Sorting Out Backplane Driver Alphabet Soup

Introduction

Our expectations of what technology can deliver are changing at an ever-increasing rate. Consumers today expect flawless performance, higher bandwidth, and better form factor from the systems they interface with—regardless of whether the system in question is a personal notebook computer or a worldwide communications network. These expectations are forcing many designers to reconsider all areas of their system designs, including the backplanes. While new, high-end, application-specific devices are being introduced every day to help in this matter, upgrading the backplane remains a far-from-simple task.

To date, the semiconductor industry has offered a confusing array of alternative technologies, in addition to the inexpensive, garden-variety, 245 Octal bus transceivers. Each of these technologies is accompanied by an abstract multi-letter acronym, such as ABT, BTL, GTL, BLVDS (see Table 1). Although there are significant differences between these technologies, there exist large areas of overlap in their application. Choosing the right technology is a matter of custom fitting the capabilities of the driver to the specific needs of the system design. Loading, power, bus configuration, and speed requirements must be weighed to meet today's needs and also provide a long-term upgrade path for the system.

The engineer's definition of a backplane can be divided into two general perceptions. The first is that of a generic pc board with multiple connectors. The second is that of a special-purpose, transmission-line network that provides high data-signal integrity, control of potentially damaging crosstalk, minimization of radiated emission (EMI), distribution of data, and even distribution of power and ground references to all cards in the backplane. Both backplane definitions are correct; which one applies to your design depends on the length of the backplane, the data rate, and the signal characteristics of the selected backplane-driver technology. The most common mistake when analyzing the backplane and its cards is to consider them only as a lumped-capacitance load. Because there are two types of backplane models, a determination of which model applies to the application must be made.

The first model applies to backplanes that typically run at a lower speed, such as 5 MHz or less. At these lower speeds, the bit width is relatively long and the edge rates (transition times or rise or fall times) are slow. The key here is the rise time, as this is the main parameter that determines if a lumped-load or a transmission-line model should be used. If the edge rates are slow, then the cards inserted into the backplane may be treated as one lumped load, since the transmission-line effects (reflections) that occur will die out in a short period of time compared to the signal's pulse width (unit interval). This gives sufficient time for the signal to settle out into a stable state before sampling occurs.

A general guideline is to compare the unit interval to six flight times (or to three round trips). A flight time is the electrical length of the backplane; in other words, it is the time it takes the signal to travel from one end of the backplane to the other. A round trip is simply two flight times. Six flight times should be less than 30% of the unit interval to generate a stable state at the 50% point.

Another way to avoid transmission-line problems is through the use of specially designed trapezoidal drivers such as the DS3862 Octal bus transceiver. These drivers feature slow edge rates which are greater than the electrical length of the backplane; thus, the backplane can again be modeled as a lumped load.

A common TTL backplane driver (F245 Octal bus transceiver) driving a 21-slot unloaded backplane is shown (*Figure 1*). Transmission-line problems are evident. The waveforms show overshoot, undershoot, and reflections. However, at a relatively low speed (1 MHz), these problems may be ignored because the unit interval is very large and settle out relatively quickly, but can cause other system issues such as EMI (because of overshoot and undershoot, and ringing). As the unit interval is decreased (10 MHz), the width of valid sample area also is decreased. Now the transmission-line effects take up a significant portion of the unit interval.


FIGURE 1. A common TTL backplane driver (F245 Octal bus transceiver) driving a 21-slot unloaded backplane shows transmission-line problems such as overshoot, undershoot, and reflections. At a relatively low speed (1 MHz, for example) these problems may be ignored, though they can cause other system issues such as EMI.

"Reprinted with written permission from Electronic Design, (Vol. 46, No. 25) Nov. 2, 1998 Copyright 1998 Penton Publishing".

The second model is for higher-speed applications (>33 MHz) in which the backplane must be treated as a transmission line. If incident-wave switching is desired and the round-trip delay is greater than the edge rate (rise time) of the signal, then you have a transmission line. The simple, lumped-capacitance model no longer applies, and now a distributed model must be used. Incident-wave switching is generally desired at higher data rates.

Such switching requires a clean signal environment to allow the receivers to properly detect the correct state as the signal travels down the backplane. There is not enough time to wait for reflections to step up the voltage, as the pulse widths are very short. For these reasons, a properly terminated bus is very important, as it will prevent the generation of undesired reflections. If you are not sure which model to apply, treat the system as a transmission line.

The first step in working with the high-speed model is to calculate the true bus (loaded backplane) impedance. The bus impedance is not the generic pc-board impedance. The generic impedance is a function of the stripline layout dimensions and the dielectric insulation and is typically in the range of 40Ω to 60Ω. This impedance can be verified with a TDR when all cards are removed, but the real system has cards installed that add capacitance, and this alters the impedance of the backplane. The net result is that the fully loaded impedance will be lower than the "unloaded" impedance. This loaded characteristic impedance of the backplane, termed Z_L, is calculated by the following equation:

$$Z_L = Z_0 / \sqrt{1 + C_L/C_0}$$

where:

$$Z_0 = \sqrt{(L_0/C_0)}$$

Z_L = Impedance of loaded line.

Z_O = Bus characteristic impedance (known).

L_O = Distributed intrinsic inductance per unit length.

C_O = Distributed intrinsic capacitance per unit length (known).

C_L = Distributed load capacitance per unit length (known; includes capacitance of the cards, connectors, vias and output capacitance of the chosen backplane IC).

Knowing the correct loaded impedance is very important, as a mismatched backplane will have major signal-integrity issues, such as negative reflections and undershoot, which may prevent incident-wave switching or overshoot. It also can cause ringing that may lead to EMI problems. These can not be tolerated due to the high speed and small unit intervals required.

The distributed capacitance not only affects the impedance, but it also affects the line propagation delay. The delay is calculated by the equation:

$$t_{PL} = t_{PO} \sqrt{(1 + C_L/C_0)}$$

where:

t_{PO} = Unloaded line delay (propagation delay).

C_O = Distributed intrinsic capacitance per unit length.

C_L = Distributed load capacitance per unit length.

t_{PL} = Loaded line delay.

In addition, a BTL signal was monitored as more cards were added to the backplane (*Figure 2*). The delays added to the line are shown to be considerable, and must be added to the overall evaluation of the backplane and used in system timing equations.

To examine this phenomenon, a 21-slot backplane was used. Two of the slots were occupied by termination cards (V_t) which were located at slots 1 and 21 (the two ends of the backplane). A BTL driver was installed in slot 2; no other cards were installed, and the BTL signal was monitored at slots 2 and 19. The observed delay is the time for the signal to travel from slot 2 to slot 19. We then populated the even-numbered slots (4, 6, 8,...18) and monitored the signal at slot number 19. The result shows that by installing more cards (capacitive load) into the backplane, the flight time for the same signal increases by approximately 2 ns. The fourth waveform is our signal monitored at slot 19 in the backplane with all 19 slots populated.

We observed two points from this evaluation: first, the flight time increased as we added cards to the backplane; second, when the backplane was fully loaded, the existing terminations were matched to the loaded backplane impedance. A better termination match was made and the amount of overshoot was far less.

It is important to remember that as the loading is increased (capacitive load), the resulting f_{MAX} (maximum switching speed) is decreased and the propagation time of the signal (flight time) is increased. This is the challenge presented to the bus driving technologies and the system designer.

FIGURE 2. A BTL signal was monitored as more cards were added to a 21-slot backplane and two observations were made. First, the flight time increased as cards were added to the backplane. Second, when the backplane was fully loaded the existing terminations were matched to the loaded backplane impedance. Thus, a better termination match was made and the amount of overshoot was far less.

TABLE 1. Backplane Driver Alphabet Soup

| Acronym	Technology
ABT	Advanced biCMOS Technology
BTL	Backplane Transceiver Logic
ABTE	Advanced biCMOS Technology Enhanced
Bus LVDS (BLVDS)	Bus Low-Voltage Differential Signaling
CBTL	CMOS Backplane Transceiver Logic
CML	Current Mode Logic
CMOS	Complementary Metal-Oxide Semiconductor
ECL	Emitter Coupled Logic
ETL	Enhanced Transceiver Logic
FAST	Fairchild Advanced Schottky TTL
FACT	Fairchild Advanced CMOS Technology
GTL	Gunning Transceiver Logic
LCX/LVC	Low-Voltage CMOS
LVT	Low-Voltage Technology
LVDS	Low-Voltage Differential Signaling
TTL	Transistor-Transistor Logic
PECL	Positive Emitter-Coupled Logic

Reflections and Terminations

Reflections are caused by mismatched impedance (changes of impedance along the line), which may occur as the result of stubs, board layer changes, or incorrect termination values. If the backplane has been determined to be a transmission line, the use of terminations is typically required. When the signal travels down the backplane and encounters a matched termination $R_T = Z_L$, no reflections occur. This is the best case for signal quality. If the termination is not matched to the backplane's loaded impedance, reflections will occur and degrade signal quality.

Fully loaded, single-ended backplanes typically have a Z_L value of around 30Ω to 45Ω as their ac (loaded) impedance. Again, for the matched case, the termination resistor value is equal to Z_L, and the resistors are located at both ends of the backplane.

The effect of terminating a BTL signal on one end only is shown (*Figure 3a*). This termination method prevents the signal from operating as an incident edge. The reflections from the unterminated end cause a step on the waveform, causing signal-integrity problems such as signal delay, mis-triggering, or double triggering by a clock. Since the waveform is not operating as an incident edge and must wait for reflections to fully change state, it adds to the system's overall delay time.

The step caused by signal reflection occurs within the threshold region; therefore, a clock may trigger at more than one point on a single edge. The waveform of a BTL signal terminated correctly on both sides is also shown (*Figure 3b*). The termination value is matched with the loaded backplane's impedance, and thus does not allow reflections; therefore, no additional time delay or false triggering occurs. Note that the dc level is reduced with two terminations, this is due to the driver operating with a larger sink current (I_{OL}).

10106303

FIGURE 3. Terminating a BTL signal on one end prevents the signal from operating as an incident edge. The reflections from the unterminated end produce a step on the waveform, causing signal-integrity problems such as signal delay, mis-triggering, or double triggering by a clock (a). A signal terminated correctly on both sides does not exhibit this step, which eliminates the problem (b).

When low-impedance (30Ω) termination resistors have been selected, we also must consider the resulting amount of load current (I_{OL}) that the backplane ICs will have to sink (per channel). The worst case is when driving a backplane with 30Ω loaded impedance and a signal with a 3V swing (for example with a F244 device). This will require 200 mA of I_{OL}, but if the signal swing is reduced to 1V (BTL), the current required decreases to only 67 mA.

Most TTL backplane ICs do not have the required current drive to pull the signal out of their wide threshold area; therefore, they must rely on reflections to change state. This is not the case with BTL, as it supports an 80 mA sink capability, thus providing incident-wave switching. The loaded bus impedance is the reason backplane drivers are required to sink large currents. This must be done while maintaining their V_{OL} ratings to maintain noise margins.

Backplane Driver Technologies

Having determined which backplane model applies, we can start to narrow down the selection of upgrade technologies. In addition to drive capability, other factors such as noise margin, backward compatibility, bus configuration, and live insertion should be considered.

Technologies for driving backplanes have developed along two paths. The first path we will examine improves the overall performance of the system while maintaining standard TTL signaling levels. This may be desired when it has been determined that the chassis (backplane) of a deployed system must remain in service due to capital costs, remaining life, or other infrastructure reasons.

Therefore, of great concern to many system designers is the ability to maintain backward compatibility with previous equipment. This simplifies many facets of system support, especially inventory and maintenance. By maintaining a compatible backplane interface it is possible to allow customers to add new boards to existing systems.

However, mixing boards with higher performance drivers that are compatible with the existing technology into a system will not improve the overall system performance. An example of this is installing a card with FCT drivers into an existing backplane using standard TTL drivers. Only by converting the entire backplane to a new technology will the end user see performance improvements in the system.

These improvements are achieved by decreasing the width of the signal's threshold region and by improving the skew specification. Standard TTL signaling maintains a signal swing of approximately 0.55V to 2.4V and a wide threshold region from 0.8V to 2.0V. Most CMOS devices have a wider signal swing, but still maintain the same threshold levels, thereby providing additional noise margin. The 74FCT245T/AT is a good example of this category of device.

Advanced biCMOS technology, or ABT, was introduced to improve on this type of device. While ABT maintains the same threshold region as TTL logic, it specifies a much tighter channel-to-channel skew specification than most TTL/CMOS-compliant families. Typical skew between channels can be up to 5 ns on standard logic families, whereas ABT guarantees a much tighter 2 ns to 3.5 ns, depending upon the output load. ABT also offers a very fast propagation delay, with a maximum of 3.5 ns to 4.5 ns specified depending on the manufacturer.

All in all, this offers a significant performance boost over standard (older) logic families. ABT also reduces concerns over bus contention by offering a much shorter disable delay than enable delay. This means that devices sharing a common bus can be switched on-to-off and off-to-on immediately following the other, because the device with preceding control of the bus is guaranteed to be shut down prior to the new device being enabled. However, contention can also be timed out from system timing equations. Low-voltage technology (LVT) is the 3.3V version of ABT and offer similar features.

A recent enhancement to ABT has been introduced by Texas Instruments (TI). This technology is known both as ABT enhanced (ABTE) and enhanced transceiver logic (ETL). This technology was developed and specified by the VME64bus committee to provide extended performance to VMEbus systems. ETL improves the noise margin by reducing the width of the threshold range from the 1.2V of stan-

dard TTL down to a tight 200 mV. This has two immediate effects: the noise margin is dramatically increased, and the window for sampling a valid signal is increased while still maintaining compatibility with older TTL systems.

Noise margin is calculated by comparing that portion of the total swing in which the signal is in a known state (outside the threshold region) to the total signal swing. By maintaining the same overall signal swing (0.5V for low, 2.4V for high) but reducing the threshold region to 200 mV, ETL raises the overall noise margin to 90%, compared to 35% for a standard TTL family. As a result, signal integrity will be maintained even if the environment is significantly more noisy. The magnitude of the noise must be much greater to push a valid signal into the threshold region in an ETL system.

The other path to upgrading backplane technology involves a more dramatic change in signal levels. While this precludes backward compatibility with TTL systems, it does offer great advancements in performance. Backplane transceiver logic (BTL), Gunning transceiver logic (GTL), emitter-coupled logic (ECL), low-voltage differential signaling (LVDS), and Bus LVDS (BLVDS) fall into this category, with each offering unique capabilities to address specific application issues. However, all trend towards narrower thresholds and reduced signal swings.

The first two technologies are single-ended, which means the logic state is indicated by the signal voltage referenced to ground alone. The last three are differential data transmission technologies (see appendix *"Single-Ended vs. Differential Transmission,"* . Using two active signal lines, the logic state is the differential voltage between the two. This improves the noise margin by a multiple of the signal swing. Recall that single-ended noise margins are only a fraction of the signal swing.

BTL was invented by National Semiconductor in 1984 in support of the initial Futurebus protocol specifications. Although Futurebus and its enhancement Futurebus+ have yet to attain significant market share, the underlying physical-layer technology, BTL, has enjoyed significant market success. High bandwidth in a heavily loaded environment became a key requirement for enterprise LAN hubs and large telecommunications systems. BTL, used in conjunction with proprietary protocols, offered performance and benefits similar to those of a full Futurebus implementation at a substantially lower cost. BTL is commonly used in 20 MHz to 66 MHz systems with as many as 20 cards.

The BTL signal structure offers a compressed signal swing—almost half that of TTL. The threshold region has been reduced as well, down to 150 mV. It also has a high drive capability (sink) of 80 mA and is an open-collector design. For the high level, a termination to 2.1V is required. What really makes BTL unique for backplane applications is its extremely low output capacitance—typically below 5 pF, whereas most TTL-compliant technologies are two to four times higher. This reduced loading has a significant impact on the maximum performance of the bus, and allows BTL-based systems to have either increased performance or significantly better timing margins than comparable ABT or ETL systems.

For this reason, BTL is often considered when a combination of high speed and a heavily distributed load of many boards are required. For dense parallel applications, National Semiconductor has recently announced CMOS BTL (CBTL). This

Backplane Driver Technologies

(Continued)

is a pure CMOS technology that provides the full feature set of BTL with reduced I_{CC} currents, thus easing system power-supply design and distribution.

Both BTL and ETL feature additional protection in the form of live-insertion circuitry. These devices offer designers the capability to pre-bias the driver before insertion into the backplane, thereby avoiding system glitching and damage to the inserted board. Live insertion, or hot swapping, is of particular importance to the telecommunications marketplace. In these applications it is critical that maintenance and repair be performed without shutting down the entire system, or causing disruption to the traffic on the backplane.

Many of these same features are offered by GTL. Invented in 1991 by Xerox, GTL further reduces the overall signal swing (0.4V to 1.2V) and threshold (100 mV). This technology was specifically created to address very fast chip-to-chip interfacing issues, such as those between microprocessors and memory devices. As such, the drive capability of GTL is specified at 40 mA, half that of BTL. Some devices on the market today have been specified to have 60 mA drive capability to boost performance.

Although originally intended for low-voltage, high-speed computer applications, GTL is finding acceptance in some areas of small backplane design. Due to the low drive capability (40 mA), it is not appropriate for a heavily loaded environment, but for systems that require high-speed (30 MHz to over 50 MHz) performance between fewer than a dozen cards, it is a very good fit.

The GTL specification has been modified slightly and reintroduced in some applications as GTL Plus, or GTLP. The difference here is that the entire signal swing has been widened slightly to 0.95V (0.55V to 1.5V) while maintaining the 100 mV threshold. This results in a slightly higher noise margin and pushes the threshold region slightly further away from ground—a consideration implemented to avoid potential ground-bounce issues. Both varieties of GTL are available on the market, and it is likely that the personal computing marketplace will determine the future of this standard.

In applications where extremely high backplane speed is required, ECL has been adopted. This technology can provide backplane clocking speeds beyond 100 MHz, far greater than single-ended bus-driving technologies. This performance requires trade-offs in terms of both power consumption and power-supply design. ECL achieves its high performance (>100 MHz) by making use of a low output impedance and a reduced threshold (120 mV).

Commonly operating at negative voltages, an ECL backplane places additional cost in the design and routing of the system power supply. Interfacing the rest of a system to the ECL backplane can also be a design issue; some ECL backplane designs require associated ECL logic in the rest of the system in order to maintain system throughput. Translator devices may be needed to interface between TTL and ECL such as the 100328 devices.

Positive ECL (PECL) offers yet another choice, this version of ECL supports positive-voltage, power-supply operation. Even with its power and complexity limitations, ECL is uniquely suited to meet the requirements of high-bandwidth systems. These are obtained by designing the drivers as Class A amplifiers operating in the linear region, thus providing fast balanced ac specifications, and an extremely low output impedance for high-speed data transmission.

LVDS is a high-speed (hundreds of megabits per second) differential data transmission technology that operates at very low power-dissipation levels from common power-supply rails (5V or 3.3V). Being differential, and supporting a ±1V common-mode range, LVDS provides about twice the noise margin of GTL or BTL. Termination is greatly simplified, as no active pull-up voltages are required (as in the case of BTL and GTL technology). A single surface-mount resistor is all that is required. LVDS drivers swing from 250 mV to 450 mV, centered around 1.25V, while the receivers support thresholds less than 100 mV. Standard LVDS drivers and receivers are commonly employed in point-to-point or multidrop (multiple receivers) applications, thus they can be used in switched-backplane applications, or on other special links across a backplane. LVDS has been standardized by the TIA (Telecommunications Industry Association) as a Electrical layer standard and published as TIA/EIA-644-A.

BLVDS also is a high-speed (hundreds of megabits per second) differential data transmission technology that extends the benefits of standard LVDS into multipoint bus and or other special configurations. The multipoint configuration supports bidirectional half-duplex bus communication. It differs from standard LVDS by providing a higher drive, which provides similar small-signal swings (about ±250-300 mV) while loaded with two terminations (one at both ends of the bus).

Since the signal swing is greatly reduced, fast transition times are possible, thus allowing the drivers to address high data rates ranging from hundreds of megabits/s to over 1 Gbps. The differential data transmission scheme provides a ±1V common-mode range and live insertion (hot plug) of devices into an active bus. Additionally, the low voltage swing minimizes power dissipation and noise generation (crosstalk and EMI). BLVDS greatly simplifies the area of bus termination as it does not require special active termination devices, nor does it require a unique termination rail (such as 2.1V for BTL) to be supplied. It simply requires a single surface-mount resistor across the pair at each end of the bus.

BLVDS also utilizes common power-supply rails (3.3V or 5V), minimizes power dissipation in the interface devices, generates little noise, supports live insertion of cards, and drives heavily loaded multipoint busses at hundreds of Megabits/s. BLVDS addresses many of the challenges faced in a high-speed bus design and products are available as simple transceiver devices, optimized parallel bus transceivers with ultra low skew, and 10-bit serializer/deserializer devices. The multipoint subcategory of BLVDS has been standardized by the TIA as M-LVDS and published as TIA/EIA-899.

CML (Current Mode Logic) is a high-speed, simple, point-to-point interface used in switched multi-gigabit backplanes where speed is the number one concern. A unique feature of CML is that it typically does not require any external resistors as termination is provided internally by both the driver and receiver devices. CML interfaces support line speeds in the 1 to 12 Gigabit per second range depending upon process and encoding. CML may be DC or AC coupled if encoding is used. CML uses a passive pull up to the supply rail which is typically 50 ohms in magnitude. CML I/O tend to be used on SERDES (Serializer / Deserializer) devices and are currently operating in the 1 to 5 Gbps range.

Fitting it All Together

Having examined the benefits and features of each technology, and applying transmission line theory, a comparison of backplane performance with respect to loading can be produced (*Figure 4*). This comparison is purely relative, as the capabilities of the common technologies incorporate a significant degree of overlap. Given enough design and debug time, each technology can be pushed beyond the limits described here. A comparison of levels and the resulting noise margins are shown in *Figures 5, 6* and *Figure 7*. *Figure 5* and *Figure 6* are common, single-ended technologies, while *Figure 7* is specific for LVDS, Bus LVDS and CML.

Note that the common-mode range in differential data transmission technologies is what should be compared to the standard noise margins of single-ended technologies. Therefore BLVDS and LVDS, with their 250 mV swings, both provide about twice the noise margin of GTL- or BTL-based

systems. For lower-speed systems, regardless of load conditions, a standard TTL family such as LCX or FACT may be used.

If a performance improvement is required, but backward compatibility is necessary, LVT or ABT or ETL may be considered. Very fast systems with a few boards (light loading) could find GTL a good design choice, provided live insertion is not required. For more heavily loaded systems running at high speed, BTL or even ECL/PECL may be required. If ultra-high performance is required, and ultra low power dissipation is a must, then Bus LVDS or CML is the driver technology of choice. In each case, regardless of the technology chosen, proper design rules should be followed to minimize reflections, crosstalk, and other transmission-line related issues. Upgrading driver technologies can help eliminate these problems, but no transceiver can mask a fundamentally poor design.

10106307

FIGURE 4. A comparison of backplane performance with respect to loading can only be relative, as the capabilities of the common technologies incorporate a signal degree of overlap. Given enough design and debug time, each technology can be pushed beyond the limits described here.

Fitting it All Together (Continued)

FIGURE 5. The noise margin for common 5V and 3.3V standard logic technologies may be calculated by subtracting V_{IH} from V_{OH} and V_{OL} from V_{IL}. Among these technologies, ETL/ABTE has a greatly decreased threshold region to increase the noise margin.
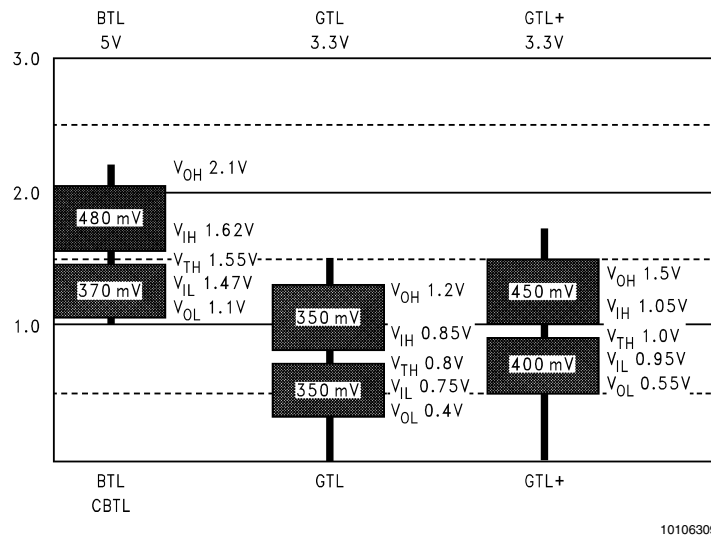

FIGURE 6. From a noise-margin/level comparison of reduced-swing technologies such as BTL, GTL and GTL+, it can be seen that although the output swings have been reduced the noise margins have actually been improved—relative to some full-swing logic technologies such as TTL, LVC, and ABT.

www.national.com

Fitting it All Together (Continued)

10106310

FIGURE 7. The active signal swings of LVDS, and Bus LVDS are reduced even further to about 1/4 of the reduced-swing technologies illustrated in *Figure 6* along with CML. However, because these are differential interfaces, the common-mode range should be compared to single-ended noise margins. The effective noise margin is two to four times better using LVDS.

Appendix

SINGLE ENDED VS. DIFFERENTIAL TRANSMISSION

In many applications, the use of differential transmission technologies is ruled out by myths alone. When the word differential is spoken, it immediately generates a vision of two pins per signal and gigantic buses. It is true that differential transmission uses two lines per signal, as the logic state is denoted by the difference voltage, whereas single-ended transmission relies on a voltage level and only one active signal line. However, in designing large single-ended backplane buses, the large return current must be taken into account. To provide a low-impedance path, it is common to assign many ground pins. The ratio of grounds to signals is application dependent but ranges on the low side commonly as 3:1 to as high as 1:1. If the 1:1 ratio is selected, then the "pins required" for a differential bus and the single-ended bus come close to par.

Differential can even beat the 3:1 applications, and reduce pins required even further. Since differential technologies use small swings to enable high-speed operation, faster signal paths are possible. Combining differential bus driving and a serializer/deserializer function reduces pin count to almost zero. Two pins for the serial signal, and a single common (GND) reference. This reduction in bus width requires less pc-board real estate, allows for smaller connectors and smaller interconnect media, and even eases the termination design.

Differential data transmission provides higher noise rejection than single-ended technologies, especially low-swing families such as GTL. With differential transmission, noise is coupled onto both lines, thus is seen as common by the receivers and rejected. For this reason, the common-mode range of the differential technologies should be compared to the noise margin of the single-ended technologies. In general, differential systems will provide twice the noise rejection of single-ended systems. This is also the reason that differential transmission works best on closely-coupled interconnects (pc-board traces close together and twisted pair cable) as it helps to ensure that noise is coupled common.

Notes

This article first appeared in Electronic Design (Vol. 46, No. 25) Nov. 2, 1998 and is reprinted with permission from Penton Publishing. The article was revised in the fall of 2002 for current trends and a discussion on CML was added.