

QUANTITATIVE PERFORMANCE-DRIVEN PC-BASED IMAGE GENERATOR EVALUATION FOR VISUAL INTEGRATED DISPLAY SYSTEMS

Budimir Zvolanek, Training and Support Systems - The Boeing Company, St. Louis
William Paone, PureLogix Division - Westar Corporation, St. Louis
Ed Elking, Training and Support Systems - The Boeing Company, St. Louis
Tim Dwyer, Training Systems Product Group - USAF ASC/YW, Dayton

Abstract

Affordability and performance of training systems devices have become key factors in improving availability of training to a broader military community. An example of this is the Target Projection System (TPS) embedded within Boeing's Visual Integrated Display System (VIDS) – a product delivered to a number of military training programs such as the T-38 and the USAF Distributed Mission Training (DMT) F-15C and DMT F-16 programs. Driven by a dedicated image generator (IG), the TPS simultaneously projects multiple high-resolution images of aircraft onto VIDS screens. Until recently, only a desk-side workstation or a full-featured, high-performance IG has had the polygon and pixel-fill performance necessary to generate the TPS aircraft imagery. However, rapid advances in PC-based three-dimensional (3-D) graphics technology have finally offered such performance at a much lower cost. This paper describes Boeing's effort in transitioning PC-based IG technology into the VIDS product by quantitative measurements of PCIG performance using TPS-specific benchmarks. IG performance requirements and their embodiment in benchmark databases and test software are described. Available PC-based IG descriptions are provided followed by a comparison of the benchmark test results, as well as a discussion of issues with real-time image generation hardware and software integration. Finally, a recommendation of the TPS PC-based IG is presented based on the observed performance, as well as IG features and other 'non-performance' factors. Suggested PCIG applications conclude the paper.

Biographical Sketches:

Budimir Zvolanek is a Boeing Associate Technical Fellow responsible for visual, NVG, IR, and radar image and database generation, image auto tracking, and image exploitation development at Boeing Training and Support (TSS). At Boeing Phantom Works, he developed innovative automatic target recognition and image processing algorithms for missile guidance. He also developed image processing hardware and software at Washington University Biomedical Computer Laboratory, where he received his MSEE and BSEE degrees. Budimir has published extensively, receiving two I/ITSEC Best Paper Award Nominations.

William Paone is Senior Development Engineer at Westar Corporation, PureLogix Division, in St. Louis, Missouri, developing image generation software and databases for the Boeing VIDS* Product Center. Bill has developed real-time imaging hardware and software systems for cockpit and out-the-window display generation. He also developed telecommunication system hardware and software for American and European markets, managed system production and test, and led hybrid microcircuit production engineering. Bill received his BSEE degree from the University of Missouri at Rolla.

Ed Elking is a Boeing Associate Technical Fellow responsible for visual image generator, database, and display software development at Boeing Training & Support Systems (TSS). He has led numerous software development projects under the A-6/F-14, Finnish F/A-18 WTSAT, DMT F-15C, DMT F-16 MTC, T-38, and other programs. His contributions range from laser target projection to high-resolution target projection and background visuals for the Boeing VIDS system as well as traditional flight simulation controls and coordinate transform theory applications. Ed received his MSEE and BSEE degrees from St. Louis University.

Tim Dwyer is a Systems Engineer and a Visual / Sensor Systems and Flight Simulation Technical Advisor at the USAF ASC/YW Training Systems Product Group at Wright-Patterson AFB. Some of his most recent assignments have been as a technology evaluator and customer for the DMT F-15C and F-16 Contracted Training Systems and Services programs. Tim received his BSEE degree from the University of Dayton.

Quantitative Performance-Driven PC-Based Image Generator Evaluation for Visual Integrated Display Systems

Budimir Zvolanek, Training and Support Systems - The Boeing Company, St. Louis
William Paone, PureLogix Division - Westar Corporation, St. Louis
Ed Elking, Training and Support Systems - The Boeing Company, St. Louis
Tim Dwyer, Training Systems Product Group - USAF ASC/YW, Dayton

INTRODUCTION

Product affordability has driven the training systems community to continually reduce component and maintenance costs while improving performance of image generation equipment. Fortunately, due to recent radical advances in three-dimensional (3-D) graphics technology, migration of high-end image generation to low-cost personal computer (PC) based image generators (IGs) (PCIGs) has started to show promise. We have actually realized this migration potential in our application, replacing a “high-end” desk-side workstation by a much lower cost desktop PCIG. However, this was accomplished only after a thorough performance evaluation of candidate PCIG platforms, procurement based on several candidate PCIG proposals, and resolution of a number of hardware and software integration problems arising from ‘peculiarities’ of the PCIG.

In particular, a systematic, quantitative as well as qualitative, comparison was implemented by carefully selecting a set of application-meaningful performance parameters, developing a benchmark database, and implementing behavior-emulating benchmark software. Then after selecting and acquiring the PCIG, major technical challenges still had to be overcome to port our application from the high-end IG to the PC-based IG. Nevertheless, the challenge has been met successfully.

The Background section of this paper provides a context of the application within the Visual Integrated Display System (VIDS) – the Target Projection System (TPS) - for Distributed Mission Training (DMT) and other programs. Also provided is an overview of the PCIG technology in terms of price / performance categories and in comparison to the high-end IGs. Next, the PCIG evaluation process is developed, including the testing parameters and

benchmarks. Results of quantitative PCIG testing are discussed, including qualitative PCIG characterization. Next, issues with real-time IG software and hardware migration are discussed. Finally, recommendations are made on critical PCIG evaluation steps, PCIG selection, and impact of PCIG on affordable image generation.

BACKGROUND

Before describing the evaluation process, the need for a low-cost IG within the VIDS is described, followed by an overview of the state-of-art in PCIG technology and the requirements for a PCIG-based TPS design.

PCIG Application for the Visual Integrated Display System

The F-15C and F-16 Distributed Mission Training (DMT) and T-38 Programs have been underway to augment existing US Air Force (USAF) training capability in both local-site and long haul networked DMT modes. All of these applications are based on simulator cockpits placed within a 216 or 360-degree Field Of View (FOV) Visual Integrated Display System (VIDS*) developed by Boeing.

Also part of each DMT site are the desktop PC-based Threat Station (TS), Instructor-Operator Station (IOS), and Brief-Debrief Station (BDS) simulators networked with the cockpits. The latter are, in effect, “low-end” PCIGs based on commercially available 3-D graphics accelerator cards. However, of particular interest here is the “high-end” IG used to drive the VIDS TPS for real-time visualization of multiple aircraft models at many orientations and levels of detail (LODs).

What Is a PCIG?

A PCIG is a 3-D hardware and software device utilizing chipsets, board-sets, busses, and graphics (scene) Application Programmer's

Interfaces (APIs) that borrow heavily from the commercial Intel-based PC architecture. Representative examples of hardware and software for PCIGs include:

- Intel/AMD CPU processors;
- SDRAM, RAMBUS, DDRRAM, and SGRAM memory;
- ASICS and field-programmable logic;
- PCI, IDE, AGP, SLI, and AFR busses;
- OpenGL, DirectX, and GLide APIs.
- 3-D sound digital signal processors;

Companies currently marketing PCIG component and system products include:

- Graphics ASIC Chips - ATI, 3Dfx, NVidia, Videologic, Evans & Sutherland, Intergraph, and Silicon Graphics;
- Graphics Boards – Intense3D, Evans & Sutherland, Quantum 3D, 3dfx, ATI, Elsa;
- Graphics Systems - Evans & Sutherland, Quantum3D, Primary Image, Silicon Graphics, and Intergraph.

Most, if not all, PCIGs utilized for virtual vehicle simulation purposes are specifically designed with field scalability in mind, i.e. field expansion in terms of pixels, polygons, database objects, mission size, number of video channels, number of eye-points, update rate, display support, etc. In general, the more field-scalable a PCIG is, the more it will cost. The same is true, of course, for high-end IGs, too.

In addition to scalability, PCIGs share other key characteristics associated with high-end IGs, including multi-chip parallel graphics processing, high-quality spatial and temporal anti-aliasing, IG process determinism, large fields-of-view, and rapid database generation support. It is these key characteristics that distinguish PCIGs from graphics workstations, PC game machines, and game consoles.

Cost per pixel is different between high-end IGs, PCIGs, graphics workstations, PC game machines, and game consoles. In rough terms, there is at least a 2x reduction in cost per pixel when transitioning between technologies taken in sequence from most sophisticated (high-end) to least sophisticated (game console). This cost

paradigm will probably hold true for a time to come, at least until high-end IGs meet all customer expectations for simulating visual reality. When high-end IG vendors cease to innovate, lower price-point device vendors will penetrate the IG market, and will drive the high-end vendors to either exit the market, switch to marketing systems based on PCIG and open graphics architectures, or license their proprietary graphics algorithms for re-use by others. The evolution of CompuScene / Real3D / SEView IG product line at Lockheed-Martin has been a recent, salient example of this ongoing phenomenon.

PCIG Requirements for the TPS design

The general requirements for the TPS PCIG and scene API were that the PCIG system featured “open” architecture, the graphics hardware had “performance scalability”, and the PCIG run-time IG software could be ported easily among different PCIG platforms. The system had to support multiprocessing capability to separate and execute the Application/Cull (App/Cull) and Draw tasks on two PCIG CPUs. The API also had to provide hooks to manipulate multiple views into a scene (e.g., VegaNT pfChannel and vgChannel) at 60 Hertz. This requirement alone immediately eliminated several systems from consideration.

Also, other goals established were to:

- Provide increased rendering performance of 12 target models.
- Provide a qualified second source for the PCIG system. (E.g., two vendors PCIGs have been successfully integrated into the TPS simulation.)
- Implement a PCIG drop-in replacement for the original IG (i.e., this replacement will not require VIDS host software changes.)

The TPS scene generation requirements are unique, as the TPS IG is not used in a traditional way where visual channels and FOVs are static during the simulation. Instead, the TPS IG delivers video to the projection electronics and optics where high-resolution targets are merged with the background terrain imagery. Thus, the TPS IG must be genlocked (horizontally phased locked) to the host system and background IG. Furthermore, the TPS scene is composed of multiple channels (views into a scene) running at 60 Hertz. The system must also provide off-axis perspective correction, viewport (FOV) rotation, and target-to-target occulting. Therefore, the IG must not

only provide imagery at 60 Hertz (controlling eye-point and target positions) but also change channel position, channel sizing, frustum (field of view) definitions, and viewport skew - also at 60 hertz. The background image generator generates the targets at close range and the system will transition the targets to the TPS system at a fixed range.

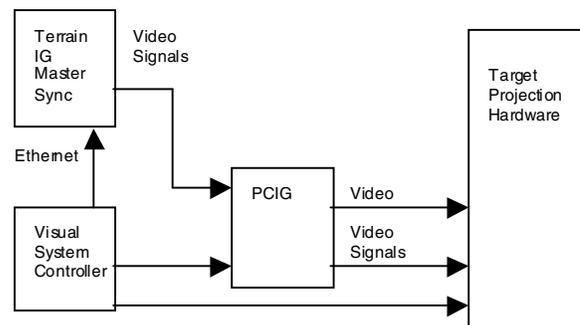
Figure 1 shows a block diagram of the TPS system. The TPS IG receives control packets over the network at 60 Hz. The VIDS Visual System Controller (VSC) delivers messages to the PCIG at 60 Hz. This interface is accomplished through a proprietary message-passing interface. The interface consists of channel and target definitions. The VSC controls the activation, screen position, FOV, and viewport rotation for all active channels each frame. Also target data for target activation, 6-DOF position, external configuration, and part articulation are also sent. The VSC also sends similar viewport definitions to the Target Projection Hardware synchronized with the PCIG at 60 Hz. This hardware is responsible for correctly positioning and sizing the mini-raster (from a high-resolution projector) on the appropriate screen.

Each channel can include multiple targets but has one target of interest (TOI). The TOI will be rendered normally and other targets will be rendered as a black mask to accomplish proper occulting. This requires the implementation of callbacks in the system. A Pre-Cull callback is installed for each object enabled into the scene. This callback is used to decide which models in a particular channel are to be rendered normally or as a mask. This allowed two types of target-to-target occulting using the z-buffering capability of the image generator. Occulting between targets in the TPS system is accomplished as well as background target to TPS target occulting. Therefore, when targets are being displayed in the background image generator, they are also still active in the PCIG TPS scene but will only be rendered as masks.

The system must run in multiprocessing (multithreaded) mode using two processors to accomplish the task of keeping the application and cull (in one processor) and the draw process (in second processor) within the 60-Hertz boundary. Keeping the Application and Cull in the same thread also allows proper communication between the application and pre-Cull callbacks. Test patterns are also required in this system. These are

accomplished with a Post-Draw callback for each channel. This callback, written in OpenGL, allows crosshairs, grids, focus patterns, etc. to be laid over the scene.

Hardware performance and interface requirements for the TPS IG are also very specific. For geometry processing the system has to provide nine models with 574 triangles, which equates to 311,204 simulation triangles per second. The screen image resolution of the TPS system is 1,280 pixels x 1,024 lines. The system also has to Genlock to an external sync signal source. The system must receive separate single-ended Vertical and Horizontal syncs and also output the same sync signals with appropriate voltage polarity. The system has to provide full-screen sub-pixel hardware anti-aliasing and per-pixel texturing, lighting, and z buffering. And the texture memory has to be at least 8 Mbytes.



The PCIG GenLocks to the input Video timing and receives commands from the Network

Figure 1. Target Projector System Block diagram

EVALUATION PROCESS DEVELOPMENT

In order to select a PCIG system, the runtime scene API was first selected. This way, we could design a benchmark program and database that could be easily installed on any open-architecture PCIG system and obtain performance results quickly to decide on potential PCIG hardware. MultiGen-Paradigm, Inc. Vega NT was thus selected that supported Multi-threaded applications and is based on Performer that was the original API for the TPS.

To test the API-plus-hardware IG system performance, a benchmark exemplifying fully integrated TPS application was conceived. At first there was no concern with the pixel-fill rate but during integration performance decrease occurred with targets that were rotated vertically. This occurred when targets were at closer range of 400 to 500 feet. The benchmark, however, was designed to calculate only shaded simulation triangles per second with full scene anti-aliasing and with the system stressed to the 60-Hertz limit in App/Cull and draw time. Also, in the original TPS IG the two processors (App/Cull and Draw threads) run out of time coincidentally. Using the PCIG architecture, we had some room in the processor performance for the App/Cull thread to improve the performance of the system.

The TPS benchmark was created for the candidate PCIG machines using Vega NT. This benchmark placed target channels into the display window at the user's discretion with the frustum and target positions and orientations all changing and updated at 60 Hertz while displaying appropriate real-time performance statistics.

This benchmark was used for the PCIG evaluation and selection process but another was created later which took a snapshot of (worst case) data of an actual simulation through a series of frames to better simulate the target image generation environment. The latter benchmark software had to be designed anyway to help solve some problems during TPS hardware/software integration.

In addition, master OpenFlight database files were created for PCIG manufacturers who could not support running the selected run-time IG software API. These database files thus allowed the manufacturers to measure image generation performance data by activating 1 to

12 models while using the scene viewers of their choice.

QUANTITATIVE AND QUALITATIVE EVALUATION RESULTS

Our evaluation included PCIGs from several suppliers including Quantum3D, Intergraph, Evans and Sutherland, Primary Image, and SGI.

In terms of raw hardware performance, the polygon (geometry) throughput of the system selected for TPS was over 500,000 triangles/sec while most of the other IGs measured at above 400,000 triangles/sec (Figure 2).

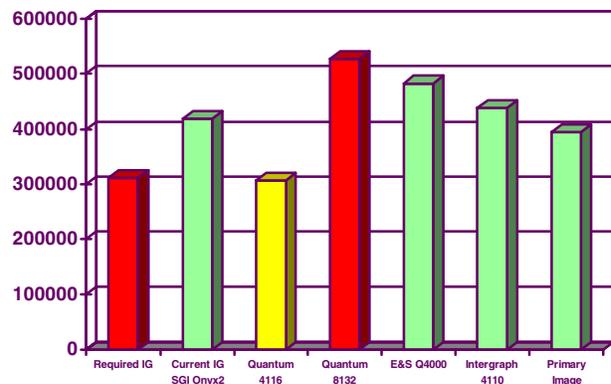


Figure 2. Polygon Performance Benchmark (Simulation Triangles/Sec)

The hardware Fill rate was measured using an Indy3D benchmark across most of the systems as noted (Figure 3). When the systems were benchmarked, there was some reduction of performance due to fill rate on two systems: the Intergraph 4110 and the Quantum 4116.

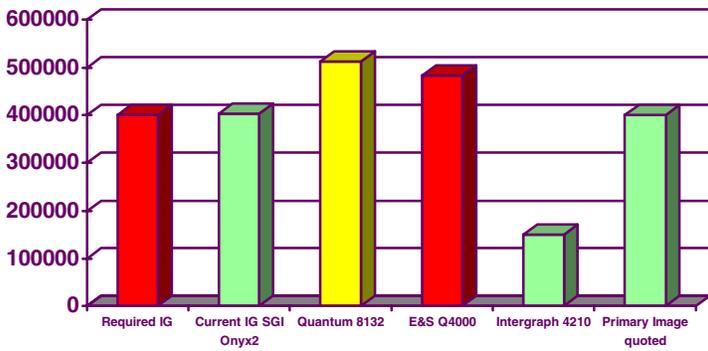


Figure 3. Fill Rate Performance Benchmark (Pixels / Second)

In terms of TPS application software performance, the current TPS system's rendering performance is 9 benchmark targets in 9 frustum FOVs within 1,280 x 1,024 at 60 Hz target database update rate. This information has been observed on the systems in simulation use. During the evaluation, three PCIG systems were integrated with the VIDS. The first was an interim PCIG, Quantum3D 4116 (Aalchemy series) offering. This system demonstrated 6 targets at 60 Hz with some fill-rate limitations at close target ranges. The second, E&S SimFusion 4000q generated 9 targets at 60 Hz and demonstrated much improved fill rate capability at close target ranges. The last, Quantum3D 8132 (Aalchemy series) (the selected production system) generated 11 integrated targets at 60 Hz.

After the initial software and hardware problems were corrected in design (see the next section), we were able to integrate a new PCIG in less than a day with the simulator. This shows the promise of open architecture for our future development of simulation systems. Not only does this provide faster and simpler integration in complex simulation environments, but also provides a development environment that allows the team to develop and test offline with minimal capital expenses. Most of the development for this system was done on an OpenGL card in a standard PC. Also a host emulator was designed to emulate a complete set of host commands for initial development. But having a complete development system attached to the simulator during integration was very helpful nevertheless.

During the time of this study, all of these products were very early in their production cycle; some were not yet in production. The

numbers captured from these benchmarks would continue to improve, as the vendors' graphics drivers would mature. Also, the vendors continually upgrade their motherboards with faster CPUs, having starting with us using PCs based on P-III 450-MHz CPUs and now (only 8 months later) shipping PCs with P-IIIs over 900 MHz (double speed). These PCIG systems will continue to improve because of these factors and because of the scalable architecture that is coming this year.

RUN-TIME IMAGE GENERATION APPLICATION DEVELOPMENT ISSUES

Integration was accomplished using a (non-Quantum3D) software development PC with an OpenGL graphics card and downloading the software to the selected PCIG system.

The initial integration was accomplished on the development machine connected to the main network hub on the simulator so that real time integration could be optimized using the compile and run option in Visual C++. This allowed main system problems fixes. But the final test was to move the TPS PCIG executable program to the Quantum3D where its flight simulation performance could be accurately evaluated rather than by simply looking at the output of the integration machine monitor.

Software Development Issues

One general problem in the API application that we observed was that even when the API was provided for open PCIG architecture, the PCIG vendors had generally worked with only one API. They had optimized their graphics engine drivers for one particular API that they had favored (with whom they had a re-seller agreement in place). This became clearly evident during our testing on systems that did not re-sell that one particular API: we saw a slight reduction of graphics performance in the benchmarks using "our" chosen API.

Software integration problems that were identified also stemmed from the multiprocessing configuration. The App and Draw stages in TPS PCIG have to align exactly because of the dynamic viewport requirements. Here the frustum and viewport rotation parameters had to be delayed to align the App/Cull and Draw processes to the frame.

And because of the infancy of the PCIG hardware and scene APIs (E.g. Vega NT) there were several problems when running in multithreaded modes. In particular, there were two problems with Vega NT in multi-process mode that are now engineering bugs at Multigen-Paradigm (now Viewpoint Visualization).

First, during flight simulation, the image generator must change the position of VegaNT vgChannel at 60 Hz. The original system could place a new vgChannel in the same area where an old channel was placed last frame. This

created a blink in multithreaded mode. The application had to be rewritten to remember the old channel and not allow a channel swap from one frame to the next from the VSC.

Second, Performer pfUserData was used in the callbacks to capture node information from the Traverser. Unfortunately, objects must be deleted occasionally from the system and in multithreaded mode, the object and callback could not be removed when pfUserData was used. An expensive workaround was written to accommodate this bug.

Finally, Quantum3D's Glide was originally not written for multithreaded support; Quantum3D implemented VegaNT multithreading during our TPS integration.

Hardware Development Issues

Hardware related problems encountered were due to the input video requirements of TPS system requirement from this IG. The selected vendor Quantum3D Aalchemy series was not originally designed to GenLock to external single ended video sources. They however did have GenLock capability for multi-channel systems. Quantum3D had to make firmware or hardware design changes to accept the separate input Vertical and Horizontal sync signals. Another requirement for this IG was to have a stable output Video timing during and after application startup and shutdown. These systems will initialize the video circuitry when applications are started which will disrupt the output video timing. For TPS we had to create an external design, which would receive the output syncs from the PCIG, then detect the "disruption" and substitute correct syncs until the system stabilizes. During initial integration, it was necessary to perform workarounds to provide and receive appropriate video timing signals. It is recommended that external GenLock requirements are thoroughly reviewed and GenLock testing is performed (in house) prior to selection of a PCIG supplier.

Figures 4 and 5 show scenes from the development system PC captured during an actual target simulation scenario. In Figure 4, each target shown is being displayed in a channel so in this case there are 9 channels active. The off-axis perspective correction is evident in the targets on the right, which are being presented to right and left screens from the ownship. If the targets were near each other or behind one another, they would be also be rendered in the channel as a mask (a black

representation of the model). Also, target configurations can be controlled. External configuration, animations, and articulated part interfaces have been designed into the application.



Figure 4. Sample TPS Visual Scene

Figure 5 shows an example of targets at closer range. Models on the right are being displayed on the right screen from the eyepoint and are flying in formation with the ownship. The most difficult simulation for the PCIG is flying with wingman at close range. In this case, there could be 10 or more channels open with more than one target rendered in each channel.

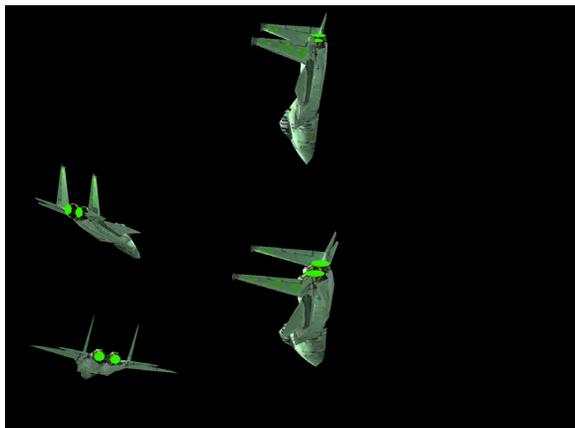


Figure 5. TPS models at close range

SUMMARY

The creation of a target image generation benchmark that is representative of the aircraft training scenario has permitted Boeing Training Systems to individually assess each PCIG vendor's product. This benchmark not only performs a series of tests on each target image

generator that are characteristic of the flight mission, it has the ability to gather statistical information required to measure the system performance of each vendor's recommended product. This approach has defined the rules of engagement, and leveled the playing field for all of the suppliers. Furthermore, it has required the evaluations to consist of performing testing in the laboratory, as opposed to comparing viewgraphs in the conference room.

Based on the comprehensive PCIG evaluation process and hardware/software integration experience, Boeing Training and Support Systems has chosen a Quantum3D Aalchemy 8132 system for the VIDS TPS production at this time. However, because of the IG software portability, other platforms have been and will be continually assessed and weighed against future production cost/performance needs.

We have demonstrated an integrated system with the performance of the high-end IG while still hoping to improve the performance in production. The unit cost is less than \$20k with the runtime license. Maintenance costs for the system are zero as we are determined to reduce this overhead and simply furnish hot spares. The production systems will be fielded with low cost hot spares and we will have approved second source suppliers. Figure 6 shows the cost comparison the systems.

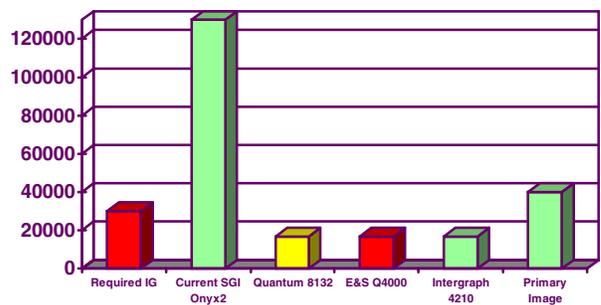


Figure 6. Cost Comparison

Future applications could benefit by applications of the newly emerging PC-based IGs with significant lower price-performance points. The additional hardware features designed into these systems such as GenLock capability and full screen anti-aliasing allowed us to complete this design. The open architecture allowed a simpler integration effort in using standard PC configurations when necessary. We also see an opportunity to reduce development costs in the future by

creating a development-and-test simulator on each developer's desk.

ACKNOWLEDGMENTS

The authors would like to acknowledge Phil Graeff and Dave Renfroe, the VIDS* Product Center Managers at Boeing TSS in St. Louis, for supporting the described effort and paper preparation.

The authors also acknowledge the following names as actual or implied trademarks belonging to their respective companies: VIDS* and TPS* to The Boeing Company; Multigen* and VegaNT* to Computer-Associates/Multigen-Paradigm; Wildcat* to Intense3D/Intergraph; Barracuda* and Cruncher* to Primary Image Ltd.; ESIG* and simFUSION to Evans and Sutherland Computer Corporation; Heavy Metal* and Aalchemy* to Quantum3D; and Onyx*, Onyx²*, and InfiniteReality* to SGI.