



UNCLASSIFIED



AMRDEC HWIL

PC Scene Generation

James A. Buford, Jr.¹, David Cosby²,
*U.S. Army Aviation and Missile Research, Development and Engineering Center, Redstone Arsenal,
Alabama, 35898*

and

Dennis H. Bunfield³, Anthony J. Mayhall⁴, Darian E. Trimble⁵
Davidson Technologies, Inc., Huntsville, AL 35806

AMRDEC has successfully tested hardware and software for Real-Time Scene Generation for IR and SAL Sensors on COTS PC based hardware and video cards. AMRDEC personnel worked with nVidia and Concurrent Computer Corporation to develop a Scene Generation system capable of frame rates of at least 120Hz while frame locked to an external source (such as a missile seeker) with no dropped frames. Latency measurements and image validation were performed using COTS and in-house developed hardware and software. Software for the Scene Generation system was developed using OpenSceneGraph.

INTRODUCTION

RDECOM AMRDEC has many years of experience developing and applying hardware and software solutions to provide dynamic electro-optic and infra-red (EO/IR) scene generation for testing weapon system sensor components. After successful custom development and utilization of the SensorVision I, II, & III systems, AMRDEC began to pursue more cost-effective Commercial Off The Shelf (COTS) image generation solution. In the early 1990's AMRDEC transitioned to the first implementation of COTS infrared scene generation technology using Silicon Graphics Incorporated (SGI) Onyx systems running COTS real-time image generation software. This approach, still in use, has proved capable, with high but known costs. As Personal Computer (PC) architectures and COTS Graphics Processing Units (GPU) began to demonstrate greater fidelity, higher performance and greater innovation, AMRDEC undertook studies to identify where and how they could be employed. In that time, the clock frequency of "PC" processors has more than doubled to the multi-Gigahertz rate, and data paths have expanded from 32 to 64 bits. GPUs have implemented additional programmability, such as vertex and pixel shaders through the OpenGL Shading Language (GLSL), as well as 32 bit floating point per color precision frame buffers.

What this paper will outline is what steps were taken to show how the Concurrent system (and a real-time PC based system in general) is capable of performing real-time Hardware-In-The-Loop (HWIL) Scene Generation required by AMRDEC labs. COTS PCs with nVidia G-sync cards were used to develop 16 bit output capability and discover the PCs potential as a viable real-time scene generation platform.

¹ Engineer, System Simulation and Development Directorate.

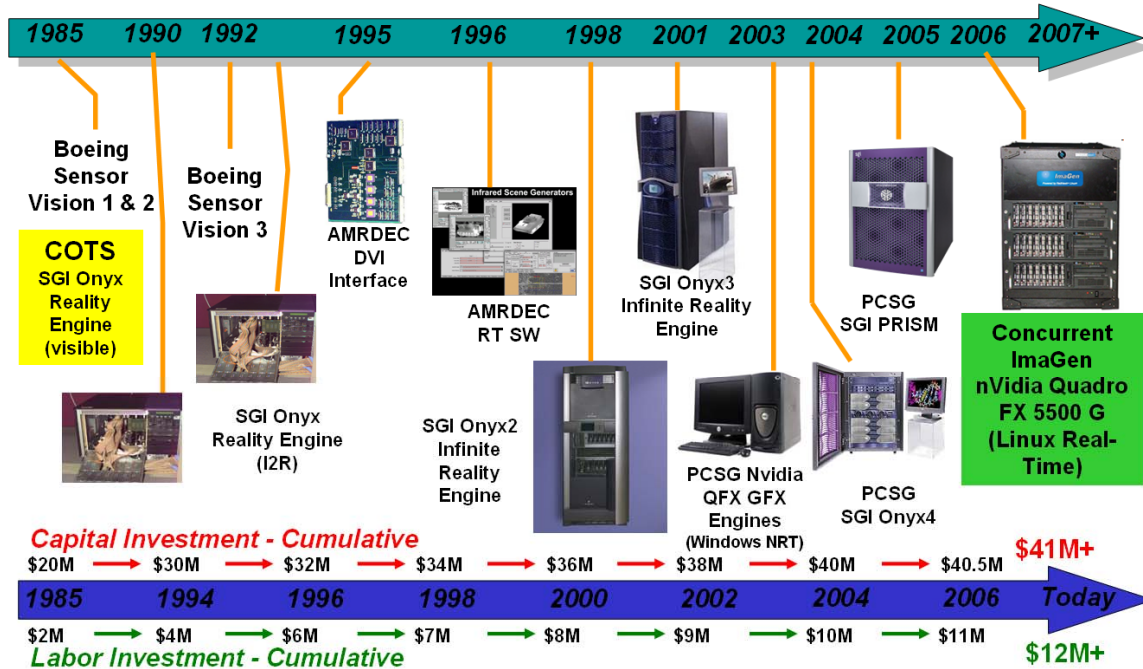
² Engineer, System Simulation and Development Directorate.

³ Engineer, Davidson Technologies, Inc.

⁴ Engineer, Davidson Technologies, Inc.

⁵ Engineer, Davidson Technologies, Inc.

UNCLASSIFIED



*Does not include SAL, LADAR or Visible projector efforts

Figure 1: AMRDEC Scene Generation Timeline

1. SCENE GENERATION REQUIREMENTS

The requirements for real-time scene generation at AMRDEC are as follows: running the windowing system in real-time along with the graphics application, the ability to lock a process in memory and to a processor in a real-time fashion, the ability to genlock or framelock the video format to an external source, the ability to generate the required scenes at 120Hz or faster, the ability to output 16 bit values, the ability to use custom video formats to drive projector electronics, and the ability to run open source software.

2. PCSG HARDWARE/SOFTWARE EVALUATED

Several systems were evaluated during this process. The first evaluation started several years ago with the SGI ONYX 4 systems with ATI graphics. This system was the first COTS PC platform with specifications to meet all the AMRDEC requirements. Upon evaluation, the ONYX 4 met most of the requirements. However, it was immediately apparent this platform would not meet AMRDEC 16 bit fidelity requirements, because there were significant problems with the ATI graphics hardware outputting noise in the final rendering. SGI indicated new hardware would be available to fix the problem.

Because maintaining frame rates is a requirement for a real-time system, only two choices remained: Silicon Graphics PRISM with ATI graphics and Concurrent Imagen with nVidia G-sync cards. Although SGI had dropped the price of this graphics systems significantly compared to previous product lines, the PRISM failed to perform as expected. SGI had not resolved the issue noted in rendering noise on the ONYX 4. SGI was contacted about the problem, but indicated that they were rewriting the entire ATI driver. Therefore, a near term solution was not forthcoming. An additional problem appeared with unexplained interrupts in the X Window System, thus causing dropped frames. Further, the ATI cards only supported NTSC broadcast standard genlock. It did not support framelock needed to meet the synchronization requirement.

UNCLASSIFIED

The overall goal was to fulfill all of the benchmarks outlined in Section 1 and not have any dropped frames during a simulated flight. The Concurrent Imagen system with the nVidia Quadro FX G line of cards was able to meet all of these requirements. These systems provided:

- real-time use of the X window system and video card drivers
- real-time operation of OpenGL were demonstrated with no anomalies
- high performance
- framelock/genlock
- a strong interest from Concurrent and nVidia to develop a working relationship with AMRDEC.

3. HARDWARE VERIFICATION METHODOLOGY

The first test was to ensure the Concurrent system could run a graphics application in real-time without interfering with the rest of the operating system's functionality. The Concurrent system allows for applications to be easily run in real-time without having to make modifications to the kernel. Using the Concurrent "run" command and adding "memory locking" into the code, it was shown a graphics application could be run in real-time on a selected CPU without affecting the rest of the system.

The next test was to generate custom video formats. Because PC video cards use the VESA standard formats to drive monitors, custom formats had to be generated while staying within the VESA constraints. To drive a set of custom projector electronics, it is routine to frequently modify the video formats to have faster pixel clocks or faster horizontal rates than default formats use. These specific formats were successfully generated at various rates and synchronized with monitors and projectors. 120Hz formats were used for a significant amount of the testing because of the COTS DVI frame grabber hardware specifications. Higher frame rates (i.e. > 120Hz) were successfully tested with the digital I/O board used in the latency testing mentioned below.

During testing of these modified video formats, and leading into genlock testing, it was discovered that the nVidia boards would lock to an external pulse and synchronize the video formats to that pulse instead of using an actual genlock signal. This feature was a substantial advantage, because it avoided the need to build a custom board for testing the systems synchronization. It also allowed for an existing in-house design to be used for future genlock boards. The next step was to test the genlock capability of the boards. For example, it was found that if a 120 Hz video format was generated, the system would successfully synchronize to an external 120Hz input pulse.

A simple OpenGL sample program was used to test the real-time and genlock capability of the graphics system. It was found that there were real-time limitations with several nVidia library calls, the driver, and the hardware architecture that would interfere with the "no dropped frames allowed" requirements. With tremendous help and knowledge from both Concurrent and nVidia, it was determined which library calls not to make. They were also instrumental in working around other issues in order to run without dropping frames, maintaining the proper frame rates, and remaining synchronized to an external source. The AMRDEC OpenSceneGraph based scene generation software required several modifications using the above knowledge. These scene generator modifications were used for the remainder of the tests and evaluations of the system.

To ensure the PC scene generation hardware could perform HWIL scene generation, the system latency needed to be determined. To test the system's latency, a digital I/O board, an in-house designed and built DVI frame grabber board, and a logic analyzer were used. A frame counter was encoded into the first red pixel of the image while the scene was rendered, and at the same time, that value was written to the digital I/O board. The DVI frame grabber board was programmed to put the first red pixel value on the first 8 red output pins. The 8 output pins from the digital I/O board and the 8 output pins from the frame grabber were attached to the logic analyzer. Using the "state mode" on the analyzer, both sets of numbers were shown. For instance, the digital I/O board would have a value of 16, and the red pixel would have a value of 14. Two frames had elapsed since the encoded frame count was placed into the red pixel value. This obviously shows the two frames of latency. It was later found that the two frames of latency were

UNCLASSIFIED

caused by the video card vendors' heavy use of pipelining and frame buffering to increase the performance of the cards. In nVidia's case, a command was used to change a certain parameter to see what effect it would have on performance. When running synchronized to the VBLANK signal (as stated in the requirements), the performance of the graphics application suffers greatly when instructing the card to buffer less than 2 frames.

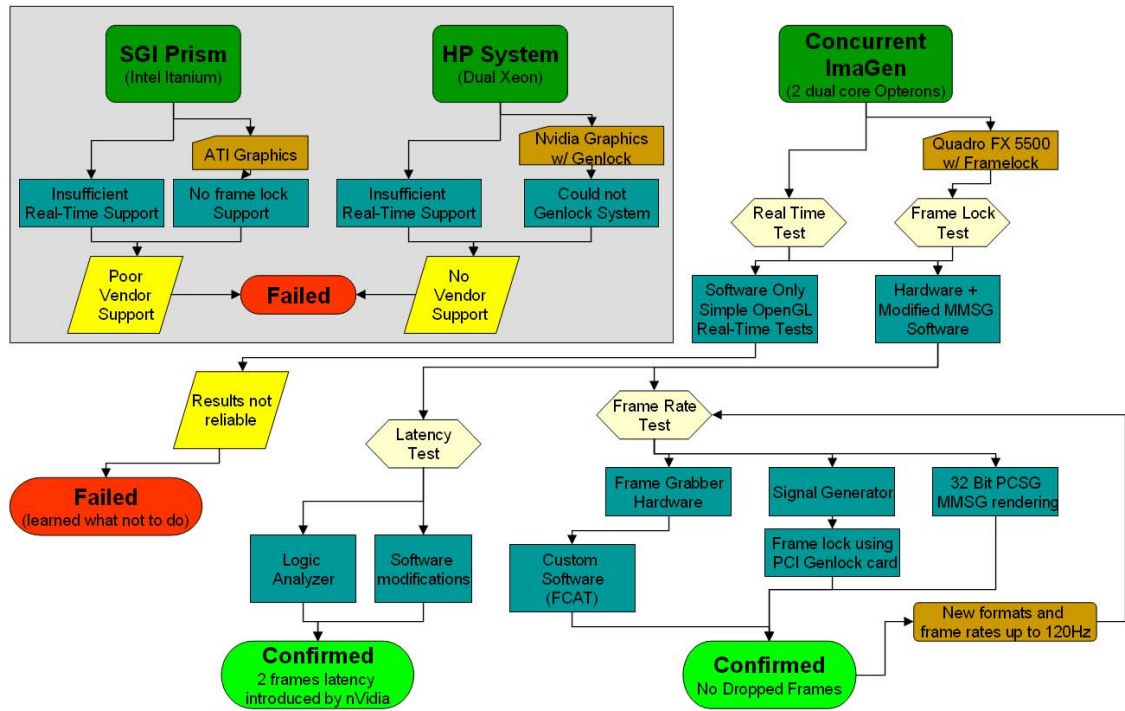


Figure 2: Platform Verification Methodology

The next test was to ensure that no frames were dropped. Because for normal 16 bit rendering, the red and green channels are used to produce the scene, the blue channel was used to encode the frame count and other diagnostic information to help debug and time the scene generator output. For diagnostic analysis, a COTS DVI frame grabber was purchased and code was written for it. This allowed the blue frame information to be captured and reported if any dropped frames occurred. This successful test was able to show that no dropped frames at 120Hz occurred using an IR tactical 16 bit scene output consisting of 400,000 polygons and 20MBs of textures while synchronizing to an external pulse.

A custom genlock board had to be built to synchronize the system to an actual seeker and still run at 120Hz. Therefore, a custom PCI board to receive an external pulse from a seeker and send a signal to the video card's genlock port at a specified rate (usually a multiple of the input rate) was designed and developed. The card was built and tested in-house, and it was successfully able to demonstrate its capability using pulse generators, oscilloscopes, logic analyzers, and the nvidia-settings panel.

Lastly, putting it all together, the scene generation software rendered a 16 bit output image with the PC attached to a custom set of projector electronics. The output was captured by the electronics, and the data was used for verification and validation of the scene. The system has since been tested while synchronized to two different missile systems currently in use in the AMRDEC HWIL labs using the custom built genlock board.

4. LAB CAPABILITIES

UNCLASSIFIED

The goal of the PCSG system is to maximize utilization of the available hardware in the entire simulation system. In most HWIL cases, the scene projector itself defines the final resolution that will be projected to a sensor. AMRDEC's tactical facilities were equipped with a Large Format Resistor Array (LFRA) which can emit 1024 by 1024 resistive elements. The input into the LFRA system accepts either one 1048 by 1048, or two 524x1048 frames from the GPU(s) across DVI-D at 120Hz, and recombines the red and green channels to get the GPU calculated 16 bits of precision. This sets the baseline capability for both the tactical and strategic PCSG systems.

The tactical scene generation capabilities meet the needs of mostly subsonic systems. These sensor systems have relatively more time to scan areas of interest (AOI), and therefore demand as highly as detailed scenes as possible from the PCSG and the scene projection subsystems to match a real world scene. Therefore, utilizing two GPUs for the IR PCSG was critical to meet the geometric complexity demands of the labs. This also allows the labs to double the performance of the system each time a newer compatible graphics card is released.

The nVidia GPU greatly increases the capabilities and fidelity of the labs by utilizing the GPU shader language. Programming the nVidia GPU to calculate phenomenological IR effects using 32 bit precision surpasses legacy hardware 12 and 16 bit implementations. The shader language gives the software greater access to the full potential of the hardware by allowing the programmer to choose and optimize most steps of the IR rendering process.

5. CONCLUSION

Last generation hardware was far behind the demands of current lab requirements. Legacy hardware could barely achieve 30Hz, because of higher fidelity requirements. For example, tactical scenarios could not support the highest fidelity of targets, terrain, and atmospheric needed to meet the demands of the system. The latest GPUs and the PCSG can achieve 300Hz with this higher fidelity terrain (400K polygons), and still perform other GPU shader calculations and dynamic effects.

Total cost of hardware limited lab capability for last generation systems because of prohibitive maintenance costs and limited upgrade capability. Current PCSG hardware allows for minimal maintenance expenses and is easily upgradeable to newer versions as hardware becomes available and compatible. PC Scene Generation hardware costs are ~1/7 the cost of legacy hardware, and maintenance costs are ~1/5 of legacy hardware.

Latency mitigation approaches will have to be evaluated and developed for each lab, because 2 frames of latency are generated by the long GPU pipelines needed to increase performance. The scene generation real-time requirements, however, were successfully met. Using the Concurrent Real-Time O/S, the scene generation application was able to be restricted to a single processor, and interrupts were also restricted from interrupting the scene generation process. It was also determined that there were also no dropped frames through analysis of the frame counter encoded in blue channel. Genlock / Framelock was also successfully tested up to 200Hz using the designed and built in-house Framelock PCI board.

PC COTS development presents promising opportunities to lower HWIL operational costs and expand capability. This new PCSG capability developed by AMRDEC far exceeds current lab scene generation capabilities. This next generation capability allows for higher fidelity models and new rendering algorithms and techniques to be developed in the future.

UNCLASSIFIED