WHY YOUR NEXT PCI FRAME GRABBER DOESN'T NEED ON-BOARD MEMORY

Frame grabbers for demanding applications such as scientific and industrial imaging and machine vision need high-speed access to the host processor's memory (PC). They get the necessary bandwidth through the PCI (Peripheral Component Interconnect) Bus. Because the PCI Bus feeds data directly into the host's memory, the frame grabber no longer needs the large memory buffers typical of older boards designed for the considerably slower ISA (Industry Standard Architecture) Bus.

Despite the fact that the PCI Bus is built for high-speed access and streaming data directly to the system's RAM, some manufacturers claim frame grabbers still require on-board memory to prevent data corruption or loss. A frame grabber without on-board memory, however, does not hurt PCI Bus or PC performance, nor does it leave you with gaps or corrupted data. In fact, having an on-board memory buffer does not guarantee your data will get through to the host's memory any better than with a board lacking memory. Extra memory is hardly the insurance policy some board manufacturers say, and it can lead to a false sense of security if you push the PCI Bus beyond its limits.

Here we will explain the memory alternatives for frame grabbers, how memory issues relate to your application, and ways you can assure your application runs at peak performance.

Bus Architecture: PCI vs. ISA

The ISA Bus, used in older systems and as a peripheral on the PCI Bus, has transfer rates around 3 MB/s. This is below the 9.2 MB/s (30 frames/s) rate of live video data traveling from the camera to produce 640 x 480 images (or the 11.1 MB/s (25 frames/s) rate for 768 x 576 images for the CCIR European standard). As a result, frame grabbers for the ISA bus rely on on-board memory to hold the image while they feed it to the host's memory at a rate the bus can handle. Without on-board memory, the data would overflow.

The 32-bit PCI Bus has replaced the ISA Bus on Pentium® and other high-performance processor boards and high-speed peripherals. Its 132 MB/s (theoretical) 95 MB/s (typical) burst rate is considerably faster than the required 9.2 MB/s, eliminating the need for frame grabber on-board memory. Data doesn't overflow because the bus stays ahead of it.

Frame Grabber Memory Alternatives: FIFOs vs On-board Memory

All PCI frame grabbers have some memory. The most popular kind is the FIFO (first in, first out) memory that holds one horizontal line of video. Data Translation's PCI frame grabbers use FIFOs, which serve as "way stations" for holding data while the frame grabber gets permission to send it on to the CPU's memory. With FIFOs, data streaming (or bursting) has this one-line delay to allow for any latency, or sluggishness, on the part of the host in giving the board the go-ahead to transfer data to the host's memory.

Frame grabbers with on-board RAM rather than FIFOs store the entire image before sending it to the host memory. In this way, the on-board memory actually delays the transfer so that it is no longer in real time.

What happens if, because of overloading or other systemic problems, the PCI Bus fails to accept data from the frame grabber? Data Translation's FIFO-based frame grabbers will not lose data, and will not spontaneously corrupt files. Instead, the board's driver will generate a software dialog box alerting you to the error. While an on-board memory buffer might store data for a short while longer, it cannot assure you that the PCI problem will correct itself before its memory overflows. So if the PCI Bus fails, having on-board memory will not ensure that data keeps moving to the host's memory.

Handling Application Demands

The PCI Bus can handle a burst of 132 MB/s, but usually apportions its bandwidth between the peripherals it is servicing. Imaging and machine vision applications can share the bus with other peripherals because they do not demand continuous bursting of video. They generally work using individual "snapshots" of data that go to the host's memory for analysis or processing.

For instance, a machine vision system might use video to inspect whether a part meets its specifications or has the correct serial number. This is unlike dedicated video systems, which rely on some kind of compression algorithm to reduce data volume and the bandwidth burden on the computer system.

But how can the frame grabber be sure to get the bandwidth it needs? The bus master architecture gives the frame grabber the upper hand, conserving CPU power and assuring data gets where it is headed.

Bus Master for Conserving Resources

The PCI Bus works with both master and slave devices. Bus masters take control of the bus to make sure data gets through smoothly. Bus slave devices take instructions from the host CPU, which can become preoccupied acting as "traffic cop" for the data transfers. Bus masters are preferable because they can take possession of the memory they need without CPU intervention. A PCI interface takes over the CPU monitoring and coordination tasks, freeing the CPU for all the heavy-duty processing needed in data acquisition and imaging. For this reason, Data Translation's PCI frame grabbers are all bus masters, with automatic data transfer that bypasses the CPU.

Manufacturers of boards with on-board memory claim that frame grabbers without it drain system resources and use excess CPU cycles. This could occur with a bus slave design, but Data Translation's bus-mastering frame grabbers bypass the CPU. They are able to use memory segments ready mapped, enabling the CPU to continue processing.

The Core Problem: Bus Overload

If used according to its specifications, the PCI Bus will transfer data reliably to the host's memory. However, piling too much onto the bus – such as control cards, video cards, networking cards and more – can overload it.

It's much like blowing a fuse. For instance, if you keep adding appliances to an outlet, sooner or later it will overload the circuit and all appliances will go down. This is not due to any inherent problem with the electrical service, but is caused by pushing the service beyond what it is designed to handle.

Similarly, when a PCI Bus gets overloaded, it can no longer operate to its specifications. The frame grabber, and any other peripheral on the bus, can get cut off from the host's memory. An on-board memory can store and hold an image for some period of time, but may also be left without any access to the host's memory.

Weighing Cost vs. Benefits

Some manufacturers market on-board memory as a security measure against losing data when the PCI Bus gets congested. However, with proper usage, the PCI Bus should not present any data loss with even rigorous applications, making on-board memory very rarely of value. Since the choice of a frame grabber must take into consideration the cost of each feature and its benefits, it is difficult to justify the additional expense of on-board memory.

In reality, on-board memory is an insignificant feature when determining what board to purchase. Other factors – such as accuracy, software features, and manufacturer support – far outweigh the usefulness of on-board memory. In addition, some manufacturers imply that on-board memory is a kind of panacea for PC system problems, helping it run at peak performance. In fact, it has no impact on the ability of the host's memory to analyze and process data.

OEMs are particularly sensitive to costs of their product components, and frame grabbers can be a significant portion of their cost of goods. Fortunately, OEMs producing dedicated instruments or devices can control how the PCI Bus gets used and can design against overload. In that way, they can assure uninterrupted data flow – without concern for on-board memory.

Maintaining Image Integrity

Some manufacturers also claim that without on-board memory, a frame grabber can cause "hiccupping" or staggered images. Again, this is false. These problems are not a function of memory, but rather of the display software. Frame grabbers with software using the Microsoft® DirectDraw (DDI) standard protect image integrity. This set of calls within Windows® allows for smooth, real-time video display with features like overlays. These low-level calls use few system resources and keep the host's memory busy with processing rather than with display. If the PCI Bus starts to get overloaded or sluggish, the frame grabber will trigger a warning dialog box to alert you that the bus is running at near capacity. But the display will not degrade.

Avoiding Congestion

A straightforward way to keep your imaging or machine vision system running uninterrupted is to keep the PCI Bus within its specifications. Stringing several high-speed devices onto the bus reduces its performance and leads to system overload.

Another method of avoiding congestion is the use of a motherboard with an Accelerated Graphics Port (AGP). The AGP is a dedicated high-speed bus connecting the PC's graphics controller and system memory, freeing up the PCI Bus for other functions. It offloads video traffic and handles bandwidth-intensive 3D graphics so they no longer have to compete on the PCI Bus. This leaves your system more stable and less prone to congestion. Available on the Pentium® II (and higher) processor, AGP is a more reliable and economical means of preventing bus congestion than the use of on-board frame grabber memory.

Conclusion

Even though the PCI Bus is built to handle the demands of high date-rate applications, it can get overloaded. Some manufacturers claim that lack of memory on-board the frame grabber can cause lost data and bus congestion. Frame grabbers with FIFOs, however, perform without data degradation or loss. Adding on-board memory can add costs without guaranteeing the data gets through in the event of bus congestion. Keeping the bus within its specification, and using processors with AGP, are more dependable ways of keeping applications running smoothly. The choice of a frame grabber is better determined by the accuracy, features, and performance of the board than by whether or not it has on-board memory.