Real-time graphics that do it all: Visualization, image processing, and network-centric transport

By Eddy Vermeulen

When it comes to visualization and processing of electro-optical sensors, military surveillance and reconnaissance applications have a very specific set of requirements and related technical challenges. Recent technology improvements such as compression, FPGAs, and high-speed networks now make it possible to provide real-time visualization, networked transport, recording, and image processing, all in a single-board implementation instead of racks or boxes of equipment.

Requirement	Challenges	Technologies	Implementation
Real-time visualization	High-resolution video requires high bandwidth	Switch fabrics	Framegrabber
	Unstable analog signals	Automatic phase adjust	
Networked transport	High data load Minimize delay	JPEG2000 compression Gigabit Ethernet	JPEG2000 codec Onboard CPU, Ethernet and storage interface
Recording and playback	High data load	JPEG2000 compression	
Real-time processing	Real-time implementation of complex algorithms	High-capacity FPGAs	Programmable video processor

Table 1

Video and imaging applications come with a variety of requirements, ranging from displaying and storing high-resolution video to routing image information quickly from sensor to operator. And increasingly, that information is shared and processed by multiple operators often located in separate locations.

Beyond these requirements, which are somewhat "classical" in nature, it's often desirable to do more video and image processing right at the sensor itself. This puts intelligence at the point of origin and makes operators' lives easier by displaying only relevant information, or taking some of the strain off of data networks. Additionally, real-time imaging and visualization is a mandate in mission-critical applications such as military or civilian transportation systems.

Fortunately, technology has kept pace with application requirements. Gigabit Ethernet, switch fabrics, the increasing capacity of FPGAs, and modern compression standards are all new technologies that act as the finely honed tools in modern image processing applications. Instead of consuming large spaces and huge amounts of power, it is now possible to combine all of these video and imaging requirements, plus all of these technologies, together in single-card COTS solutions.

Real-time visualization requirements

The specific market requirements and associated technical challenges for several imaging and visualization categories are shown in Table 1. Most critical to an Operational Display System (ODS) is the real-time visualization of a sensor. A delay of only a few frames in a defense application or a search and rescue mission can have life-critical consequences. For example, interactive applications (such as remote control of unmanned vehicles) imply another kind of delay restriction. Applications should remain below human perception in order to not confuse the operator. This requires a prerequisite amount of technology "intelligence" and connectivity.

Many of these applications prefer a networked approach. Not only does this adhere to the generic network-centric way of system implementation, but also it provides cost benefits by relying on a single type of COTS cabling for all data communication. The networked approach allows for remote operation and observation in a safe and centralized environment and makes it possible to fuse information from different remote sensors that used to be too far apart to benefit from combined information content.

Recording and archiving sensor information and operator actions is also a desirable feature. With information stored for later retrieval, operators can review missions for analysis or training purposes. Additionally, in-mission replay allows observers to look back in time for things they might have missed in real time only moments ago. The key advantage is that the mission is still in process and it's not too late to take additional action.

Having a large quantity of sensors multiplies the amount of information, and higher sensor resolutions increase the coverage of one sensor. But these trends also put more load on the operator, requiring more human effort in visually processing the picture. Image enhancement techniques such as image fusion, video stabilization, and others can help to extract information in an easier way. Motion detection algorithms can further help in avoiding operator fatigue and improving effectiveness. All of these requirements drive the need for advanced real-time image processing power.

Capturing, transporting, processing, and storing huge quantities of electro-optical

Reprinted from ***VMEbus Systems*** / October 2005

sensor data are typical tasks to be performed in the ODS chain. New technologies that act as well-honed tools in modern image processing applications include Gigabit Ethernet, switch fabrics, the increasing capacity of FPGAs, and modern compression standards.

The art of digital framegrabbing

Without doubt, a digital approach in handling sensor data is the preferred way to drive digital flat panels and optionally process, store, and transport dynamic imagery. While digital sensors are clearly the wave of the near future, many cameras today still provide the classic analog video output, hence the need for digitization. When it comes to full motion video, this is commonly referred to as *framegrabbing.*

This analog-to-digital conversion may sound like common technology, but the huge variety in video standards – ranging from simple NTSC or PAL cameras to high-resolution computer graphics – requires a flexible front-end input that's capable of easy adaptation to any of these input types. Typical auto-detection algorithms work on the digital sample data, measuring and comparing to known timing entries in a database for parameters such as frequency, the length of active video inside the signal, and refresh rate. Algorithms become more complex when nonstandard timing is involved, requiring more advanced FPGA solutions at the input stage.

Moreover, the real world often provides signals that are far from the ideal video-generator patterns. Environmental variations such as temperature and component specification margins affect the stability in time of a video signal. When sampling such signals, a varying phase shows up as a blurry image. A technique called *Automatic Phase Adjust (APA)* prevents these kinds of artifacts by dynamically adapting the sampling process to the periodic shift in signal timings. Figure 1 (a) shows the image without APA, while Figure 1 (b) shows how APA compensates for unstable analog signals through digital enhancement.

Once the signal is digitized, a further challenge is imposed in moving around the large amount of digital data to the next processing block. A 1600 x 1200 graphics source, for example, generates more than a 2 Gbps data load. Switch fabric technology provides high-speed pathways for digital data transfer between components, boards, and/or devices.

Figure 1 (a)

Figure 1 (b)

Data transport and storage

Now what about moving the data over long distances? Here's where the networking part comes into the picture. Gigabit Ethernet continues to become more commonplace. Ethernet already provides an impressive bandwith on traditional physical layer transports ranging from copper to optical media, but it is still not sufficient for transferring high-resolution motion imagery. So we need compression. The same applies when archiving high-resolution graphics: Even with the ever-increasing storage capacity and transfer speed of modern storage devices, image compression is required to handle recordings that may last hours or even days.

Most popular motion video compression schemes such as MPEG use techniques that encode the difference between subsequent frames instead of encoding each frame individually. Relying on the fact that natural imagery mostly changes in a smooth fashion rather than through sudden disruptions, this technique is indeed an effective way of strongly compressing motion video.

A predefined sequence of predictive frames (encoding the differences to the preceding frame), bidirectional frames (relying on both preceding and following frame), and independent intra-frames form a group of pictures. Due to frame interrelation, the compression and decompression process requires a complete unit of frames, commonly consisting of 15 frames. This buffering process introduces a serious delay between source and visualization in a streaming solution. While this delay is not disturbing in multimedia or even in security applications, it is unacceptable in life-critical applications where life and death action happens in only seconds.

When compressing still pictures, JPEG compression is a well-known standard, popular through its Internet usage. The more recent JPEG2000 introduces wavelet (predefined mathematical functions of time) transformation instead of the earlier discrete cosine transformation. When considering motion video as a sequence of still frames, JPEG2000 can be used for encoding video with low latency as there is no interrelation between subsequent encoded frames.

If one difficulty with JPEG2000 encoding remains, it is the complexity of the algorithm, requiring large processing power even out of reach for contemporary COTS processors. It is only the recent availability of dedicated encoding chips that makes this a practical solution.

Real-time video processing

Running the video stream through a programmable video processor adds functionality that is restricted only by the imagination and the ingenuity of algorithm programmers. While contemporary computers are very capable of performing these actions on archived motion video, they still require significant computing time, which is not suitable for real-time observation. The use of high-capacity FPGAs such as the Xilinx Virtex-II Pro can implement complex image enhancement features in real time on a frame-by-frame basis.

This needs to be a customizable feature, as requirements are very specific for each application. A rolling picture presented to an airborne camera, for example, is not ideally compensated in the same way as the typical vertical movement of a moving vehicle-mounted camera. The flexibility of an FPGA implementation enables this customization.

Implementation

A practical example of applying these technologies is found in Barco's Flexivision III product platform as illustrated in Figure 2. Framegrabber modules, JPEG2000 codecs, and programmable video processor mezzanines form function blocks to be plugged on a carrier board that controls overall functionality and internal routing of the video signals. The acceptance of an eventual digital video standard is anticipated by space for a small personality module interface to capture whatever future sensor output standard may emerge.

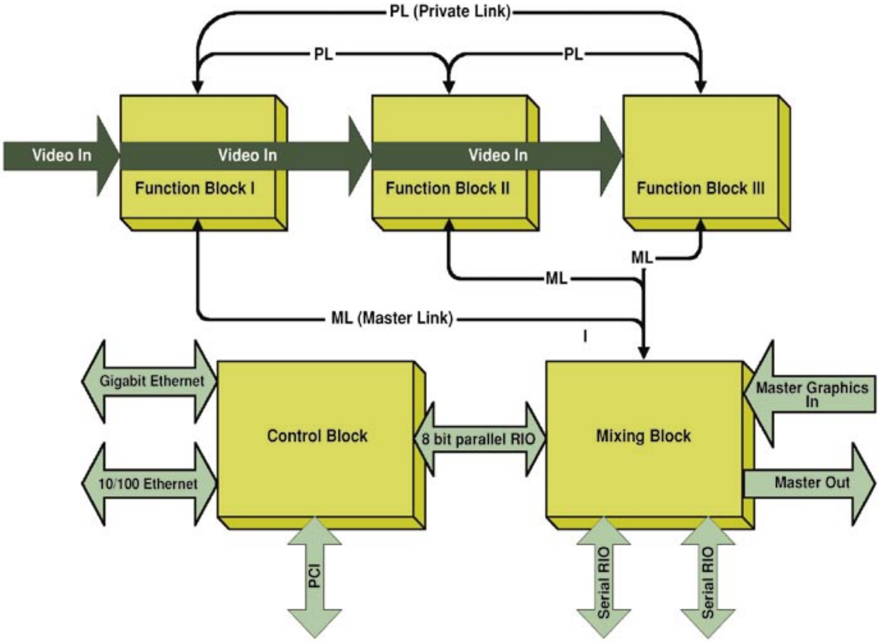

Figure 2

RapidIO switch fabric technology is used for the onboard data path because of its suitability for chip-to-chip and board-to-board connections and because of the maturity of available solutions. Serial RapidIO provides a path for board interconnections that may be required to build more complex configurations.

The mixing block combines high-resolution graphics, referred to as the *master graphics,* on a pixel-by-pixel basis with the captured video input(s). Decision criteria such as the color of the master graphics allow defining what appears under or above the inserted video, thus providing graphics overlay functionality.

The control block holds a PowerPC running embedded Linux and a full-featured Web server that allows for a Web-based graphical interface to set up a specific configuration without writing a single line of code. The lower-level API libraries are available to embed in a customized software application.

We end up with a fully self-contained solution capable of solving a wide range of real-time visualization problems through modularity. While still being a board-level solution, the integrated aspect makes the solution modular for VME, PCI, and even embedded form factors, depending on the environmental requirements of a project. The networked control aspect of the solution makes implementation easier for different target platforms, shortens the integration time, and allows for fast prototyping. An innovative combination and implementation of new technologies results in a flexible solution for a variety of visualization problems. Mission accomplished.

Eddy Vermeulen *has worked as a product manager for the international market in the defense and security market unit of Barco since 2004. As product manager, he supports and helps define imaging and network-centric products. He has worked at Barco in various product support positions in the defense and security and air traffic control markets since 1999. Prior to Barco, he worked in technical support positions in the commercial computing and networking markets.*

To learn more, contact Eddy at:

BarcoView
Kennedypark 35, B-8500
Kortrijk, Belgium
Tel: 32-56-233-202
Fax: 32-56-233-013
E-mail: eddy.vermeulen@barco.com
Website: www.barco.com/defense