

Serial Busses Take Over

Low Voltage Differential Signaling (LVDS) connections have been gaining in popularity over the last few years. They address many of the performance restrictions inherent in traditional parallel bus designs.

FOREWORD

The increasing speeds of processors and I/O devices demand corresponding increases in bus speeds to tie these devices together. Unfortunately, traditional parallel busses are reaching their performance limits. PCI-X will achieve a burst bandwidth (not sustained) of perhaps 1 GB/s, but this is likely the practical limit for a traditional style parallel bus. Signal noise, cross talk, signal skew, high power consumption and cost combine to create this limit. The only ways to make parallel busses faster are to increase the clock rate or make the bus wider. A faster clock only makes the cross talk and power consumption worse, and making the bus wider increases the cost, power consumption and skew problems. Creating a faster interconnect requires a different approach - a serial interconnect using LVDS.

Low Voltage Differential Signaling (LVDS) connections have been gaining in popularity over the last few years. They address many of the performance restrictions inherent in traditional parallel bus designs.

WHAT IS LVDS?

As the name states, LVDS involves low voltages. Voltage swings are only about 350 mV, compared to the 2.4 V swings required for TTL. This certainly reduces power consumption. But, even more important, the edge rates are dramatically slower because of the reduced voltage swing. The result is the potential to clock data significantly faster, as long as noise and cross talk can be minimized. The differential signaling is the magic for controlling noise.

By using two wires with opposite current/voltage swings instead of just one as is done with conventional busses, common mode noise can easily be rejected. Common mode noise would appear equally on both wires. Because the receiver looks only at the difference in voltage levels, not their absolute values, this noise is rejected. In addition, less noise is radiated by these pairs since the opposite voltage levels tends to cancel out radiated emissions. Another benefit is that while a fast parallel bus may be limited to a few inches in length because of skew problems, a serial interconnect can go a couple of feet across a backplane and through backplane connectors, or even several meters through an appropriate cable.

LVDS does require two wires per signal compared to only one with traditional busses. However, because LVDS can go so much faster, fewer wires are actually required for a given bandwidth. In real applications, bandwidth can be tripled while reducing the number of wires by a factor of 3 compared to the technology

LVDS is replacing. This translates into a performance boost while simultaneously saving cost in component size, PCB layout and manufacturing costs, and connector sizes.

The most common use for LVDS is in point-to-point connections. One of the early, popular uses was in driving the LCD display on laptop computers. Low cost, low power, high data rates, and the ability to use a simple, small flat cable without radiated emissions were all factors in this choice. Even higher performance and more general purpose applications are now starting to appear. LVDS support chips have been available for a few years now to support semi-custom applications. Now there are two major efforts to provide "standard" LVDS busses. The first was InfiniBand. A more recent example of an LVDS bus is RapidIO.

RapidIO

RapidIO is an example of how LVDS can be used to implement a high-performance parallel bus. RapidIO is a proposed specification that came out of work done by Mercury Computer Systems and Motorola. The RapidIO Trade Association has recently been formed with representatives from several companies to finalize the specification, and to promote it. RapidIO is intended as a low cost, high performance processor interconnect.

A RapidIO link consists of 8 or 16 data signals, a clock signal and a frame signal. These signals are duplicated - one set for transmit, one set for receive, for a total of 20 signals for the 8-bit interface. Because LVDS is used, each signal requires 2 wires. As a result, the minimum 8-bit interface requires 40 wires. This is comparable to some older high-performance interconnects like SKYchannel and RACEway, but better than other busses like PCI and VME. While the wire count is comparable to these older busses, the clock rate is significantly higher.

Encoded onto the link is both data and control information. The frame signal is used to distinguish between the two types of information. The clock runs continuously, so there is always information on the bus, even if the information is only IDLE characters. However, control information can indicate whether there is room at the receiving end for additional data, provide an acknowledge for successful data transmission, request retries on errors, and much more.

The minimum RapidIO implementation is a single connection, perhaps between 2 processors. The more interesting case is where it is used to connect a large number of devices. For example, a system with 256 processors is possible. Crossbar switches are used to

maintain the point-to-point nature of LVDS and simultaneously scale up to these larger systems.

RapidIO is a packet switch architecture. Transmitting a large buffer requires breaking the buffer up into smaller pieces called packets that can be managed efficiently by the RapidIO fabric. Each packet contains a header with packet type, source and destination. At the end of each packet is an ECC word for reliable data delivery.

RapidIO is simple enough that it can be implemented in a variety of available technologies, from FPGA to full custom. Using available FPGA technology, clock rates of up to 622 MHz are possible, producing a full-duplex data rate of 622 MB/s for an 8-bit implementation, or double that for a 16-bit implementation. This compares favorably with the half-duplex 1GB/s speed of PCI-X, and RapidIO has significantly better connectivity. Speeds of 2.5 GHz are possible using full custom CMOS.

INFINIBAND

InfiniBand combines LVDS technology with capabilities developed for telecommunications to provide high performance over a small number of wires. The minimum implementation of InfiniBand uses only four wires - two for transmit and two for receive. The signal is clocked at 2.5 GHz. In order to provide a useful connection, a mechanism must be provided to deliver data, clock and control over a single differential pair. IBM's 8B/10B encoding is used to provide this.

8B/10B encoding was originally developed to support fiber optic data transmission. There are a number of serial transmission issues addressed by this encoding.

The first problem is maintaining phase locked loop (PLL) clock synchronization on the receiving end. Without encoding, a string of zeroes or ones wouldn't have any signal transitions, and the receiving PLL clock would drift. By taking each 8-bit data element and encoding it into a 10-bit data frame, enough signal transitions can be guaranteed to make sure that a PLL clock can maintain synchronization with the data stream.

The next problem is transmitting control information as well as data. Having 10-bit characters allows a significant number of control characters, enabling data synchronization and control over the same signal as the data.

It is also possible to implement some basic error detection because of the rules involved in implementing 8B/10B.

The final advantage of 8B/10B encoding is that it has good DC balance. The low frequency content of the data stream is low. There is a limit of 5 bits in 8B/10B before a signal state transition will occur. This isn't very important for LVDS implementations, but it does make it possible to use low-cost fiber optic lasers if longer distances are needed. Using fiber optics is simply a matter of adding a laser to the LVDS pair at the sending end, and a photo diode to the LVDS pair at the receiving end.

By using 8B/10B encoding, InfiniBand is able to

achieve a full-duplex 250 Mbyte/s bandwidth over 4 wires clocked at 2.5 GHz. InfiniBand also defines a x4 and a x12 link, using 4 or 12 of the basic connections in parallel to increase the link bandwidth. A x4 link achieves 1Gbyte/s using only 16 wires, and a x12 link gets 3Gbyte/s of full-duplex bandwidth using 48 wires.

InfiniBand is also a packet switch architecture. It was originally targeted at providing a connection between servers, disks and the Internet in a computer room. As a result, cables - both copper and fiber-optic - have been defined for connections. Recently, a backplane and packaging specification have also been introduced, making it attractive for use within a system.

As with RapidIO, the InfiniBand specification defines the packet formats, as well as the control information that must be sent over the links. In addition, InfiniBand has built on the work that went into VIA to provide secure message passing with minimal processor overhead. In addition, extensive work has gone into defining reliable data transfer, live insertion, "discovery" so that devices can be detected and configured when inserted, and redundant data paths in the fabric.

SEMI-CUSTOM

InfiniBand chips are starting to be announced - IBM was the first - with product availability next year. RapidIO is farther into the future. You don't have to wait for these technologies to start using LVDS. Several companies, including National, TI, AMCC and Vitesse offer SERDES (Serializer/de-serializer) chips that will convert parallel data into serial connections. Crossbar chips are also available to scale the interconnect size up if required. As is done with InfiniBand, 8B/10B is used to encode user data with the control information required by the parts over the serial links. The initial market for these devices is for switches supporting GigaBit Ethernet, Fibre Channel and other protocols, but they are general purpose enough that they can be used in a wide range of applications.

LVDS technology is enabling a dramatic improvement in interconnect bandwidth. It is in a rapid ramp-up in popularity and availability. The support chips exist today to enable application development, and it will be incorporated into even more chips in the future. Standards development will make it economically viable to include LVDS in processors and standard I/O devices. Traditional parallel busses will disappear from use as the interconnect between processors and I/O devices ■

Stephen Paavola is a product manager at SKY Computers, Inc. of Chelmsford, MA. He has over 25 years experience in senior engineering, sales and marketing roles at Digital Equipment Corporation, Cetia, and SKY. Stephen Paavola holds a Bachelor of Science from Cal Tech in Pasadena, California.