

Accelerated-X™ Summit Series SingleLogicalScreen Feature

Introduction

The SingleLogicalScreen ("SLS") feature - sometimes referred to as "Xinerama" or "Stretched Desktop" - is the capability of being able to have two or more monitors (projectors) display their respective portions of an image generated by a single x screen ("screen") in the X Window System ("X"). The feature allows a window opened on one portion of the display (the x screen) to be dragged across boundaries of the monitors, and enables an array of monitors arranged in a wall display to be viewed at a distance.

Some Definitions

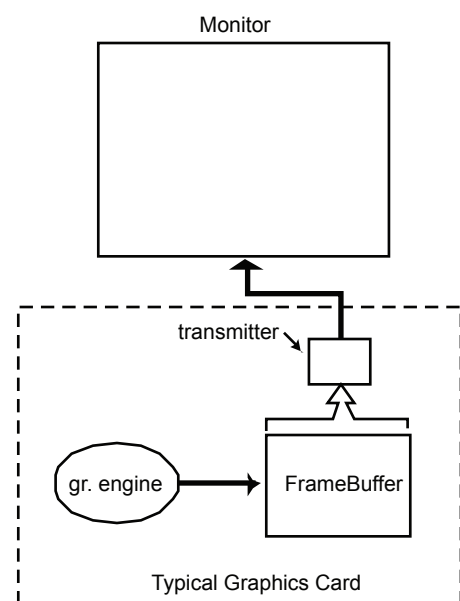
In an effort to avoid confusion over terms, the following list is presented:

- **head** - usually refers to a monitor; a "headless" server has no monitor;
- **screen** or "xscreen" - a root window in X;
- **display** - container for one or more xscreens;
- **channel** - (not official X term) the path from the graphics chip's FrameBuffer to a monitor (through the transmitter and cable);
- **FrameBuffer** - the block of memory on the graphics chip or in RAM that contains the image data to be displayed by one or several monitors - capital letters used for style.
- **graphics chip** - the hardware that incorporates a particular graphics architecture used to create (rasterize) an image into a FrameBuffer for subsequent display;
- **monitor** - a display device used to display all or a portion of a display image;
- **rotate** - the process of re-arranging the image in a FrameBuffer so it can be displayed in a portrait mode on a display device (monitor) that has been physically rotated from a horizontal to a vertical orientation;
- **overlay** - the practice of having a (usually small) window (usually) temporarily appear in front of an underlying (usually) complex image without changing the underlying image data in the FrameBuffer;
- **dix/ddx** - (graphics) device independent and device dependent pieces of an X server;
- **X Protocol Packets** - drawing (and other) commands are sent to the X server by an applications program that is drawing an image for display via packets.

Basic Configuration - Single Head

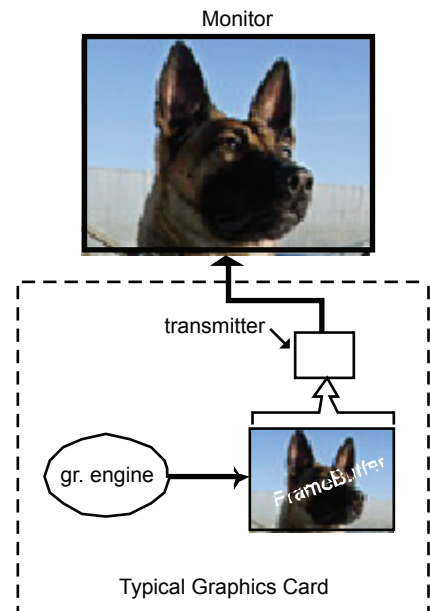
The illustration at the right shows a simplified block diagram of a graphics card with a single graphics chip driving a single monitor. The FrameBuffer may be on the graphics chip/card, or it may use the computer system's RAM memory. Similarly, the transmitter that transmits image data to the monitor may be a separate chip or may be incorporated into the graphics chip.

Note the horizontal orientation of the display and FrameBuffer. The transmitter is programmed to fetch the image data in the FrameBuffer and send it to the monitor synchronously with the sweep of the pixel gun (figuratively speaking) in the monitor. In this case, the whole FrameBuffer is being "serviced" by the single transmitter and monitor.



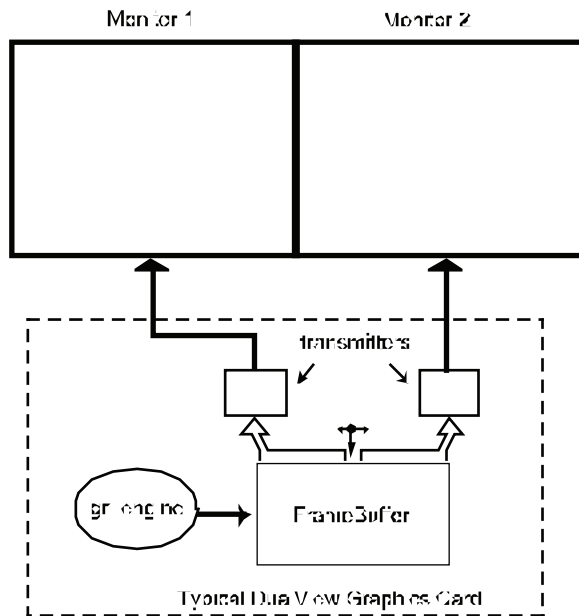
This configuration, although very simple, is (very technically speaking) a single logical screen configuration - one xscreen is being displayed on one monitor. The SLS term, however, is usually thought of as applying only when there is more than one monitor involved.

Before we move to more involved configurations, on the right is the same single head configuration illustration with some color added.



Dual Head Configuration

The Matrox G400 graphics card brought the idea of a single card driving two monitors to the masses. They called the feature "DualHead," and it quickly became popular. Below is a block diagram of a DualHead (Xi Graphics calls it DualView) configuration. By increasing the FrameBuffer size and adding another



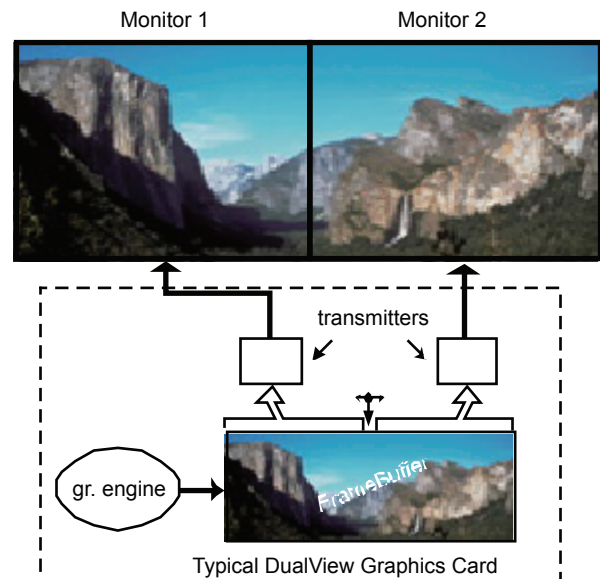
another channel (the second transmitter) and other appropriate hardware to allow fast access to the FrameBuffer by the several paths, two monitors could be serviced by one graphics card/slot. Note that the way the illustration is drawn, the image in the FrameBuffer is "stretched" across the two monitors creating, by default, the stretched desktop, or SLS configuration, since the single image is the result of the single xscreen in a single FrameBuffer.

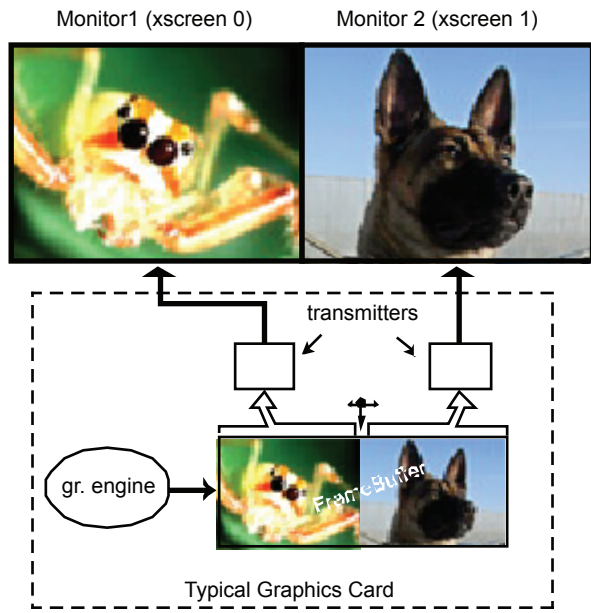
This SLS configuration is essentially a "freebie", in the Summit Series software since a single FrameBuffer is involved and the graphics hardware is doing most of the work. SLS

configurations are available in Summit Series Desktop (DX) and Workstation (WX) products when the graphics card can drive more than one monitor from a single FrameBuffer. (When there is more than one FrameBuffer involved, then Summit Series SingleLogicalScreen (HX) products are required).

Although not shown, SLS can be configured in a "stacked monitor" (vertical) configuration.

Summit Series products also allow both monitors display the same portion (the "SimulView" configuration) of the FrameBuffer image. Thus one monitor can be remote from the other, with each displaying the same image (useful in hospitals, for example).





Two Screens Per FrameBuffer

If the software complexity is increased a bit more, it is possible to have more than one xscreen generating an image into the FrameBuffer. For example, half of the FrameBuffer could have the image created by one xscreen and the other half could hold the image created by the second xscreen. Xi Graphics' Accelerated-X Summit Series allows two fully hardware-accelerated xscreens per FrameBuffer to be configured. Two per FrameBuffer is the limit imposed upon Summit Series products currently, even if the graphics card can support three or four monitors.

Quad View

Skipping over Triple View (which some of the new Matrox Parhelia cards feature), we jump to

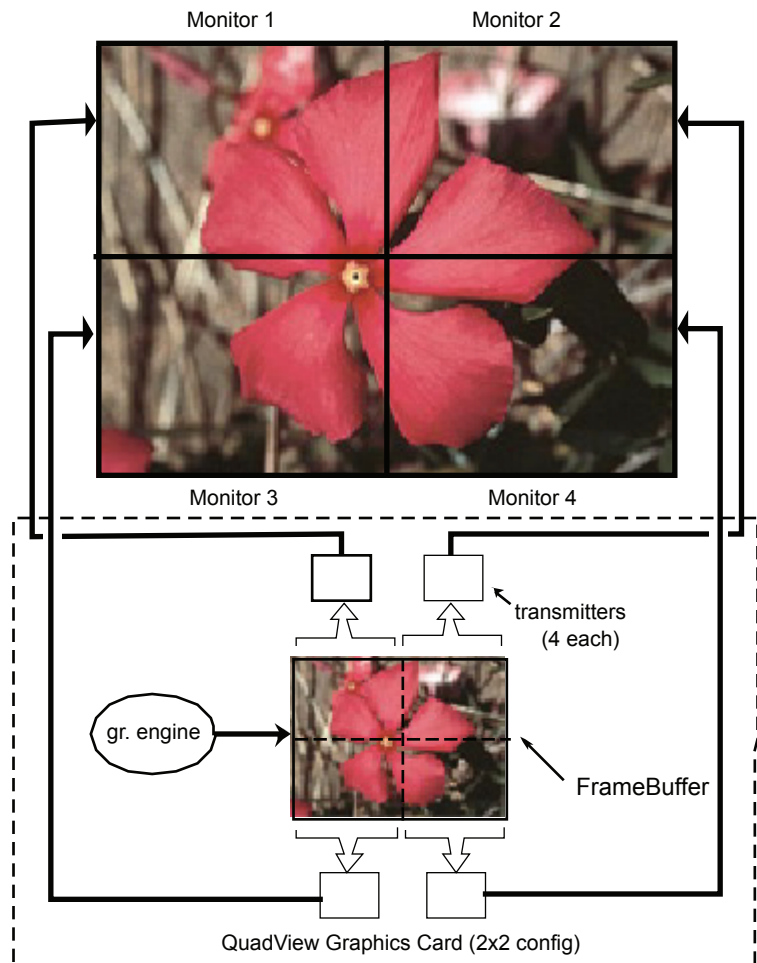
the Matrox QID card as an example of a card that can drive four monitors from a single graphics engine. The "back end" of the card has four transmitters servicing the single FrameBuffer. The 2x2 configuration shown below is the only QuadView configuration that Summit Series currently allows with high resolution monitors, but it is pretty impressive as an SLS display which is available in the Workstation (WX) Series (including OpenGL 3D images when using the QID Pro card).

The four monitors can be arranged in a 1 high x 4 wide configuration, but the horizontal resolution of each monitor can not exceed 1024 (a current software limitation in the Summit Series) when used in a desktop system.

For a wall display, however, the card really shines. Summit Series software hardware-accelerates all four pieces of the image being stretched across the four monitors (something that the XFree86/x.org software seems not able to do correctly, even with 2D images).

And two QID cards makes a sizable wall display. That gets us to the matter of multiple FrameBuffers and SLS and why the Accelerated-X Summit Series has the new SLS (HX) Series.

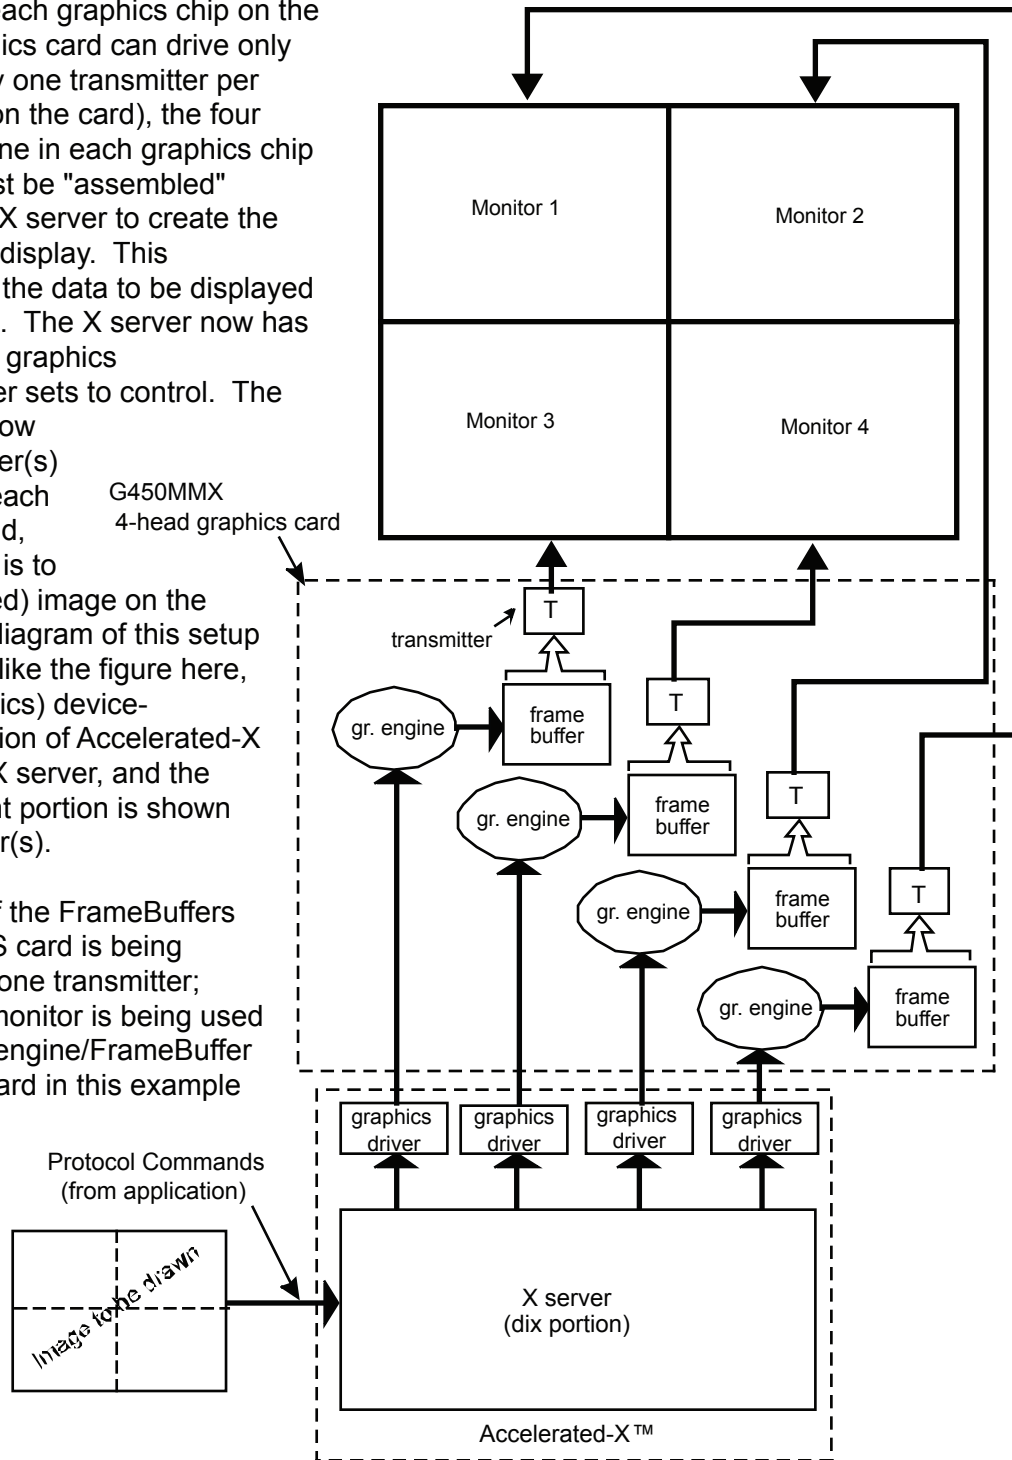
Using a Matrox G450MMX card with four graphics chips/engines, each capable of driving one monitor from its FrameBuffer,



the same wall display can be configured. However, since each graphics chip on the G450MMS graphics card can drive only one monitor (only one transmitter per graphics chip is on the card), the four FrameBuffers - one in each graphics chip in this case - must be "assembled" (logically) by the X server to create the desired array for display. This "assembling" of the data to be displayed is rather involved. The X server now has four independent graphics chips/FrameBuffer sets to control. The X server must know which FrameBuffer(s) are involved for each drawing command, since the display is to be one big (unified) image on the wall. The block diagram of this setup looks something like the figure here, where the (graphics) device-independent portion of Accelerated-X is shown as the X server, and the device-dependent portion is shown as graphics driver(s).

Note that each of the FrameBuffers on the G450MMS card is being serviced by only one transmitter; hence only one monitor is being used by one graphics engine/FrameBuffer pair. While the card in this example is using four chipsets to produce the image on four monitors, and the QID card in the previous example used only one graphics chip (graphics

engine), some ATI cards (and Colorgraphic cards which use ATI graphics engines) use two graphics engine/FrameBuffer pairs to drive four monitors. The G450MMS example, though, brings out pretty clearly the complexity of the SLS configuration with multiple chips. With the Colorgraphic Xentera GT8, which can drive eight monitors from one card with four graphics chip/FrameBuffer pairs, is even more complex (and takes up too much space to illustrate here).



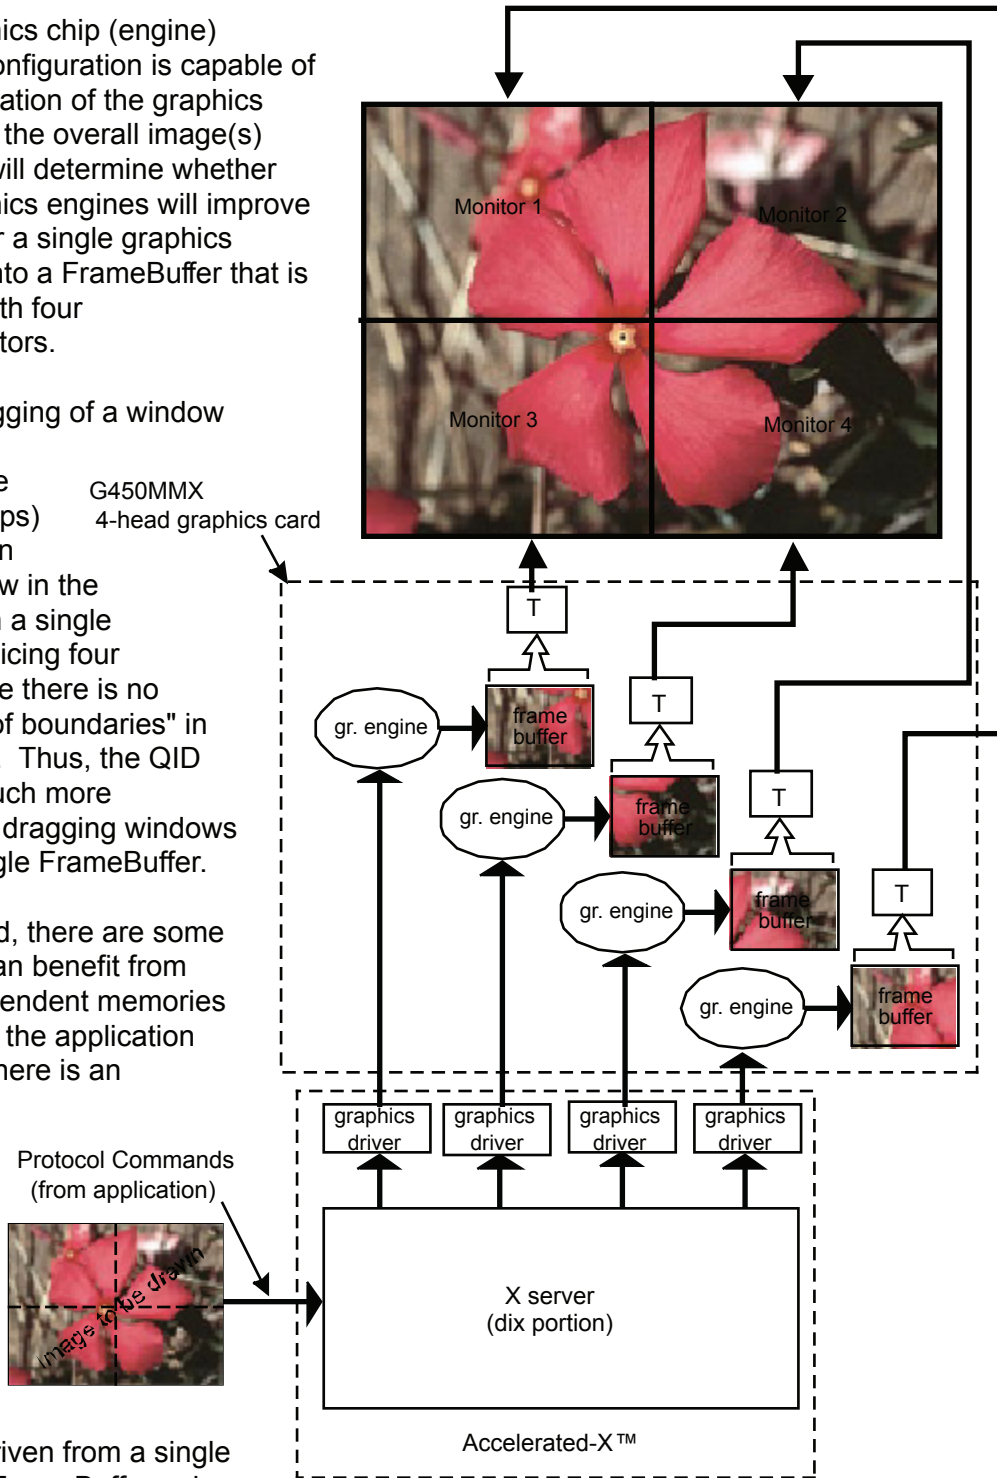
Here is the previous block diagram with some "color" added. (Clearly, no one would generate such an image using X11 Protocol drawing commands - but the jpeg images are pretty and show the desired effect).

While each graphics chip (engine) involved in this configuration is capable of hardware acceleration of the graphics drawing process, the overall image(s) being produced will determine whether having four graphics engines will improve performance over a single graphics engine drawing into a FrameBuffer that is being serviced with four transmitters/monitors.

Certainly the dragging of a window around in the configuration here (four graphics chips) will be slower than dragging a window in the configuration with a single FrameBuffer servicing four monitors, because there is no actual "crossing of boundaries" in the former config. Thus, the QID card would be much more responsive when dragging windows around in the single FrameBuffer.

On the other hand, there are some operations that can benefit from having four independent memories to work with. So, the application will determine if there is an advantage to having more than one graphics engine.

For fun, the next illustration shows a wall display with six monitors, each driven from a single graphics engine/FrameBuffer pair with a single transmitter per FrameBuffer - again using a jpeg image for effect.

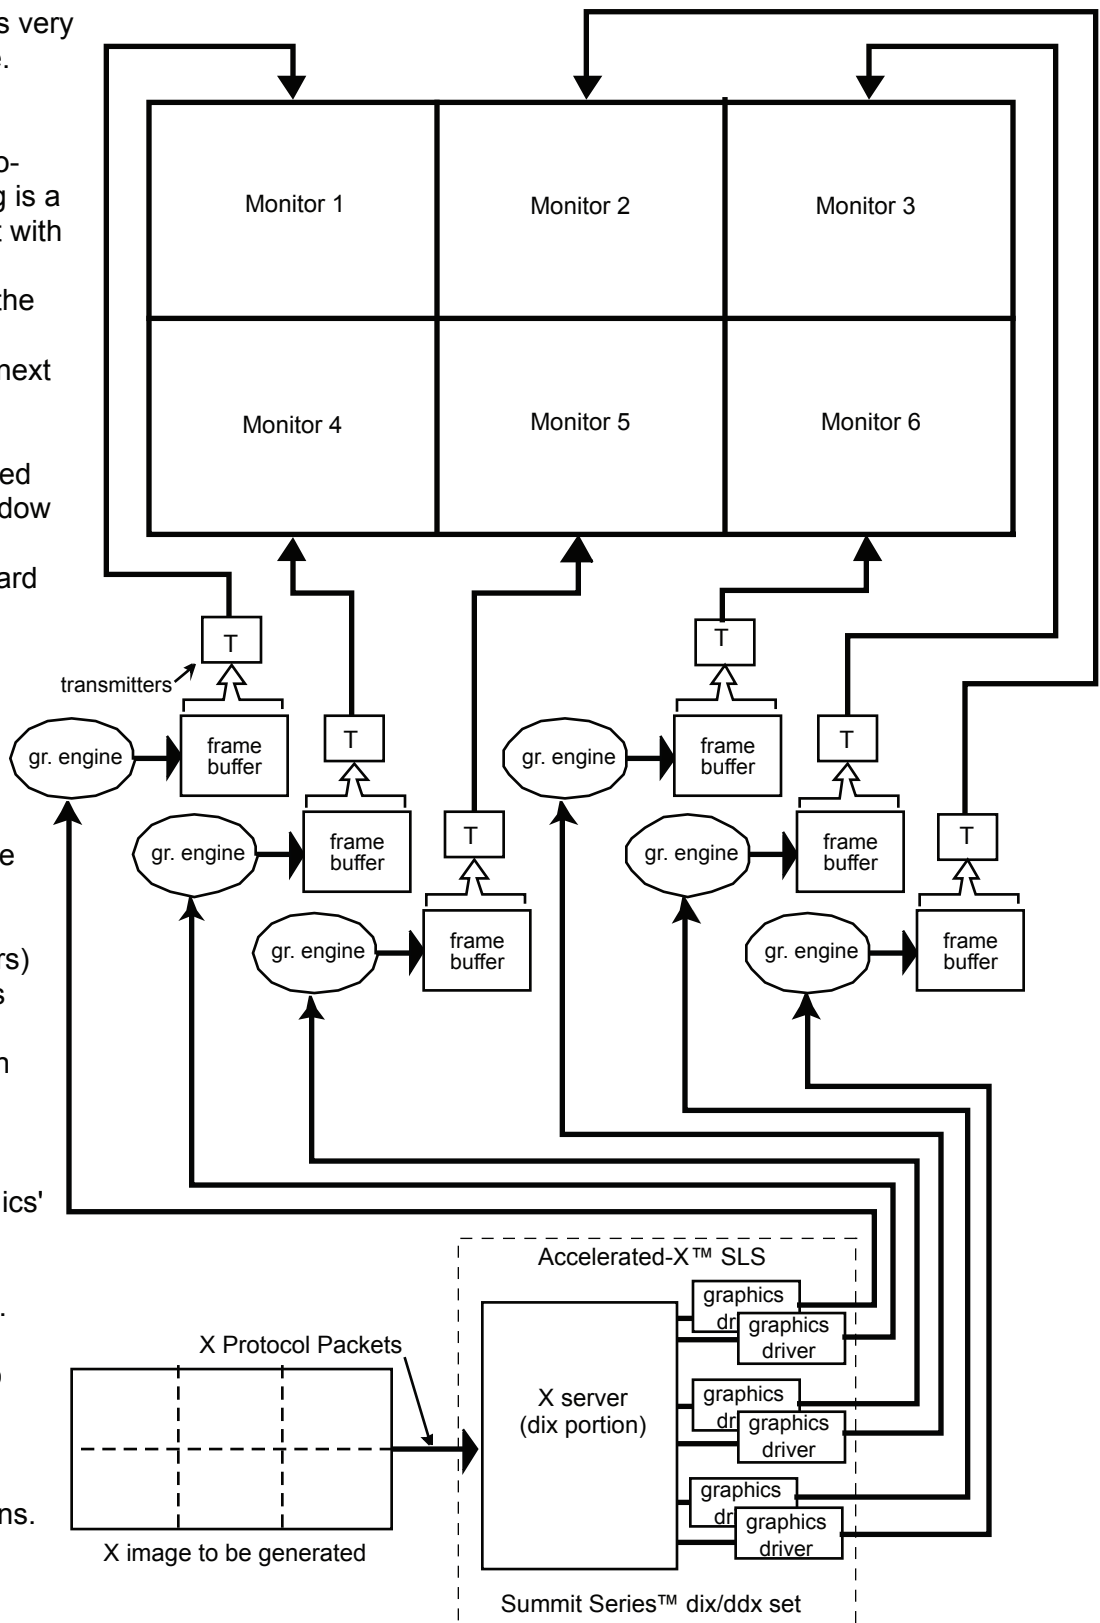


Drawing (and other) commands are sent to the X Window System X server (the device independent part of the X Window System) via "X Protocol Packets" by an application. Those packets are interpreted by the X server and executed to generate the display. Obviously, the graphics system configuration is very important here.

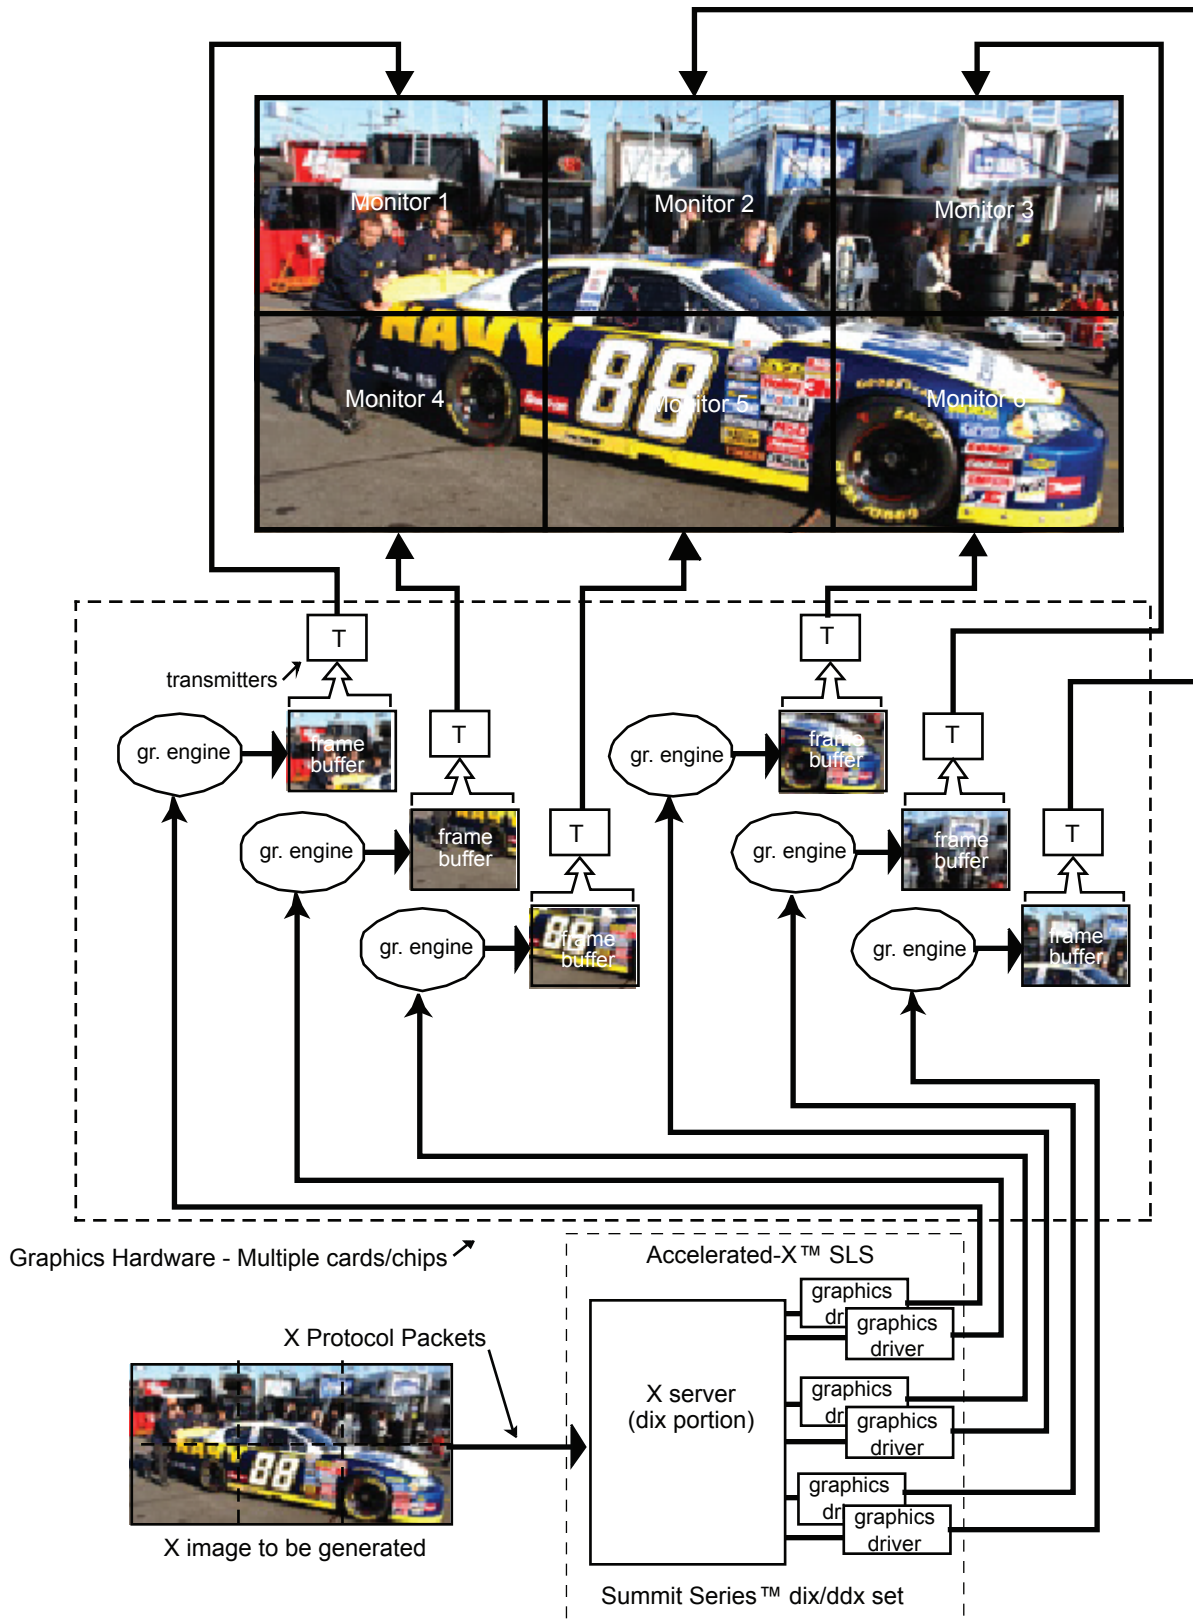
Note that the FrameBuffer-to-monitor pairing is a bit chaotic, but with the proper assignments, the results are as shown on the next page.

It might be noted that the X Window System is an industry standard graphics sub-system for UNIX operating systems, and Accelerated-X Summit Series software products (X servers and graphics drivers) conform to this standard, and perform to high commercial standards.

Also, Xi Graphics' standard maintenance support is free. Customized support is also available for medical and other mission critical situations.

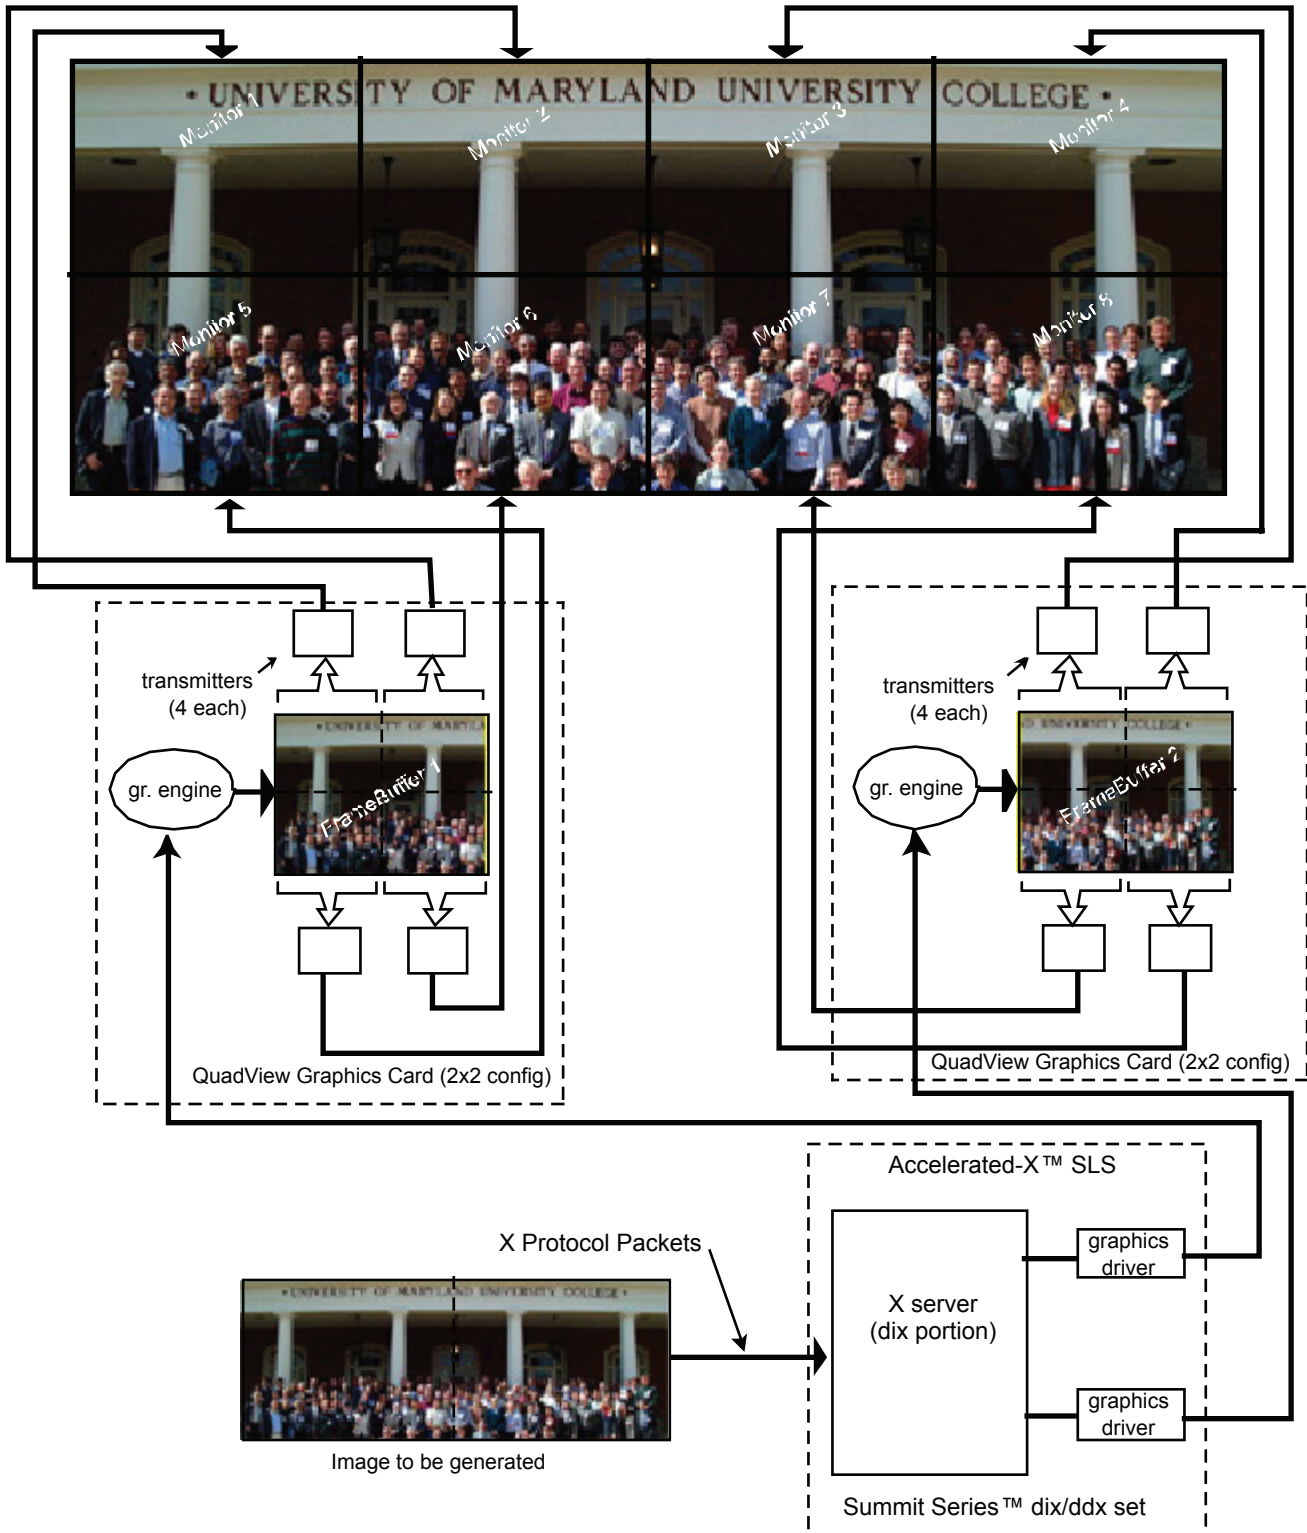


Because Accelerated-X is in use on US Navy ships, the color for this illustration is fitting. Fitting graphics will have to be found for the US Army and Marine Corps, too, since they also use Accelerated-X graphics software.



Dual Quad View Example

To illustrate the versatility of SLS with graphics cards that employ multiple transmitters per graphics chip, we chose a configuration of two Matrox QID cards. Each card has a single graphics engine rasterizing into a single FrameBuffer serviced by four transmitters. Thus the two cards can drive a 2x4 wall display as shown in the diagram below. There is only one boundary to be crossed (vertical in the middle), so dragging a window around is very responsive until that boundary is crossed. Then it slows down a bit.



Contrast that with a 2x4 wall display that uses two Matrox G540MMS cards, where each monitor is driven from a separate FrameBuffer. This means that there are many more FrameBuffer boundaries to be crossed as a window is dragged around the display, placing a heavier load on the X server to manage all the data transfers and keeping track of things.

Standard & Custom Wall Display Support Available

The many possible configurations and system requirements encountered with graphics display applications preclude standard graphics sub-system software offerings that can cover all, or even a majority, of applications. Hence, Xi Graphics offers both standard off-the-shelf configurations and can customize systems to fit specialized needs. Systems with Accelerated-X are in use around the World in small and large systems, including Air Traffic Control/Management, Medical, Manufacturing, Military, Space, Education, Energy, etc., and, of course, in the home.

dhm April 2006