Proton Testing of AMD e9173 GPU

E.J. Wyrwas[1]

NASA Goddard Space Flight Center
Code 561.4, Radiation Effects and Analysis Group
8800 Greenbelt RD
Greenbelt, MD 20771

Test Date: 2 June 2019
Report Date: 22 July 2019

[1]SSAI Inc, Greenbelt, MD USA

1. Acronyms

ATE	Automated Test Equipment
BGA	Ball Grid Array
BSOD	Blue Screen of Death (Windows crash message)
Cat5e	Category 5e (enhanced) specification
COTS	Commercial Off the Shelf
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
CUFFT	CUDA Fast Fourier Transform library
DHCP	Dynamic Host Configuration Protocol
DRAM	Dynamic random-access memory
DUT	Device Under Test
EGL	Embedded-System Graphics Library
ES	Embedded Systems
GPU	Graphical Processing Unit
GUI	Graphical User Interface
HDMI	High-Definition Multimedia Interface
IPv6	Internet Protocol version
MGH	Massachusetts General Hospital
OpenGL	Open Graphics Library
OpenCL	Open Computing Language
RAM	Random Access Memory
RJ45	Registered Jack #45
SDK	Software Development Kit
SEE	Single Event Effects
SEFI	Single Event Functional Interrupt
SKU	Stock Keeping Unit
SNTP	Simple Network Time Protocol
SOC	System on Chip
SOM	System on Module
SRAM	Static Random Access Memory

2. Introduction and Summary of Test Results

Single-Event Effects (SEE) testing was conducted on the AMD e9173 Graphics Processor Unit (GPU); herein referred to as device under test (DUT). A product brief[1] is provided below.

AMD Embedded Radeon™ E9170 Series (MCM, MXM, and PCIe®)

- 14nm FinFET "Polaris" architecture
- Eight Compute Units[2]; 1.2 TFLOPS
- 2 or 4GB GDDR5 Memory; 64- or 128-bit wide
- 35-50W Total Board Power; 35W Total Graphics Power for MCM
- Graphics Clock 1124 or 1219MHz
- Memory Clock 1500MHz
- AMD Eyefinity technology for up to five display outputs[3]
- 4K HEVC/H.265 and AVC/H.264 decode and encode[4]; 4K support at 60Hz
- Microsoft DirectX® 12 capable

Figure 1: Product Brief from AMD's Power-Efficient Embedded GPUs website

Testing was conducted at Massachusetts General Hospital's (MGH) Francis H. Burr Proton Therapy Center on June 2nd, 2019 using 200-MeV protons. This testing trip was purposed to provide a baseline on radiation susceptibility data for the DUT and from application payloads compiled in Q3FY18. While not all radiation-induced errors are critical, the effects on the application need to be considered. More so, failure of the device and an inability to reset itself should be considered detrimental to the application. Radiation effects on electronic components are a significant reliability issue for systems intended for space.

The testing that has been conducted covered two types of test vectors: Matrix arithmetic and graphics output buffer. Except in the case of a single event functional interrupt (SEFI), the test vectors employed in this round of testing were created to target the shared memory, texture memory and control logic of the DUT. Because the device was recoverable upon a power cycle of the computer system (CPU, mainboard and GPU), its use in a radiative environment may be possible given a hardware or software watchdog routine to detect an error and reset the device.

3. Device Tested

The AMD e9173 GPU is a graphic coprocessor for use in a modern commercial off the shelf (COTS) computer. It was cannibalized from a Dell Wyse 5070 thin client computer. The carrier board is connected to the computer motherboard via a PCI-e x16 slot. The GPU die, itself, is the device under test (DUT) and is located underneath the unit's heat sink. Figure 1 shows pictures of the graphics card without its heatsink. Table 1 gives information on this part.

Figure 2: AMD e9173 GPU as-tested

[1] https://www.amd.com/system/files/documents/power-efficient-gpu-product-brief.pdf

Table 1: Part Identification Information

Quantity	1
Part Model	e9170
Board Model	e9173 AMD Radeon Embedded
REAG ID	19-022
Manufacturer	TSMC
Technology[2]	14nm FF
Packaging	Flip Chip, BGA

4. Test Facility

Facility: Massachusetts General Hospital's (MGH)
Francis H. Burr Proton Therapy Center

Ion species: Proton

Energy: 200 MeV incident energy at 0° angle

Flux: $2.3 \times 10^7 – 8.2 \times 10^7$ p+/sec

5. Test Setup

The DUT relies on a typical computer setup in order to be used. Here, the following platform bill of materials (BOM) was utilized (Table 2) along with Newegg part numbers. Newegg.com part numbers are referenced here as its website retains obsolete part numbers and single unit pricing. The operating system was Windows 10 x64.

A custom tooled cooling solution was created to permit beam access to the DUT die from the obverse side while absorbing the heat through the reverse side 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 allows the device to operate under load while maintaining an ambient temperature appropriate for the test (i.e. 20°C). While not defined in Table 2, the cooling system uses 400W of thermoelectric coolers, a large aluminum fan sink and a secondary ambient blower. Temperature monitoring was performed in software and by taking point measurements with a thermocouple at the die fillet.

Table 2: Computer Platform - Bill of Materials

Newegg.com Part #	Description
N82E16813119107	ASUS TUF X470-Plus Gaming AM4 AMD X470 SATA 6Gb/s USB 3.1 HDMI ATX AMD Motherboard
N82E16819113446	AMD RYZEN 3 1200 4-Core 3.1 GHz (3.4 GHz Turbo) Socket AM4 65W YD1200BBAEBOX Desktop Processor
N82E16820236072	CORSAIR Vengeance LPX 64GB (4 x 16GB) 288-Pin DDR4 SDRAM DDR4 3000 (PC4 24000) Desktop Memory Model CMK64GX4M4C3000C15
N82E16817139084	CORSAIR HXi Series HX750i 750W 80 PLUS PLATINUM Haswell Ready Full Modular ATX12V & EPS12V SLI and Crossfire Ready Power Supply with C-Link Monitoring and Control
9SIA12K77Z5902	SAMSUNG 970 PRO M.2 2280 1TB PCIe Gen3. X4, NVMe 1.3 64L V-NAND 2-bit MLC Internal Solid State Drive (SSD) MZ-V7P1T0BW

[2] https://www.techpowerup.com/gpu-specs/radeon-e9173-pcie.c3031

A. Arbiter Setup

An external arbitration computer (laptop) operating over a closed network was used to interrogate the device, execute remote commands and monitor the DUT health. USB over Ethernet was used to access the test computer's mouse and keyboard. HDMI over Ethernet was used to view a monitor located in the operator hallway.

A video capture device was used to passively record a video stream from the DUT over USB 3.0 as received over the HDMI over Ethernet connection. This redundant-monitoring approach permitted direct control of the DUT during the test and minimized the risk of false errors recorded due to upsets in the primary networking connection itself. Payload software was placed on an FTP storage site local to the closed network's router for easy update, download and results extraction between DUT and arbitrator.

B. Test Vector Software

The following test payloads were performed using payload code developed at GSFC.

- Matrix arithmetic
- Graphics output buffer

C. Hardware

The DUT is the graphics chip located on the PCI-e carrier board. Due to the orientation of the GPU on the system board, a riser cable (Digikey part number 3M12026-ND) was used to place the GPU above the test computer. This also permitted the computer to be surrounded by lead and Lucite bricks to prevent SEFIs on the motherboard. The beam was aligned to the obverse side of the GPU card. No fan-sink is located on this side of the card because sufficient cooling is produced on the reverse side of the card from the cooling system. There was sufficient clearance around the GPU chip and no components were present on the secondary side of the system board within the z-axis of the chip. This was advantageous as it allowed some radiation shielding to other system components such as the power management and flash memory components. Lucite bricks were used to shield the power supply of the DUT from scattered neutrons which are a result of proton collisions within materials in the beam's path. A photograph of the board under test is shown in figure 3.

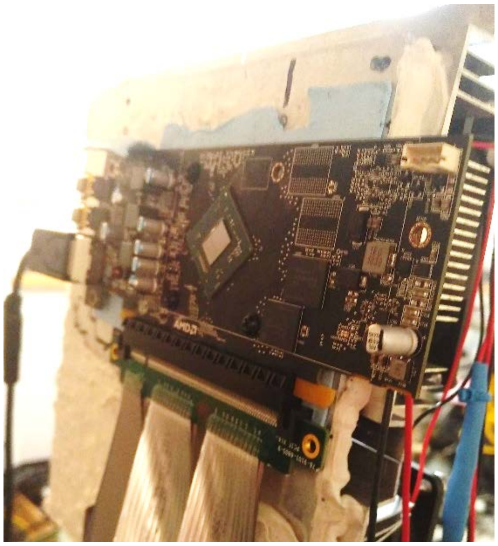

Figure 3: AMD e9173 on Cold Block aligned in beam at MGH

6. Failure Modes

Four types of failure events can be recorded during the test campaign. These events are indicative of the sensitivity the hardware to the radiation and the fault resilience of the operating system to failure of instructions, memory fetches and architecture microcode running in the background. While further analysis is required to identify the failing operating system or hardware component when the event occurred, a brief explanation of these event types is provided.

The first event type was that of a processor machine check (MC) error which was logged by the operating system. Each recorded machine check error was logged as a 64-bit value and was decoded using the vendor specific hardware manuals. The decoded value was able to indicate whether the error was "corrected" or "uncorrected" and the functional block within the microprocessor from which the error originated. A mixture of corrected and uncorrected machine checks was observed. All uncorrected machine checks led directly to a system crash. Some of the corrected machine checks produced a system crash as well, but the majority of the corrected machine checks were recorded without a system crash and with no noticeable change in operation from the OS.

The corrected machine checks logged during the tests decode to either an L1 or L2 cache error according to the documentation. Further decoding of the machine checks indicate a specific cache operation associated with each event. It is unclear however if the error resulted from a bit flip in a SRAM cell within the cache or an upset in other circuitry involved with the operation of the cache. Note that cache level naming convention is adopted from the Intel SDM which lists the levels as L0, L1, and L2 with L0 being the lowest level cache.

The second event type was a system crash where the OS would become either unresponsive, shut itself down, or reboot itself. After the system crash was observed and the system was restarted, the operating system and its idle behavior was assessed to determine if latent damage had occurred. This type of event is categorized as a Single Event Functional Interrupt (SEFI).

The third event type is hardware failure. Multiple sets of the test platform and spare hardware are present at the test facility. This enables real time debug and diagnosis when any component within the hardware bill of materials becomes suspect or exhibits "hard failure" during irradiation. In general, the computer fails to boot.

The fourth event type are pixel artifacts. When the display output to the monitor is not as expected, then the behavior is generically categorized as pixel artifacts. An example of this behavior is shown below. The display should show the Windows' desktop. Instead, it shows a rainbow colored snow pattern.

Figure 4: Examples of Pixel Artifacts

7. Test Procedure and Results

Five (5) runs were performed. Each run resulted in pixel artifacts and a Single Event Functional Interrupt (SEFI) of the test system. The electrical measurements from the fifth run reflected a latch-up type of event. Further post-processing of data is required to determine root cause of failure. It is also worth noting that upon power cycling the test system, the device behaved normally. Further, no drift in temperature was noted other than a negligible increase due to computational loading. Table 3 and Table 4 show the results of this testing campaign.

Table 3: Testing Results

Time of run (s) until SEFI	Average Flux (p+/sec)	Effective Fluence (p+)	Dose (rad(Si))	SEU Cross section (cm^2)
3.6	8.17E+07	2.94E+08	17.06	3.40E-09
30.6	5.16E+07	1.58E+09	91.61	6.33E-10
28.2	2.30E+07	6.49E+08	37.65	1.54E-09
129	2.40E+07	3.10E+09	179.52	3.23E-10
25.8	2.39E+07	6.16E+08	35.72	1.62E-09

Table 4: Summary of Results

	Average Flux (p+/sec)	Effective Fluence (p+)	SEU Cross section (cm^2)
min	2.30E+07	2.94E+08	3.23E-10
max	8.17E+07	3.10E+09	3.40E-09
average	4.09E+07	1.25E+09	1.50E-09
standard deviation	2.59E+07	1.14E+09	1.20E-09

8. Discussion

During the irradiation of this device, elements of the GPU appear to have operated improperly (incorrect clock frequency reading, degraded state) as compared to similar device architectures. The device was unable to complete a full session of the software payloads. Further characterization is required to identify the lower bounds of flux necessary to operate the device for an elongated period of time. Being able to achieve this will permit further baseline and application-level testing.

The methodology used for testing was a “best effort” method to replace traditional custom bias boards and expensive Automated Test Equipment (ATE), albeit a method that has been refined over a few investigations. The SoC manufacturer is able to afford both the ATE equipment and the manpower to develop the test vectors due to commercial sales volumes (i.e., free market economics).

We have performed a series of proton irradiations on commercial off the shelf (COTS) microprocessors, utilizing system-level tests that are conducted with commercial and free software tools. This work is a continuation of previous efforts supported by the NEPP Program and builds upon successful collaborations with NSWC Crane, Jet Propulsion Lab (JPL) and other entities. The authors look forward to future tests on these parts.