

Increasing Performance of Medical Imaging Systems

— by John Phillips, Pleora Technologies

Image-guided surgery relies on real-time video with no noticeable delay between a surgeon's visual perception and the image feed used to precisely guide their surgical instrument. This article discusses the effects of latency and jitter in medical imaging systems, and the benefits of multicasting in a networked hospital operating room. The article also introduces the use of jumbo frames to reduce CPU requirements and improve overall medical imaging system performance.



Ultra-Low Latency and Jitter in GigE Vision Systems

Ultra-low latency and jitter are important characteristics of any digital communications network, and are closely related. The levels of latency and jitter will determine the quality and timely delivery of the video. Minimizing both is highly desirable, particularly where real-time analysis is required.

Latency is the total delay between the sensor and the receiver (display or image processing application), whereas jitter is the variation in the delay. The link between latency and jitter is that a network with constant latency has no jitter, whereas a network with variable latency may have high jitter.

While systems can compensate for the effects of jitter, they generally cause an increase in the delay, making latency compensation considerably more complex. Interestingly, typical compensation techniques can actually add to the problem:

- Buffering can compensate for jitter but introduces latency;
- Interleaving, redundancy, or retransmission may compensate for lost packets but can add latency and jitter, degrading overall system performance.

Many performance characteristics that are imperative to real-time video—such as low and consistent latency, low jitter, high throughput, guaranteed data delivery, and low CPU usage—can vary greatly from system to system. Of all these, low latency and low jitter are critical factors in real-time systems, because delays or inaccurate timing over certain thresholds are generally not tolerable. Research shows that the ability to perform surgical tasks begins to decrease when glass-to-glass video latency exceeds the range of 300 to 500 milliseconds (ms). Most medical imaging system designers aim for an end-to-end latency of 200 ms or less.

When Ethernet networking was initially conceived, delays in the order of seconds were considered acceptable. Even now, security video networks experience typical delays of around two seconds because most of the video is stored for later viewing. By contrast, today's sophisticated industrial vision systems have glass-to-glass latency threshold tolerances in the order of 100 ms or less.

Where dedicated electrically connected systems, such as those based on frame grabbers, exhibit very low and consistent latency, standard Ethernet packetized

systems have inherent latency and jitter, introduced, in part, by overheads in the data packets and traffic handling protocols. High-performance GigE Vision implementations largely address these issues by optimizing the packet handling process on both the transmit and receive ends of the link. Tests show that the minimal jitter introduced by the combined network connection and GigE Vision software package over that of a Camera Link system is typically <100 microseconds (μ s) — well within the tolerances of the vast majority of imaging systems.

Benefits of Networked Video

One of the key performance advantages of GigE-based distributed network architectures is the ability to locate intelligent nodes at locations where data collection and control occurs. The Internet Group Management Protocol (IGMP), implemented universally as part of Ethernet devices, network interface cards (NICs), and switches, allows network elements to be part of one or more ad-hoc video distribution groups, where a single server can multicast its data to several clients.

In a hospital operating room, for example, the surgeon can use a single






Diagram 1: Images from an x-ray detector and lamp head camera are converted to GigE Vision, processed, and multicast to various display panels in the operating room and hospital with low, consistent “glass-to-glass” latency.

display screen as a dashboard to track real-time patient data from different imaging devices and systems. The video can also be transmitted to a display panel for the OR scheduling staff, a conference room, multiple departments, and remote specialists.

At the transport layer, the server sends only one copy of the data to a network switch. The switch replicates the data for distribution to display panels and processing platforms as required. This ensures video distribution doesn’t impact server performance. Leveraging Ethernet’s inherent multicast capabilities, display and processing functions can be offloaded from a single device to multiple devices to help ensure reliability. This also helps GigE Vision cameras to multicast high-quality, uncompressed video to multiple display screens and processing nodes simultaneously with the lowest possible latency.

Using Jumbo Frames to Increase System Performance

There are a number of methods

designers can employ to increase imaging system performance. With GigE, a significant method is through the use of jumbo frames.

A jumbo frame carries a much larger payload, typically about 9,000 bytes compared with a standard frame of 1,500 bytes. At six times the size, its overhead is just 0.55%, compared with 4% for standard frames—a significant amount considering that approximately 13,800 jumbo packets per second can be sent at maximum throughput over a 1 GigE link. The use of jumbo frames not only has a positive effect on throughput, but also on CPU usage because the CPU processes and copies each frame on arrival.

A Camera Link frame grabber assembles complete images via its FPGA and generates a single OS interrupt. GigE Vision, on the other hand, relies on the CPU to process each packet, with associated multiple interrupts per received image, and then forwards each image to the application. The CPU is not only processing the data stream, but also assembling images from received

network packets, which can add a significant burden if not managed appropriately.

A high CPU load will inevitably result in processing delays in the host PC, as well as a reduced image processing capability. The use of jumbo frames clearly diminishes the number of interrupts, lowering the CPU burden accordingly, and increasing the processing bandwidth. GigE Vision filter, stack and performance drivers can further reduce the burden. Such drivers, for example, transfer the larger video packets directly to the application software, freeing the CPU of a considerable processing load.

With the right products and proper implementation, GigE Vision can achieve exceptional performance, transmitting uncompressed video between cameras/sensors and computers or displays with low, consistent “glass-to-glass” latency that exceed requirements for medical imaging systems. *

— John Phillips
Senior Manager, Product Management
with Pleora Technologies