System Requirements for DMA interfaces

Overview
EDT DMA boards are high-performance, high-speed direct memory access (DMA) devices, and the devices they connect to range from low to very high speed/high bandwidth. As such, they can demand a little or a lot from the host computer in terms of memory speed, display performance, and bus bandwidth.

Due to the wide range of devices and applications, required bus performance can vary greatly. In many cases, a modern Intel- or AMD-based computer that meets the basic system requirements will be more than adequate. But frequently, the device being connected to and the demands of the application dictate a system that's tuned to the task at hand.

PCI Express bus considerations
PCI Express was designed to replace the older PCI / PCI-X specification, which had less speed and bandwidth. The following list represents peak bandwidth guidelines for EDT's PCI Express boards:

- Gen.1 PCIe 8-lane ~ 1.3GB/s
- Gen.2 PCIe 8-lane ~ 2.6GB/s
- Gen.3 PCIe 8-lane ~ 4.0GB/s

EDT PCIe boards will typically fit in any slot with the same or a higher number of lanes than the board has. For example a 1-lane (x1) EDT board will also work in an x4, x8 or x16 slot, and an x4 board will also work in an x8 or x16 slot.

Note that PCIe slots can have different electrical and physical designations — for example there are 8-lane physical / 4-lane electrical slots. An 8-lane board will work in such a slot, but the maximum bandwidth (speed) will be reduced by approximately 50%. The pciediag_tool utility provided with the EDT PDV or PCD installation package, can be used to find out how many PCI Express lanes your board negotiated with the system.

PCI bus considerations
Note: The following discussion was written when PCI was the standard, and PCI-66 and higher slots were present in many computer systems. Since PCI Express now has replaced PCI as the standard, legacy PCI slots (if present) are implemented via PCI Express bridge chips, which only implement PCI-33, not PCI-66 or PCI-X. Therefore, an EDT 66MHz PCI board installed in a PCI Express computer will run at only half of its rated bandwidth. The following discussion therefore pertains mainly to older systems using PCI- and PCI-X.

Parhelia B.V.
Lylantse Plein 1
2908 LH Capelle aan den IJssel
The Netherlands

Phone: +31 (0) 10 741 00 28
www.parheliabv.com

Chamber of Commerce 24366010

PCI (& PCI-X) provides far less bandwidth than PCI Express; nevertheless it is still sufficient for a large class of devices. EDT PCI boards are capable of streaming data to host memory at device speeds of up to ~100MB/sec (33Mhz boards), or 200+ MB/sec (66 MHz boards in a 66Mhz or faster PCI slot).

When connecting to a device that reads or writes data at speeds of around 80 MB/sec or less, a single standard PCI bus system is generally adequate, assuming there are no other bus "hogs", such as SCSI controllers or PCI video boards sharing the bus with the EDT board. But when approaching bandwidth limits, particularly in conjunction with other PCI I/O devices, the EDT board can saturate the host computer's PCI bus. When this happens, data will be dropped from the EDT's DMA stream, and timeouts (data gets lost and the read call times out after waiting for a predetermined period for all of the requested data to arrive) will occur. For these reasons, EDT recommends any PCI bus systems have least one 66Mhz or faster PCI or PCI-X bus, and preferably multiple busses.

Higher-speed, dual PCI bus systems are not generally a feature in desktop computers, so it's usually necessary to go with a workstation or server that has some 66Mhz or faster slots. Unfortunately, determining whether a given computer even has dual PCI busses is not always straightforward. But with many systems it is possible to infer indirectly. If the manufacturer's spec sheet does NOT specify the clock speed or number of busses (not slots) in the computer, then it's a pretty safe bet that all PCI slots are on a single, standard 33Mhz bus. Systems that have different speed slots — say, one 33Mhz conventional PCI slot and one or more 100 MHz PCI-X slots are guaranteed to have separate PCI busses, since all slots on a given bus need to run at the same speed. It is because of this dual bus "side effect" of PCI + PCI-X computers that EDT recommends a system with at least one high-speed PCI or PCI-X bus for high-demand applications, even if the EDT board is a 33MHz device. If PCI bus bandwidth is an issue due to bus contention with another PCI bus "hog", then a system with multiple busses will likely solve the problem since it allows separation of the EDT interface's PCI data stream from other PCI I/O data traffic. Putting two or more EDT PCI boards in the same system will result in the same bandwidth sharing issues, so the dual-bus solution would apply in that case as well.

EDT PCI boards designed before 2001 run at 33Mhz. All newer boards are 66Mhz or faster. See the specifications on the data sheet for your product to find out its bus clock speed. 66Mhz boards will run in 33Mhz slots but will only achieve their optimal performance in a 66Mhz or faster PCI or PCI-X slot. Note that the presence of a 33Mhz board in another slot on the same bus will throttle the whole bus back to 33Mhz.

Calculating bandwidth

To determine bus bandwidth, you should calculate the total data rate from the device or devices. Be sure to count bytes, not data words or pixels. For example a CCD camera with a 80MHz pixel clock, 10 bits per pixel and 2 taps (ports) is actually sending out 4 bytes of data for every clock cycle, or 80 x 2 x 2 = 320 megabytes per second. An EDT x4 or x8 PCI Express board (in a 4, 8 or 16-lane slot) would have bandwidth to spare, but not a PCI board (see PCI bus considerations, above)

Things get tougher when you have other significant bus activity. For example, many imaging applications require simultaneous saving of images to a SCSI Raid disk. If the SCSI interface is on the same bus as the EDT board, you may saturate the bus. Using a system with 2 PCI buses, and isolating the EDT board on its own bus is the typical solution in this situation.

Interrupts, Frame Rates and Buffer Rates

Another limitation that sometimes shows up relates to OS system interrupts. Each buffer (or frame, in the case of a camera) generates an interrupt to tell the software when the buffer is full, and the OS's interrupt scheduler needs to be able keep up. If it doesn't, some buffers (frames) may be missed. The achievable interrupt rate is highly dependent on OS as well as other interrupt traffic (mouse movement interrupts for example) but maximum reliable rates for a single EDT board/device are typically in the 1-2KHz range. If your device / application exceeds this rate, there are various workarounds, for example by grabbing two images/frames per buffer in the case of a digital video board, or streaming bytes and setting the board to ignore any gating signals that would cause interrupts, e.g. frame valid / line valid. Look at our example applications that come with the driver / SDK or contact EDT if you need help with this.