprocessor used in a Consumer-Off-The-Shelf (COTS) computer. The carrier board is connected to the computer motherboard via a PCI-express x16 slot. The GPU die itself, is the DUT and is located underneath the unit's heat sink. All GPUs that are greater than one watt need a cooling solution. Low power devices can use a passive cooler, but mid to high power devices need active cooling (i.e. heatsink with fan). Working with Michael Campola, of NASA GSFC's Radiation Effects and Analysis Group, the stopping range of heavy ions was calculated for various heat sink materials using information from The Stopping and Range of Ions in Matter (SRIM) website and spreadsheet arithmetic. For heavy ion and laser testing, the die is thinned to 150um and polished

Edward J Wyrwas, Lentech, Inc., 7500 Greenway Center Drive, MTC I, Suite 505, Greenbelt, MD 20770. Work performed for NASA Goddard Space Flight Center (GSFC) and NASA Electronic Parts and Packaging (NEPP) Program. (telephone: 301-286-5213. E-mail: edward.j.wyrwas@nasa.gov.

Martha V. O'Bryan is with ASRC Federal Space and Defense, Inc. (AS&D, Inc.), 7515 Mission Drive, Suite 200, Seabrook, MD 20706. Work performed for NASA Goddard Space Flight Center (GSFC) and NASA Electronic Parts and Packaging (NEPP) Program, e-mail: Martha.V.Obryan@nasa.gov.

Michael J. Campola, and Kenneth A. LaBel are with NASA/GSFC, Code 561.4, Greenbelt, MD 20771, e-mails: Michael.J.Campola@nasa.gov, and Kenneth.A.Label@nasa.gov

such that energy transfer and particle interaction can take place at the active transistor layers within the die. Unfortunately, the factory stock cooling solution had to be removed from the card to expose the die during operation.

Proton testing had previously been conducted using the factory stock cooling solution as the proton range is sufficient to traverse the entire device thickness with minimal energy loss. To conduct heavy ion and laser tests, a custom tooled cooling solution was created to permit access to the thinned die from the top side while absorbing the heat through the backside of the printed circuit board. This orientation permitted nominal operation from both the DUT GPU and a control GPU (with stock cooling solution) within the test bench. The cooling solution created for GPU testing is also a verified solution to test COTS CPU devices such as an AMD Ryzen microprocessor. The image shows such a setup with a 400W cooler plate connected to an AMD motherboard with an AMD Ryzen 1700X CPU being configured for radiation testing by Carl Szabo (NASA GSFC, NEPP, AS&D), which was de-lidded prior to operation (shown on the bottom of Figure 1). An alternative cooling method, from the primary side of the PCB, can also be employed using a thin (20mil) thermally conductive sapphire window and heat sink combination. The extra material, while thermally beneficial, adds unnecessary material and poses physical clearance risks with the beam line around the DUT at some test facilities.




Fig 1: Cooling solution on delidded AMD CPU

The DUT preparation described allows an ideal situation to be developed. A direct path is created to the active layers through thinning and polishing. The cooling solution allows the device to operate under load while maintaining a temperature appropriate for the test (i.e. 20°C). The die can be thermally imaged and superimposed onto an optical image of the active regions (mirrored in the case of a flip-chip device, of course) to provide a feature map. A laser test can correlate radiation response from a proton or heavy ion test to a very specific area on the die and be marked on the feature map. Each of these characterization activities can be performed in a controlled fashion - torque specification, software interface with drop down menus, automated electrical measurements, etc.

III. SOFTWARE TEST VECTORS

Unlike discrete GPU coprocessors, some GPUs take the form of an IP block or embedded engine within a System on

Chip (SoC) device. The best example of this is a smart phone's SoC such as the Qualcomm Snapdragon 820 which contains a Qualcomm Adreno 530 GPU. Within this device are various functional blocks which can be exercised with software payloads. Nvidia's Jetson TX1 SoC is provided on a modular printed circuit board connected to a main carrier board by a connector. Within it are ARM CPU cores and an nVidia GPU which can be accessed similarly to the discrete GPU coprocessor. The point here is that while the packaging is different, each one of these GPUs can be tested using the same standardized code.

The test vectors created for these GPU-related microprocessor tests exercise specific circuit structures within the GPU device such as control logic, cache and other memories, and the processor pipeline within its cores. Except in the case where a single event functional interrupt (SEFI) happens, the test vectors employed in these test plans were created to upset specific target structures, monitor any electrical anomalies (if present) and record any computational errors resulting from the upset.

Three types of payloads have been created for the GPU test bench: Neural Network, Math-Logic and Colors. The neural network is a convolutional neural network (CNN) type, which can avoid processor optimizations that recursive neural networks (RNN) primarily benefit from. Math-Logic uses mathematics and conditional logic statements to exercise memory hierarchy. The Colors payload assesses corruption in the output image presented to a monitor or display.

- Convolutional neural network (CNN) to identify land usage objects using a dataset modified from [4 for use with a "You Only Look Once" (YOLO) algorithm for object identification in still images and live stream video. The CNN was configured to be trainable on both GPU and CPU microprocessor types. Twenty one image categories were identified across the dataset. Figure 2 shows three such categories. The YOLO algorithm provides an accuracy rating and the most likely image category when presented with an image. The categories are: agricultural, airplane, baseball diamond, beach, buildings, chaparral (vegetative desert), dense residential, forest, freeway, golf course, harbor, intersection, medium residential, mobile homes, overpass, parking lot, river, runway, sparse residential, storage tanks and tennis court.
- Mathematical and logic payloads such as calculating Pi, polynomial arithmetic, Markov permutations (such as folding protein algorithms) and algebra are leveraged to fill the computational and memory components of the device while preventing hardware optimizations to manipulate the software bit-stream. These math payloads target the layers of the memory hierarchy of the device.
- Graphics, texture and color rendering tests have been developed. Graphics memory tends to be directional in that it behaves as read-only. The simplest test allows monitoring of this memory by triggering a pixel color change with automatic screen compare for pixel-change identification. The most complex of these tests performs

a burn-in to the rendering logic of the device. Pixel corruption or display artifacts are monitored and recorded using the test bench.




Fig 2: Neural Net Training Images from the Land Use Dataset

Multiple CNN configurations were tested across three different hardware configurations. The algorithm was multi-faceted in that the network needed to be smart (accurate guess), contain deep thought (the computational time greater than the accumulated data transfer duration), and be intensive (consume as many device resources as practicability allows). There are multiple knobs that can fine tune the operation of the network permitting the payload's efficacy in regards to exercising specific device structures to be scalable to the DUT which permits repeatable testing. This also allows comparable and defensible tests to take place across part numbers.

Neural networks are one type of payload that can be scaled for hardware that it supports. Unfortunately, there is not yet one neural network platform that is fully cross platform for hardware (e.g. Intel, ARM, AMD, nVidia). Therefore, a basic Math-Logic or Colors payload can test multiple generations of a device that doesn't support neural networking. The best examples are OpenCL and OpenGL which are open source computational and graphics libraries, respectively. Both of these standards are supported across most vendors' hardware (discrete components and embedded IP in system on chip) and are supported on all modern forms of Windows and Linux operating systems. The payloads using math or colors are also tuned to be scalable and efficient just like the configurations of neural networks. The payloads that have been completed have been compiled to be able to run within a Windows 2016, Ubuntu Linux, Linux for Tegra, and Android OS environments.

IV. PORTABILITY

Test portability also plays a major role in standardizing a test. It isn't beneficial to have an expansive lab setup that cannot be affordably and easily transported. Radiation testing often requires trips to other facilities. The hardware selected for the test benches are COTS computers that can run Windows and Linux. This permits a test bench computer to be procured near the test facility in case of a failure with freight shipping. The software is compiled and packaged with all its dependencies and licenses. Simply put, there is nothing to install. The test bench software also includes the software necessary to produce and retain run logs with unique

identifiers and template-based formatting of data across each source (i.e. V core, PSU V, memory maps and bit streams). Further, a Python-coded results parser was produced by Noah Burton (NASA GSFC, Code 562, AS&D) during a 2017 internship in NASA GSFC's Radiation Effects and Analysis Group. This post-processing application and others permit a near-immediate rapid analysis of results at the beam line. Lastly, the source code of each payload is compiled for cross platform usage. This allows avoidance of compiler optimizations - meaning it is the same payload code running on all DUTs. To increase confidence and reduce test variation, both a golden (statistical control) GPU and temperature-controlled DUT are operated in one system (either by carrier board or network). This is achieved by either software control on the hardware itself or on an arbiter computer located on a local network.

V. RESULTS AND FUTURE TESTING

Over 100 runs have been performed to date using the various test payloads and proton irradiation. Several different types of single event upsets (SEU) have occurred, such as memory corruption and single event functional interrupts (SEFI). The latter sometimes triggered system reset conditions. Figures 3 and Figure 4 show the cross section for these failures. In most instances, no noticeable electrical anomalies, visual artifacts or system latency took place during the test runs up to a preset fluence. No significant temperature rises were noted during the radiation exposure, which could be a possible indication of a single-event latch-up (SEL) event. Because the device was recoverable upon a power cycle of the computer system (CPU, mainboard and GPU), it could be used in a system that has a hardware or software watchdog routine to detect an error and reset the device.

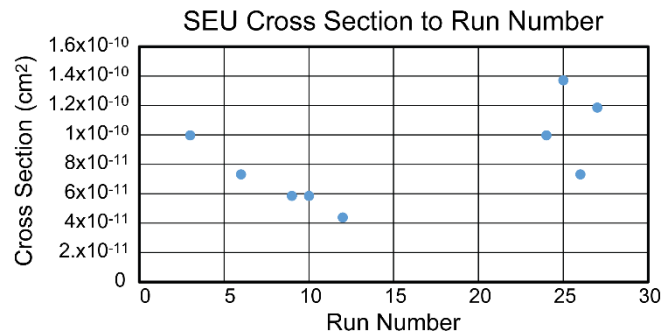


Fig 3: SEU Cross Sections (cm²) from 200MeV 1x10¹⁰ p/cm² proton irradiation testing of nVidia's GTX 1050 [1]

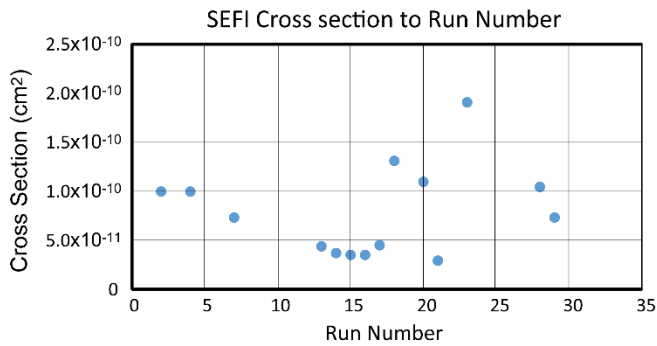


Figure 4: SEFI Cross sections (cm²) from 200MeV 1x10¹⁰ p/cm² proton Irradiation testing of nVidia GTX 1050 [1]

Laser and heavy ion testing will be performed when facility time is available. The expected payloads are GPU L1 cache, shared memory, graphics memory, math-logic, and neural network. Heavy-ion testing will determine effects of different levels of Linear Energy Transfer (LET) on the device. Because the process technology is mixed architecture and is smaller than 180nm CMOS it may be susceptible to destructive SEE in its embedded sensors. Laser testing exposes a specific area of the chip to laser pulses and the focused light (about 1 micron in diameter) moves across the surface in a controlled pattern. The system is interrogated after each laser pulse to see if there was a single event effect. For each laser pulse, we record all relevant information such as position and energy for later analysis. Proton testing evaluates SEE-induced parametric variations such as transients, SEFIs, and accessible device power-states. While some proton testing has already been conducted, more testing on other DUTs will take place. Total Ionizing Dose (TID) testing is performed in an accelerated environment and characterizes the long-term radiation effects on the device and determines whether dose-rate sensitivity exists. The cooling solution described in this methodology is radiation hardened by design so that the device can be used in open air, in a vacuum or radiation chamber.

VI. SUMMARY

The GPU test bench and its software payloads have been written with attention to open-source or public domain-sourced code snippets and hardware components such that these tests could be recreated by other engineers. This standardized approach to testing mitigates the hardware optimizations found in newer generation microprocessors whereas an apples to apples comparison would otherwise not be possible. This approach involves rapid development, quicker procurement using modular system and network components, using COTS, in house development using public domain material, and software that can be easily updated to accommodate new DUTs while maintaining the ability to test older DUTs. The goal of the test is not to confirm that a newer GPU is better than an older GPU (which optimization will most certainly do), but rather whether the fabrication technology itself is more susceptible to radiation effects. OpenCL and OpenGL code syntax allows this code to run on

most device brands and compare similar computational features across multiple device generations.

VII. ACKNOWLEDGEMENT

The Author acknowledges the sponsor of this effort: NASA Electronic Parts and Packaging Program (NEPP). The authors thank members of NASA GSFC's Radiation Effects and Analysis Group (REAG) who contributed to the creation of the test bench: Stephen R Cox, Carl Szabo, Noah Burton, Alyson D. Topper, Ray Ladbury and Martin Carts.

VIII. REFERENCES

1. Edward Wyrwas, "Proton Testing of nVidia GTX 1050 GPU," <https://nepp.nasa.gov/files/28629/NEPP-TR-2017-Wyrwas-17-039-GTX1050-2017Apr-TN45745.pdf>
2. NASA/GSFC Radiation Effects and Analysis home page, <http://radhome.gsfc.nasa.gov>
3. NASA Electronic Parts and Packaging Program home page, <http://nepp.nasa.gov>
4. Yi Yang and Shawn Newsam., "Bag-Of-Visual-Words and Spatial Extensions for Land-Use Classification," ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM GIS), 2010. The original satellite images were from the USGS National Map Urban Area Imagery collection for various urban areas in the USA. This material was based on work supported by the National Science Foundation under Grant No. 0917069
5. Joseph Redmon., "YOLO9000: Better, Faster, Stronger," arXiv preprint arXiv:1612.08242, 2016.
6. Interactions of Ions with Matter website, <http://www.srim.org/>