Military technology insider

By John Wemekamp

High-performance graphics applications migrate to PCI Express

Since the late 1990s, the Accelerated Graphics Port (AGP) has been the desktop standard for connecting high-performance graphics processing devices to a computing subsystem. The AGP was developed by Intel as an alternative to PCI, providing much higher bandwidth by means of a dedicated 32-bit data path between processor, memory, and the graphics processing device. Power Architecture (PowerPC) devices – used extensively for embedded, real-time applications in both military and commercial applications – do not offer the same level of support for AGP, most using PCI to hook up to graphics devices. This allows more flexibility than AGP in that multiple graphics ports can be easily offered but have less data bandwidth for the truly challenging applications. However, AGP is being phased out of the desktop to be replaced by PCI Express, providing new levels of performance and easy migration to Power Architecture-based products.

There are two application classes that will benefit from the move to PCI Express:

- Live video streams into multiple windows on tactical situation displays: These are typically used in Naval combat systems, airborne surveillance, and maritime patrol aircraft, where operators build a multisensor, 3-D situational picture over large areas of land or ocean.
- Rapid, full-screen updates of complex synthetic data, such as digital maps.

Tactical situation displays

Live sensor video for display in windows comes in many forms. These include TV-standard analog PAL/NTSC or RGB, digital formats such as DVI, or a custom digital format from radar, for example. Video input rates may vary from 300 Kbps for a typical 25 Hz, low-resolution streaming video channel up to 4.95 Gbps for HDTV. It is very unlikely that more than one HDTV channel will be displayed in a window on a tactical situation display, even at the typical 2K x 2K resolution used by these systems. This is because the norm is from two to four windows at TV resolution at the sensor's frame rate.

To be displayed on screen in a window, video must be captured, passed to memory, and then overlaid with synthetic graphics and output by a Graphics Processor Unit (GPU) for display. The GPU, in addition to creating the synthetic graphics of the situational display, organizes the incoming video streams to produce a single seamless display on the operator's screen. The PMC/XMC mezzanine standard is a popular choice for implementing video capture and graphical display functions. It is most often used with VMEbus modules to form embedded display computing systems in military applications. For many current implementations of VMEbus SBCs using Power Architecture, the only practical way to connect a high-performance GPU to the processing system is via PCI (in some cases through a PCI or PCI-X to AGP bridge), requiring a dedicated PCI bus to achieve the bandwidth required to system memory.

However, the GPU PCI bus is needed for both the host control of the graphics operations and as the path for the incoming captured video. Thus, for these applications, current architectures are bottlenecked by the PCI bus's available bandwidth, system memory bandwidth, and the need to dedicate one PCI bus to the GPU.

> "The eight-lane PCI Express, offering up to 2 GBps bandwidth, more than quadruples the theoretical bandwidth of the 32-bit, 66 MHz PCI interfaces usually used to support the most popular families of GPUs."

Moving map displays

Moving map displays, found in many aircraft types from large transport planes to single-seat combat aircraft, are also used in armored ground vehicles. These displays often use predictive algorithms to anticipate the platform's movement, pipelining a small number of frames to the display to provide smooth transitions during motion. The displayed map is generated synthetically by the GPU so that while normal movement just requires incremental redrawing of the display, a change of context such as a range change results in the entire screen being redrawn. As in the example mentioned, this stretches the PCI bus, resulting in noticeable delays while the display is redrawn.

PCI Express

The eight-lane PCI Express, offering up to 2 GBps bandwidth, more than quadruples the theoretical bandwidth of the 32-bit, 66 MHz PCI interfaces usually used to support the most popular families of GPUs. Based on a PCI Express-enabled GPU from NVIDIA, the newly introduced XMC-710 graphics processing mezzanine from Curtiss-Wright Controls Embedded Computing (Figure 1) is ideally suited for windowing and moving map display applications.

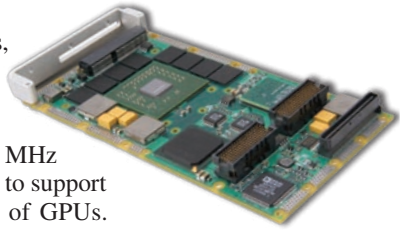

Figure 1

The introduction of PCI Express-capable video capture devices, GPUs, and the latest generation of VMEbus SBCs has leveled the playing field, ending the AGP versus PCI debate and offering new levels of flexibility and performance from Power Architecture-based display systems. PCI Express has the growth potential to remain first choice for these and many other high-performance applications for now and the foreseeable future.

For more information, e-mail John at john.wemekamp@curtisswright.com.