



Technical Brief

NVIDIA Quadro FX 3000G
Solutions for Advanced
Visualization



Solutions for Advanced Visualization

Introduction

Many interactive graphics visualization applications require the use of multiple systems to drive multiple displays or projectors (also called display channels). It is extremely important that these systems (graphics pipelines) treat the displays collectively as one continuous “virtual canvas” that behaves as though it were one large logical display. So far, the challenge of providing multiple views, or view frustums, of the same large data sets in real time with a minimum of latency has only been addressed by expensive and proprietary system clusters.

In addition, other applications for video and post-production require the ability to synchronize graphics with other input streams and recording devices for video effects, editing, and compositing.

The NVIDIA Quadro® FX 3000G graphics solution unleashes these capabilities—called “genlock,” “frame lock,” or “frame sync”—that have traditionally only been available on expensive, proprietary hardware platforms. Furthermore, the NVIDIA solution introduces revolutionary capabilities that have never been available in these proprietary platforms, including programmable shading, high-precision graphics computation, and significantly increased graphics performance.

Applications that benefit from using the NVIDIA Quadro FX 3000G include

- ❑ Room-sized, rear-projected powerwalls used for collaborative industrial design or decision support centers (Figure 1)
- ❑ Multichannel vehicle simulators used for pilot training or mission rehearsal
- ❑ Large asset stations used by geophysicists for geoseismic interpretation of large volumes and data sets
- ❑ Video editing, on-air broadcast graphics, and digital content creation

“The combination of HP’s innovative workstations and the NVIDIA Quadro FX 3000G is designed to enable teams to interact with their work in life-sized environments in real time. Auto and aerospace engineers, for instance, could review, collaborate, and visualize their designs as they exist in real life, with photorealistic quality, at a level of performance and interactivity that keeps pace with their thought processes. To visualize large assemblies at full size, they need only scale the problem across multiple systems, such as HP’s scalable visualization technology, to drive a virtually unlimited assembly size.”

Jim Zafarana, vice president of Worldwide Marketing, HP Workstation Business Unit



Figure 1. High-Performance Powerwall Display

Synchronization of these channels (each consisting of one NVIDIA Quadro FX 3000G) is critical for ensuring a continuous display between multisystem configurations, and for guaranteeing the best and consistent frame rates.

To achieve a seamless presentation, the graphics hardware must be capable of the following processes:

- ❑ Frame synchronization—synchronizing frame rendering across all displays
- ❑ Swap synchronization—synchronizing the swapping of front and back buffers

In addition, specific applications must be able to synchronize the hardware to an externally generated sync source (sometimes referred to as genlock or “house sync” signal). Examples of applications that require genlock are video compositing, editing, and broadcast.

The NVIDIA Quadro FX 3000G raises the bar by providing the frame synchronization, swap synchronization, and genlock requirements of the most demanding advanced visualization applications. At the same time, it delivers industry-leading NVIDIA Quadro FX performance, precision, programmability, and quality.

Hardware Solution

The NVIDIA Quadro FX 3000G graphics solution is based on the industry's highest performing and most advanced workstation graphics board, the NVIDIA Quadro FX 3000. With its true 128-bit floating point frame buffer, 8-pixel pipeline, and massive 256 MB/256-bit graphics memory pipeline, the NVIDIA Quadro FX 3000G retains its performance, precision, programmability, and features while adding full frame synchronization capabilities.

Interfacing through a standard AGP 8X graphics interface, the NVIDIA Quadro FX 3000G fits all standard motherboards (Figure 2). It was designed to occupy the AGP and one adjacent PCI slot (the daughterboard lies in line with the adjacent PCI slot, but uses no power nor data from the slot). And like the standard NVIDIA Quadro FX 3000, the NVIDIA Quadro FX 3000G does not require additional cooling.



Figure 2. The NVIDIA Quadro FX 3000G: An NVIDIA Quadro FX 3000 with Genlock and Swap Synchronization

The NVIDIA Quadro FX 3000G comes with this hardware:

- Two digital display ports (one single-link DVI and one dual-link DVI)
- One workstation stereo sync connector (3-pin mini-DIN)
- Two RJ-45 frame lock connectors (using standard CAT5 networking cables)
- One BNC (75 ohm, male) genlock connector

Frame Synchronization

Frame synchronization is the process of synchronizing display pixel scanning to a synchronization source. When several systems are connected, a sync signal is fed from a master system to the other systems in the network, and the displays are synchronized with each other.

The NVIDIA Quadro FX 3000G's frame synchronization lets multiple systems synchronize frame rendering across multiple displays, plus synchronize one or many systems to a common sync source. NVIDIA supports frame synchronization for both Windows and Linux operating systems, and provides full-featured Control Panels (Figures 3 and 4).



Figure 3. Frame Sync Control Panel (Microsoft Windows Operating System)

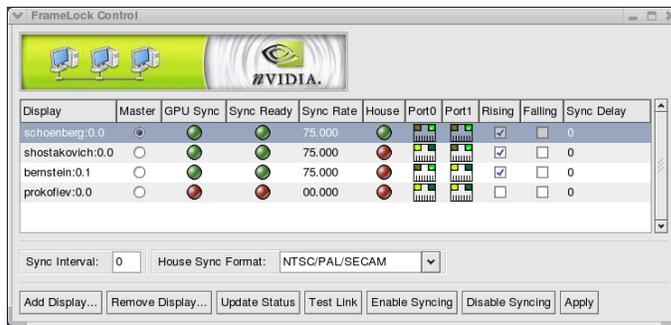


Figure 4. Frame Sync Control Panel (Linux Operating System)

NVIDIA frame synchronizing technology involves two main processes:

- ❑ Frame lock and swap—synchronizing applications across multiple displays or windows
- ❑ Genlock—synchronizing one or more systems to a common sync source

Frame Lock and Swap

Proper synchronization of an application that runs on multiple displays involves frame lock synchronization and swap synchronization.

Frame Lock

Frame lock uses hardware to synchronize the frames on each display and to redraw to multiple displays at the same time. When an application is displayed across multiple monitors, frame-locked systems help maintain image continuity to create a virtual canvas. Frame lock is especially critical for stereo viewing, where the left and right fields must be in sync across all displays.

As with frame sync, frame locking several systems requires that the systems be connected, with the sync signal fed from the master to the other systems in the group. A network of frame-locked systems can be synchronized by connecting the master system to an external synchronization source.

Swap (Swap Sync)

Swap sync refers to synchronizing buffer swaps from multiple application windows. This includes the ability to have frame-accurate synchronized displays. Using swap sync, applications running on multiple systems can synchronize the application buffer swaps between all the systems. Swap sync requires the graphics systems be frame locked, plus the buffers be swapped at the same time. The absence of swap sync introduces visual continuities that detract from the overall quality of the display.

The NVIDIA solution lets the user or application instruct one system to be the master system, while the remaining systems automatically default to “slave” mode.

Swap-sync signaling is implemented in the NVIDIA Quadro FX hardware, so latency is minimized. This ensures that all the graphics cards deployed in the solution are synchronized to the same pixel rate and that the output pixels are in complete lock-step.

Genlock

As a subset of frame synchronization, genlock is commonly used in video post-production, nonlinear editing (NLE), and broadcast studios (Figure 5). Its use ensures that a workstation's graphics output is locked (LOCK) to an externally generated (GEN) signal, guaranteeing that the studio's devices (cameras, videotape recorders, character or title generators, and so on) work together effectively. The NVIDIA Quadro FX 3000G graphics hardware ensures subscanline-level synchronization to the external sync source.



Figure 5. Genlock Used in a Typical NLE Video Environment

Through the BNC genlock connector on board the NVIDIA Quadro FX 3000G, the graphics video timing can be synchronized to an external timing source. In turn, this source provides a periodic signal to the display system to lock the vertical refresh rate.

Typically, a workstation user plugs a “house sync” signal into the graphics board—usually set at NTSC (U.S.) or PAL (Europe) rates—instructing the graphics board to trigger various edges of the external sync signal.

API Extensions

NVIDIA provides a set of API extensions for Linux (X and GLX) and Windows (WGL) operating systems to enable applications to take maximum advantage of NVIDIA's Quadro FX 3000G hardware capabilities.

Descriptions of the most common extensions are provided. For additional information—such as the complete list of extensions and capabilities, plus programming guidelines—please refer to *NVIDIA Frame Lock User's Guide*, available in late summer of 2003 through the NVIDIA Developers Program (www.developer.nvidia.com).

WGL_NV_swap_group (Windows), GLX_NV_swap_group (Linux)

This extension provides the capability to synchronize buffer swaps in a group of OpenGL windows. A “swap group” is created, and then windows are added as members to the swap group. Buffer swaps to members of the swap group can then take place concurrently. A “swap barrier” can also be created to synchronize swaps between systems.

The extension also provides a universal frame counter among systems locked together by swap groups/barriers. It is based on the internal synchronization signal that triggers the buffer swap.

WGL_I3D_genlock (Windows)

The genlock extensions allows application control of the genlock features. Genlock is used to synchronize monitor refresh to an external signal.

Five parameters can be used to control genlock:

- Enable/disable
- Source selection
- Source sample edge
- Source sample rate
- Source delay

NV_Control (Linux)

The NV-CONTROL X extension provides a mechanism for X clients to query and set configuration parameters of the NVIDIA accelerated graphics driver on XFree86. The state set by the NV-CONTROL X extension is assumed to be persistent only for the current server generation. Attributes are configurable on a per-X screen basis, and some attributes are also configurable on a per-display device basis.

Some attributes can only be queried, though most can be both queried and modified. The NV-CONTROL extension can determine what values are valid for an attribute, if an attribute is read-only, if it can be read and written, if it requires a display device qualifier, and if the attribute is available on the specified X screen. NV-CONTROL clients may also request notification when an attribute is changed by another NV-CONTROL client.

The NVIDIA Control Panel for Linux (refer back to Figure 4) is, in fact, a sample implementation of the NV Control X extension and is available open-source to be customized, configured, and modified.

Conclusion

By fully supporting and expanding upon the capabilities of the NVIDIA Quadro FX family, the NVIDIA Quadro FX 3000G provides unprecedented visualization capabilities and redefines the expectations for graphics professionals.

For more information about the NVIDIA Quadro FX 3000G or products using the APIs, please refer to www.nvidia.com/workstation or www.developer.nvidia.com.



Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, and NVIDIA Quadro are trademarks of NVIDIA Corporation.

Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

Copyright NVIDIA Corporation 2003.



NVIDIA.

NVIDIA Corporation
2701 San Tomas Expressway
Santa Clara, CA 95050
www.nvidia.com