Transmission at 200 Mbps in VME Card Cage Using LVDM

Frank Dehmelt   *High Performance Linear*

ABSTRACT

Test results of a TI LVDS multipoint (LVDM) system show that a VME backplane with an LVDM electrical layer can support signaling rates to 200 Mbps.

The performance of differential signaling with the SN65LVDM176 transceiver is characterized using a commercially-available VME card cage. Test results show how changes in bus loading influence the electrical characteristics of a backplane and the resulting requirement on maintaining proper bus termination. Test results also show the effects of different backplane loading configurations and different stub lengths on maximum signaling rates and crosstalk. Bit-error-rate (BER) test results for different system configurations are included.

Contents

Figures

Tables

1   Introduction

This report delineates the major differences between low-voltage differential signaling (LVDS) as defined in TIA/EIA-644 and low-voltage differential multipoint (LVDM). It examines the pros and cons of this new technology in the context of existing high-voltage differential (HVD) standards such as TIA/EIA-485 (RS-485). Included, are the results of a bit-error-based evaluation of the capabilities of an LVDM transmission system.

The intent of the report is to help a designer estimate the performance of an LVDM system and to suggest ways to implement the design, including design issues and layout considerations for daughter cards. Backplane termination techniques and the effects of different stub lengths are evaluated.

Testing has shown that a VME backplane can be used for LVDM, with a signaling rate of 200 Mbps. Higher signaling rates can be achieved by optimizing for a specific configuration of loads, by optimizing the backplane termination, and by keeping stub lengths as short as possible. Hot swapping cards onto the bus may create a few bit errors but it does not cause bus failures. Power consumption and dissipation are not affected by the number of loads on the bus. Decoupling measures are of relatively low importance for drivers, but they are essential at the receiver.

1.1   Comparison Between LVD and Other Signaling Techniques

The primary advantages of LVD signaling are (i) high speed and (ii) low power. Single-ended and HVD signaling schemes require a signal transition of several volts, which takes *time* and *power*. At higher signaling rates, power dissipation and electromagnetic interference (EMI) become concerns, too. In LVD, the signal swing is about 300mV and is differential (or balanced), so the transitions require low power and are fast.

1.1.1 Differential vs Single Ended (SE)

The first milestone in achieving higher signal integrity and reduced EMI was the development of differential signaling. With differential signals, two wires (or conductors) are required: one wire is a true value and the adjacent signal-return wire carries the inverted value. Ideally, this leads to a complete cancellation of emitted noise. The receiver then rejects common-mode noise, and the result is a substantial improvement in both speed and reduced emissions of EMI. Further details on EMI reduction with LVDS can be found in *Reducing Electronic Interference (EMI) With Low Voltage Differential Signaling (LVDS)*, SLAA030B [1]. In many applications, these advantages have been well worth the added cost of the two-wire requirement of differential signaling schemes, as detailed in Table 1.

Table 1.   Comparison Between Single-Ended and Differential Signaling

	SINGLE ENDED	DIFFERENTIAL
Speed	Low to moderate	Moderate to high
Susceptibility to radiated EMI	High	Low
Radiation of EMI	High	Low to moderate
Cost	Low	Moderate
Common-mode voltage range	Low	Low to high
Power consumption	Moderate to high	Low to moderate

1.1.2 LVD vs HVD Signaling

LVD takes the differential advantages one step further. With a signal swing of only about 300 mV (compared to approximately 3V for HVD-systems), the transition time from one logic state to another is also decreased. This means that the next transition can occur sooner, resulting in a faster signaling rate. The transition time for LVD signals is on the order of 600 ps.

Due to the low voltages, the power dissipation and EMI are greatly reduced compared to HVD systems. HVD and LVD performance are compared in Table 2. The drawback of LVD is a reduced common-mode voltage range of zero to 2.4 V, resulting in shorter line lengths and less tolerance to noise than HVD. Both of these disadvantages are addressed by Texas Instruments extended-range devices which, by incorporating a common mode input range (–2 V to +4.4 V) and hysteresis, overcome false triggering from noise.

Table 2.   HVD and LVD Compared

	HVD	LVD
Speed	Moderate (< 50 Mbps)	High (< 2.5G bps)
Susceptibility to radiated EMI	Low	Moderate
Radiation of EMI	Moderate	Low
Cost	Moderate	Moderate
Common-mode voltage range	High	Low
Power consumption	Moderate	Low

1.2 Comparison Between LVDS and LVDS-Multipoint (LVDM)

Figure 1 illustrates the three commonly-used interconnect configurations. Figure 1(a) shows a point-to-point system that uses a single driver and a single receiver. Figure 1(b) shows a multidrop system consisting of a single driver connected to more than one receiver on the same bus. Information on LVDS multidrop systems can be found in *LVDS Multidrop Connections*, SLLA054 [2]. In the multipoint system, as shown in Figure 1(c), there are multiple transceivers (alternatively multiple drivers and multiple receivers) on the same bus. Systems that have two or more drivers on the bus are considered multipoint busses.

There are also different types of multipoint systems, but the one discussed in this report is the basic bus configuration shown in Figure 1(c). Notice also that the multipoint system requires two termination resistors. The double termination doubles the load on the driver. To maintain the same signal levels on the bus as for a singly-terminated LVDS, the line driver current must be doubled, because it sources two transmission lines in parallel. Double termination of the bus allows for half-duplex operation. In a half-duplex system only one driver can be active at any time but, with appropriate supervision of the bus, communication is possible in both directions for bus participants. Note: Texas Instruments uses LVDS and LVDM in its part designations to distinguish between the output-current drive specifications of its LVDS and LVDM line drivers.
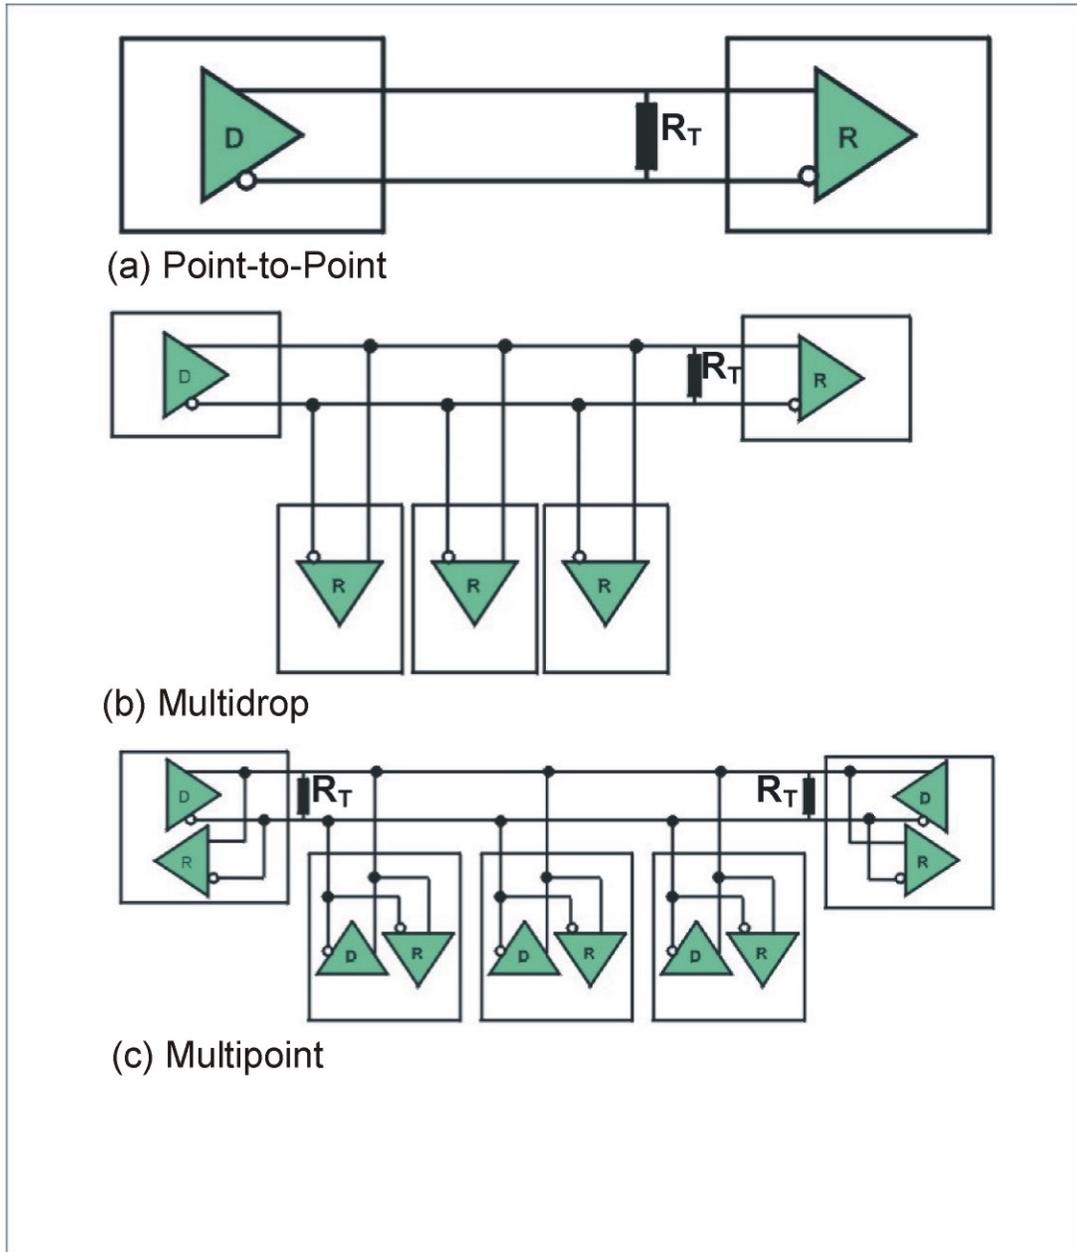
Figure 1. Common Differential Interface Configurations

2 Test System Description

2.1 Hardware

The LVDM system developed and tested in the present work consisted of a 21-slot VME backplane and small daughter cards, each containing eight Texas Instruments SN65LVDM176 LVDM transceivers. Twenty-one daughter cards were assembled for testing with the backplane bus lightly, partially, and fully loaded (all VME slots populated). The physical system is described in this section.

2.1.1 *Backplane*

The tested LVDM used a 21-slot VME backplane. Because of their high pin count (up to 96 per slot), high reliability, and relatively high signaling rates (~40 Mbps), VME boards are used widely in backplane applications. The VME backplane is a multi-layer board intended for 50-Ω to 60-Ω single-ended transmissions and, therefore, it does not provide the optimum performance for 100-Ω differential transmission. An ideal differential line has a 100% coupling between the two lines of a pair and no coupling towards ground which, of course, is not the case for a pair of two ground-referenced 50-Ω lines. Although the VME backplane and bus may not be ideal for testing an LVDM configuration, the primary reasons for selecting it were the following:

- The LVDS standard (TIA/EIA-644) specifies the electrical layer only and contains neither bus, backplane, nor connector specifications.

- Circuit cards in a VME chassis may be installed or removed quickly and easily, and the backplane/connectors are very affordable.

- VME is widely used and is a good candidate for improved bandwidth.

A 21-slot, 8-layer backplane, equipped with the well-known DIN41612, 64-pin ac female connectors was chosen. Each row of pins was routed on a separate layer; therefore, adjacent pins on one row were chosen to carry the differential signal. Between every signal pair, one pin was left open (except for one pair, which was used as a reference), and this provided the option either to leave the line floating, to tie it to ground, or to set it to a specified voltage. Figure 2 shows the pin-assignment on the plug; Figure 3 depicts the fully-equipped backplane.
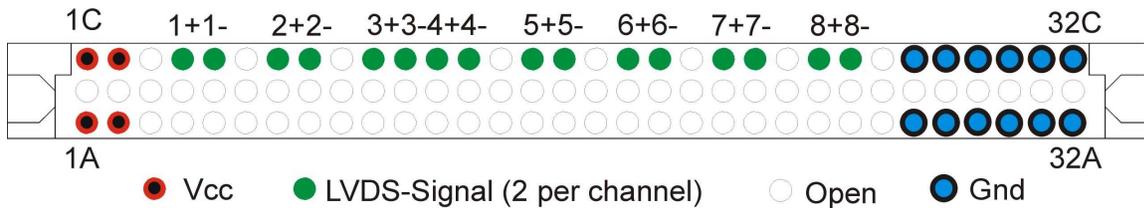

Figure 2.   DIN41612-Plug With Pin Assignment
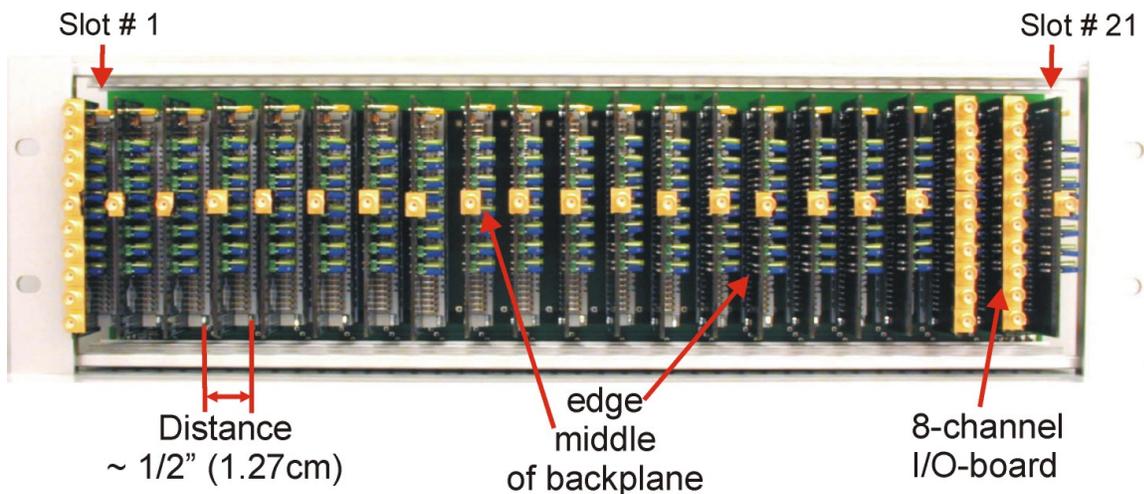

Figure 3.   Fully-Equipped Backplane

For this test setup, eight individual transceivers were used in parallel to provide a common parallel interface with a realistic amount of crosstalk. All measurements—jitter and bit-error rate testing (BERT)—were taken using a channel located down the middle of the backplane, with a floating line on either side of the signal pair. For comparison, data were also collected (not presented here) on a channel at the edge of the backplane. It was found that at the edge of the backplane, the crosstalk was reduced, as evidenced by improved BERT results and a more open eye pattern. To provide worst-case results, the data presented are for the channel that runs across the center of the backplane.

Inductance-capacitance-resistance (LRC) measurements were made on the backplane as each daughter card was installed. In this test, a step-function was applied to a non-terminated line. The amplitude of the incident wave and the delay between the first slope (incident wave) and the second allowed the derivation of the line impedance, Z_0, the inductance, L_0, and the capacitance, C_0. The test can be applied for both single-ended systems and differential lines. For more details see *Data Transmission Seminar Reference Manual '99*, SLLD005 [3]. The results of these measurements are shown in Figure 4.
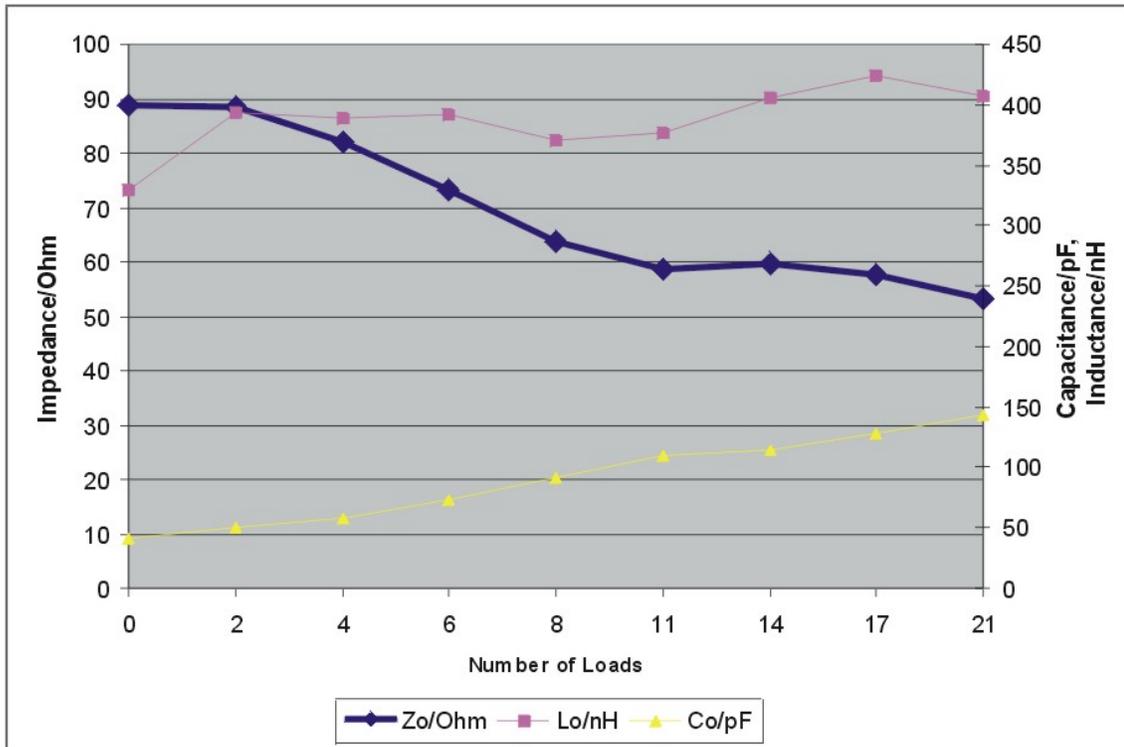
Figure 4.   Backplane Impedance as Daughter Cards Are Installed

Figure 4 shows that the inductance of the backplane remained nearly constant as cards were added to the bus, but capacitance increased and, therefore, the characteristic impedance of the bus decreased. The impedance ranged from 53 Ω to 89 Ω. A termination value of 62 Ω was selected to cover both the lightly and heavily loaded busses expected in real applications. Surface mount device (SMD) resistors were soldered directly to the connector solder pads of the first (slot 1) and last (slot 21) connector on the backplane.

Several measurements validated that 62 Ω was satisfactory for point-to-point as well as for a fully-loaded (LVDM) backplane. Of course, if the number of loads and their spacing is fixed, the designer can optimize the termination value according to the plot in Figure 4. The performance can be increased significantly if the termination is precisely matched to the specific impedance of the backplane, because reduced reflections allow a wider opening of the eye.

2.1.2 Delay Considerations

At signaling rates in the range of hundreds of megabits per second, propagation time on the transmission media is an important factor. Because the bit duration is lower than the backplane's propagation delay, correct timing of the readout at the receiver is essential. There are two important factors at work here. First, the delay increases with the physical distance between the driver and the receiver[1]. Secondly, as shown in Figure 4, the capacitance increases with increased backplane loading (more cards in the chassis), which increases the propagation delay, i.e., it slows down the backplane even more. Therefore, a transmission approach is recommended, in which a data-clocking signal is sent with the data so that the receiver can trigger on a separate clock line. Alternatively, embedding the clock into the data stream offers another solution (but one that will decrease the effective data throughput and is not sensible in a parallel configuration). Figure 5 shows the timing for the most extreme locations on the backplane, including a budget of approximately 10 ns for the cable connections between generator and driver as well as receiver and analyzer. Obviously, a signaling rate of 300Mbps is not applicable to the full-load situation in worst-case locations.
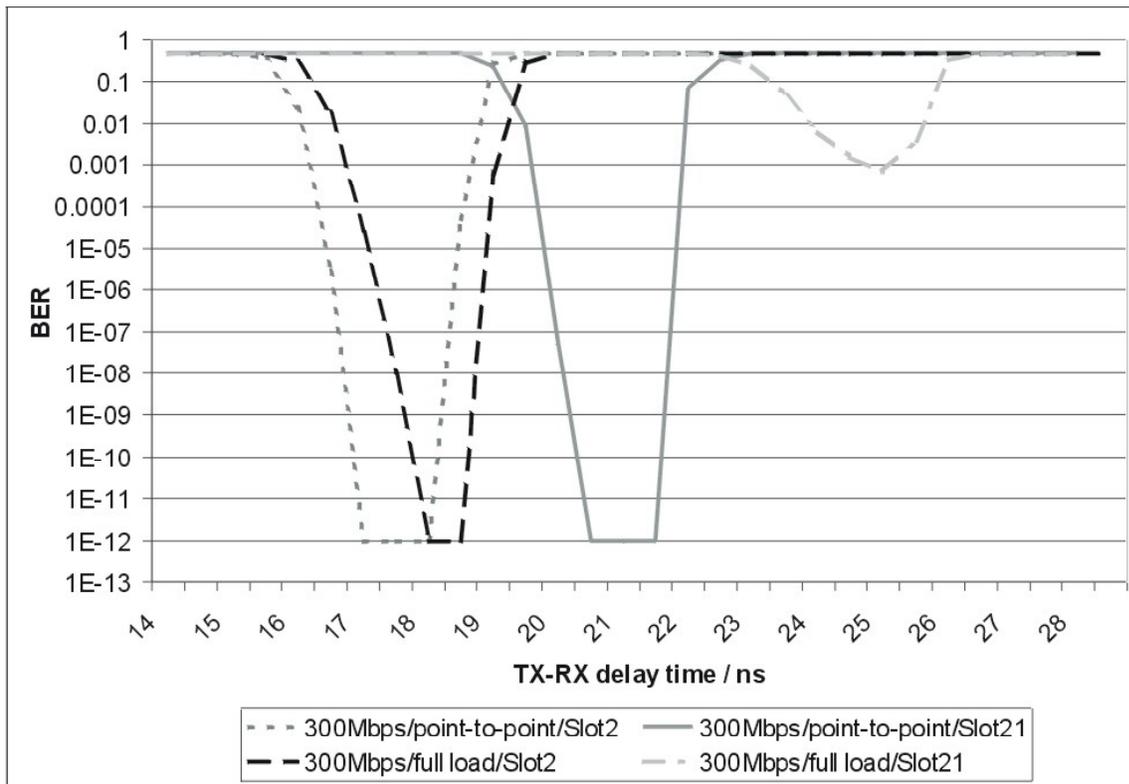
Figure 5. Propagation Time vs Load Conditions and Location of Receiver

---

[1] Propagation delay is calculated as the product of the propagation delay per unit distance and the distance.

2.1.3 Daughter Cards

The daughter cards are four-layer printed circuit boards (PCB) with the layer arrangement shown in Figure 6. The dielectric is FR4 with a thickness of 254 μm (0.01") per layer, each coated with 35-μm (0.00138") copper layers. Manufacturing mode was hot air leveling (HAL).

Signal Plane; Carry All Single Ended and Differential Signals

Ground Plane

V_{CC} Plane (3.3 Vdc)

Logic Plane, Supporting the Enable and Disable Signal Traces
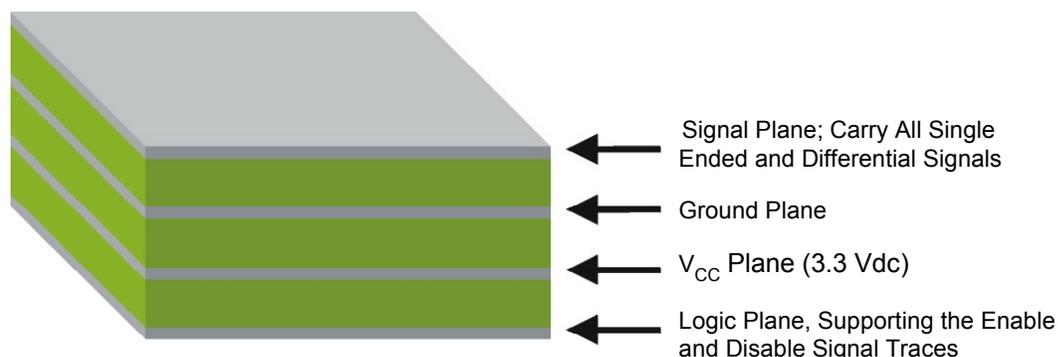
Figure 6.   Layer Stacking of Daughter Cards

The daughter cards are shown in Figure 7 and were designed especially for these tests. Each daughter card contained eight SN65LVDM176D transceivers, with separate drive- and read-enable control lines. Jumpers allowed any combination of transmission modes (such as anti-parallel communication on adjacent lines) to be set up on each device on each card.

Both the single-ended and differential lines were designed to match the intended impedance. The single-ended lines were point-to-point connections; therefore, the length is not important, because it is not considered as a stub line. The differential lines were kept as short as possible. The total stub line length, including the connector, was 30.4 mm (1.20"), of which 26.4 mm (1.04") was due to the connector. Assuming a rise-time of 500 ps, this line length already approaches the critical trace length ($l_{max} \approx t_r/2t_p$) that determines when transmission line effects are noticeable and a termination is necessary.

All input data (single-ended and TTL) were applied via SMA connectors installed at the opposite card edge (away from the backplane). The cards were designed with a 470-Ω series resistor connected to the output of each receiver, followed by a 50-Ω pulldown resistor. When the output was being measured (i.e., connected to test equipment), the additional 50-Ω input of the test equipment in parallel with the 50-Ω pulldown created a total load of 500 Ω on the receiver output and also created a 20:1 resistor divider into the measurement equipment.

Due to limited space on the edge of the daughter card, only two channels (Channel 1 near the edge of the backplane and Channel 4 in the middle of the backplane) were equipped with SMA connectors, as seen in Figure 7. The receiver outputs from the other six channels could be monitored from jumper posts on the daughter card.

The daughter cards were designed to allow V_{CC} to be provided via the backplane (pins 1 and 2 on both rows *a* and *c* carry the supply voltage while pins 27 through 32 on rows *a* and *c* are connected to ground, see Figure 2). However, jumpers on each daughter-card allow an external V_{CC} to be provided directly to the daughter card. This feature was used, for example, to evaluate the driver's power consumption individually while all receivers were powered via the backplane supply.

V_{CC} filtering on each daughter card was provided by a series inductor (100 nH) and a capacitor (4.7 μF, tantalum) on the supply line. Additionally, each individual transceiver had a decoupling capacitor (10 nF, ceramic) located at each device. For comparison, two more boards were assembled with no decoupling measures at all, i.e., no inductor, no tantalum capacitor, and no individual ceramic capacitors. One board was set up for slew-rate-controlled outputs, by soldering different capacitors to the differential output.
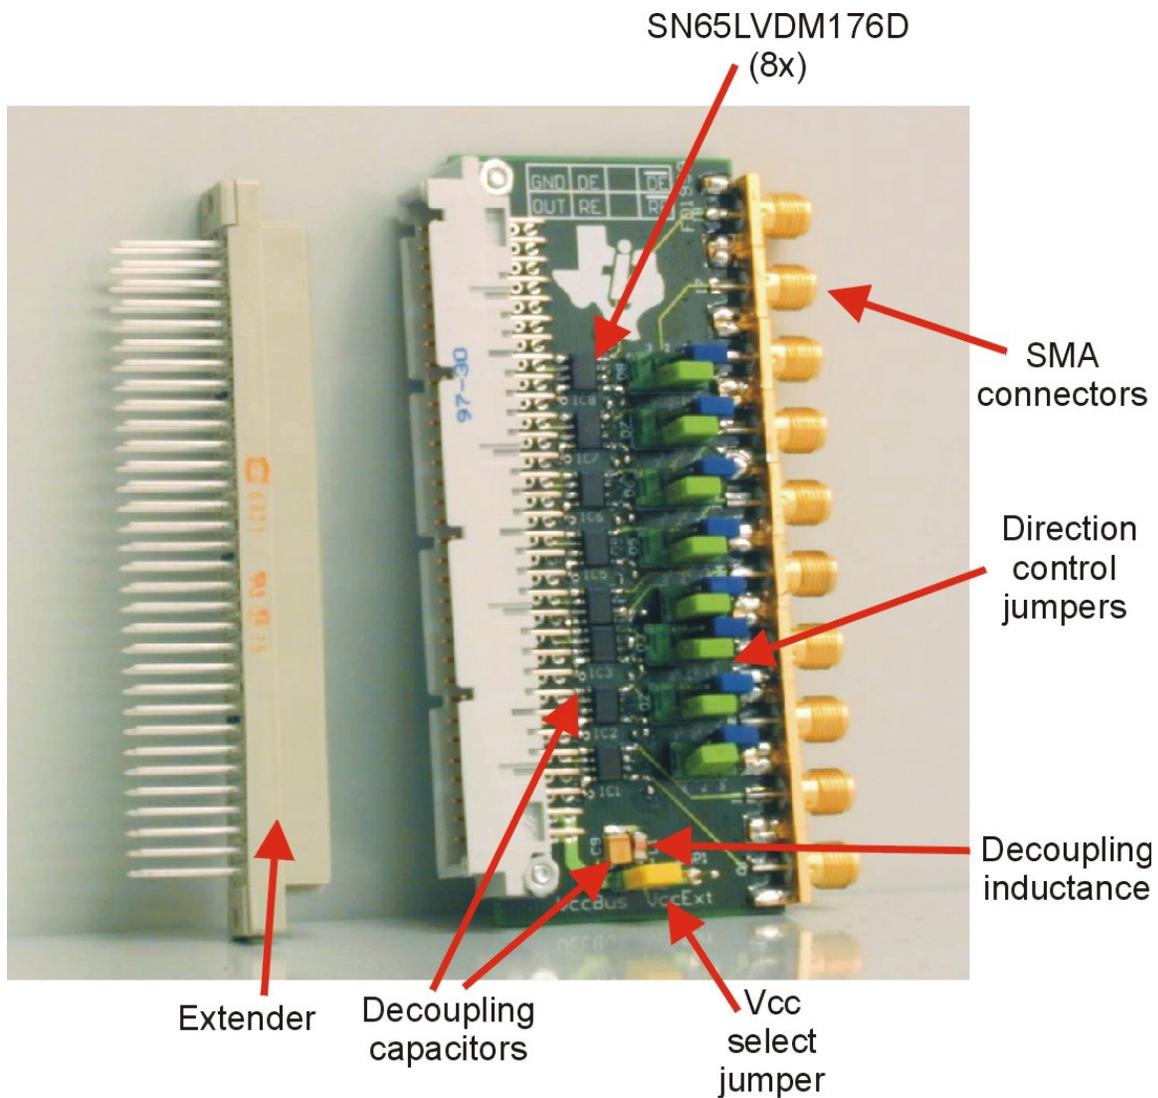

Figure 7.   Daughter Card With Stub Length Extension

2.2 Test Setup

2.2.1 *Eye Pattern*

The jitter measurements were performed by applying a pseudo-random bit sequence (PRBS) of $(2^{15} - 1)$-bit length into the driver. The differential signal was measured at the receiver inputs on the backplane for the module being monitored. The jitter values represent the peak-to-peak jitter at the zero-volt differential of the eye pattern (refer to Figure 9 for details on eye pattern). The pattern was generated by a Hewlett-Packard 81200 data generator / analyzer system for the monitored lines. To induce crosstalk, additional signals from a Tektronix HFS9009 pattern generator were applied to the adjacent lines.

Jitter on the backplane was measured with Tektronix real-time oscilloscope TDS694C and a P4267 1-GHz differential probe head[2]. The output jitter of the receiver was measured (after the 20:1 voltage divider on the daughter-card described earlier) using a coaxial cable connecting the SMA connector on the daughter card to the oscilloscope. The setup is shown in Figure 8.

---

[2] The 1-GHz bandwidth is actually not quite sufficient (it should really be capable of detecting the 5th harmonic of the data clock), but at the time of starting the tests it was the fastest differential probe commercially available.
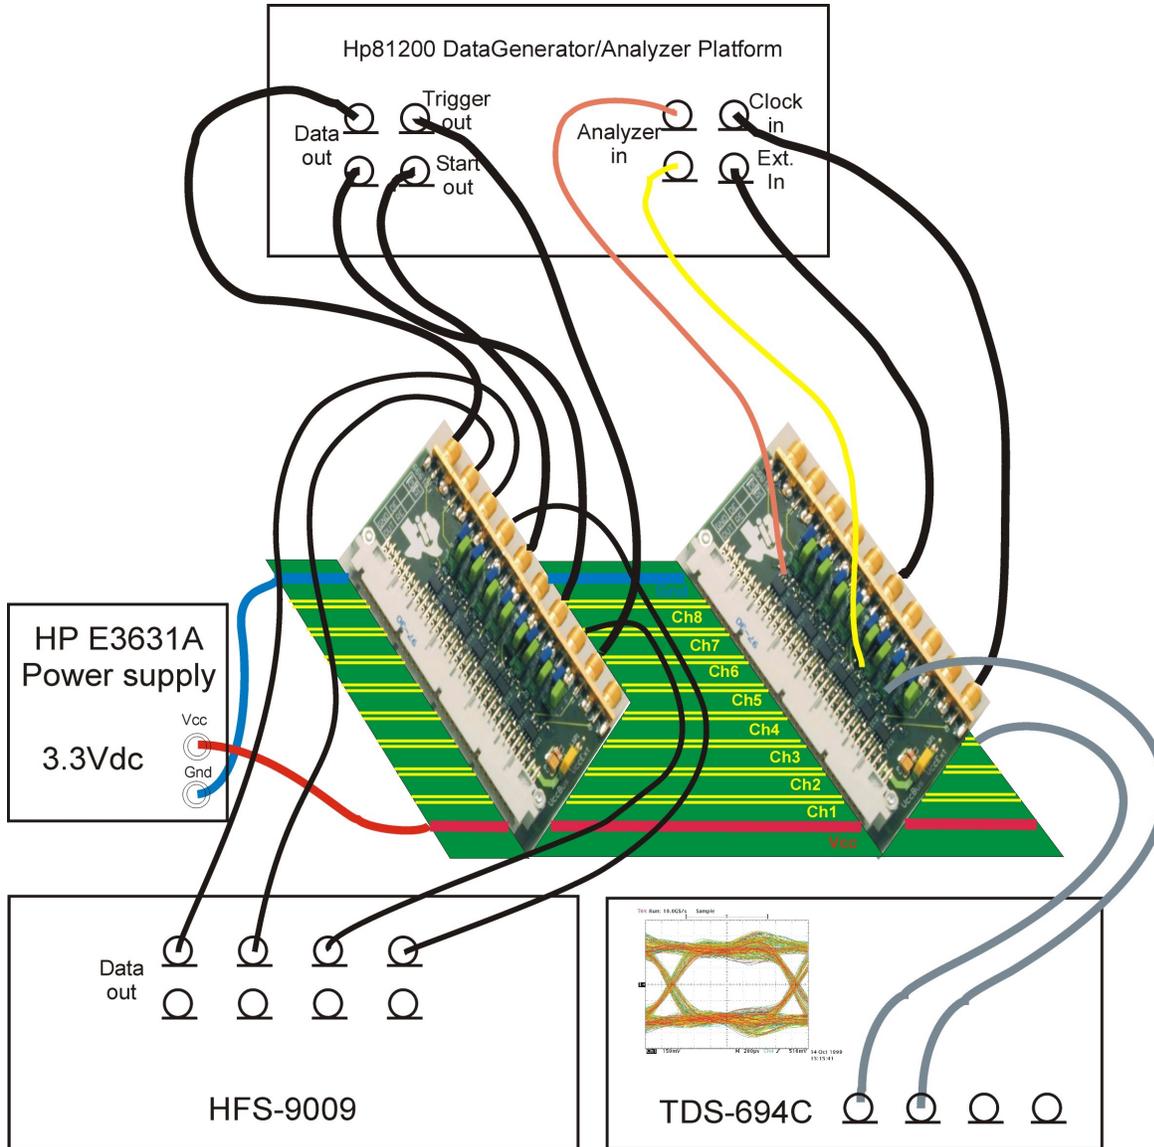
Figure 8.   Test Setup

A typical eye pattern, captured over a time interval of 10 seconds at infinite persistence, is displayed in Figure 9. Several characteristics of the eye pattern indicate the signal quality of the transmission circuit. The height or opening of the eye above and below the receiver threshold level (refer to Figure 9) at the sampling instant is the noise margin of the system. The lateral spread of the transitions across the receiver threshold measures the peak-to-peak jitter of the data signal. The signal rise and fall times can be measured relative to the 0% and 100% levels provided by the long series of low and high levels.
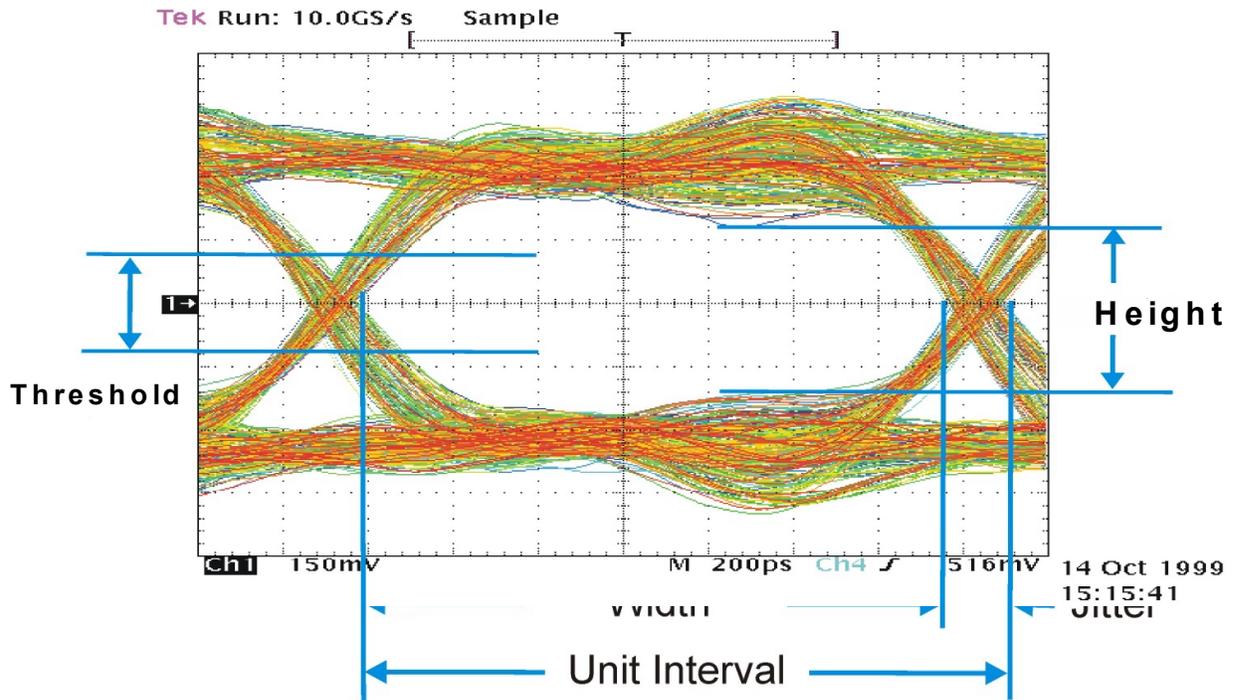
Figure 9.   Typical Eye Pattern

Jitter is the time interval over which the logic state transition of a signal occurs. The jitter may be given either as an absolute number or as a percentage with reference to the unit interval (UI). This UI or bit length is the reciprocal of the signaling rate, and the time during which a logic state is valid is just the UI less the jitter. Percent jitter (the jitter time divided by the UI multiplied by 100%) is more commonly used and represents the portion of the UI during which a logic state should be considered indeterminate.

The data applied for these tests was an NRZ-coded pseudo-random binary sequence (PRBS) with a depth of 32K bits.

2.2.2 Bit-Error Rate (BER)

The setup used to apply the data for the BER tests was identical with what has already been described in Section 2.2.1. For the measurements with the clock transmitted internal of the Bit Error Tester (referred to as *Test Set 1*), Channel 4 was monitored by the Hewlett-Packard 81200 data generator / analyzer platform. For the tests with the clock transferred across the backplane (referred to as *Test Set 2*), Channel 4 was used for the start signal, and channels 3 and 8 were monitored using the installed jumper posts at the receiver outputs (refer to Table 1 and Figure 8). As expected, the edge-channel proved to have slightly better performance here.

The threshold voltage at the analyzer input was set to 75 mV, which (with the 20:1 voltage divider) represented 1.5 V at the receiver output. The delay between the data generator and the analyzer was optimized for each test in Test Set 1 to detect the bit in the middle of its unit interval. Because the backplane delay varied with bus loading, this analyzer delay was adjusted while monitoring the error rate in order to tune it to the timing with the lowest error rate.

For measurements in Test Set 2, these fine-tunings were not necessary because the analyzer self-triggered by receiving the start signal (which had been subjected to the same delay and was exposed to the same mismatches—stubs and improper termination—as caused the reflections). Only a restart was necessary after each relocation of the boards.

The channel assignment for Test Set 2 is given in Table 3.

Table 3.   Signal Assignment for the Measurements of Test Set 2

CHANNEL	SIGNAL
1	Clock bit rate identical to data, e.g., 300 Mbps equivalent to 150-MHz clock
2	PRBS†, data source inducing crosstalk
3	PRBS, analysis point
4	Start signal, triggering the analyzer
5	PRBS, data source inducing crosstalk
6	PRBS, data source inducing crosstalk
7	PRBS, data source inducing crosstalk
8	PRBS, analysis point

† PRBS = pseudo random bit sequence

3   Measurement Results

Jitter measurements were taken for reference only, because the effect of the jitter depends on the sensitivity of the downstream hardware, and the sensitivity is often unknown.

3.1 Eye Pattern

Jitter values measured for an eye pattern provide a good indication of the quality of the interface circuit. The jitter results and oscilloscope plots were used mainly to determine the best time interval (analyzer delay) during which to trigger the receiver, because a bit-error tester was available to give a direct measure of signal integrity. The eye pattern showed a dependency on the signaling rate and the load configuration, as well as on the location of the driver and receiver.

Figure 10 shows eye patterns of the differential bus signals detected at Slots 2, 6, 9, and 20, all taken at 300 Mbps with the driver located at Slot 7 on the fully-loaded backplane. The optimum trigger-timing varies with location on the bus. Towards the end of the backplane (Slots 2 and 20) a wide-open eye is visible. At Slot 6, the eye is almost closed in the center, while at Slot 9 the widest position is right in the middle. Traces below the threshold voltage may cause bit errors; this will be discussed in Section 3.2. Because the backplane is fully loaded, the 62-Ω termination resistor is too high compared with the bus impedance. Therefore, overshoots are caused by reflections. Because the driver is not located in the center on the backplane, and the propagation delay time on the backplane is longer than the bit duration, one cannot predict which reflected signals meet each other on the bus. For example, with particular reference to slot 6, it appears that some reflections of negative traces overlay positive ones, which reduces the amplitude during the steady-state time of the bit. A second use of the eye pattern was to determine the impact of driver slew rate control on reflections and amplitude (see Section 3.8, *Slew Rate Control*).
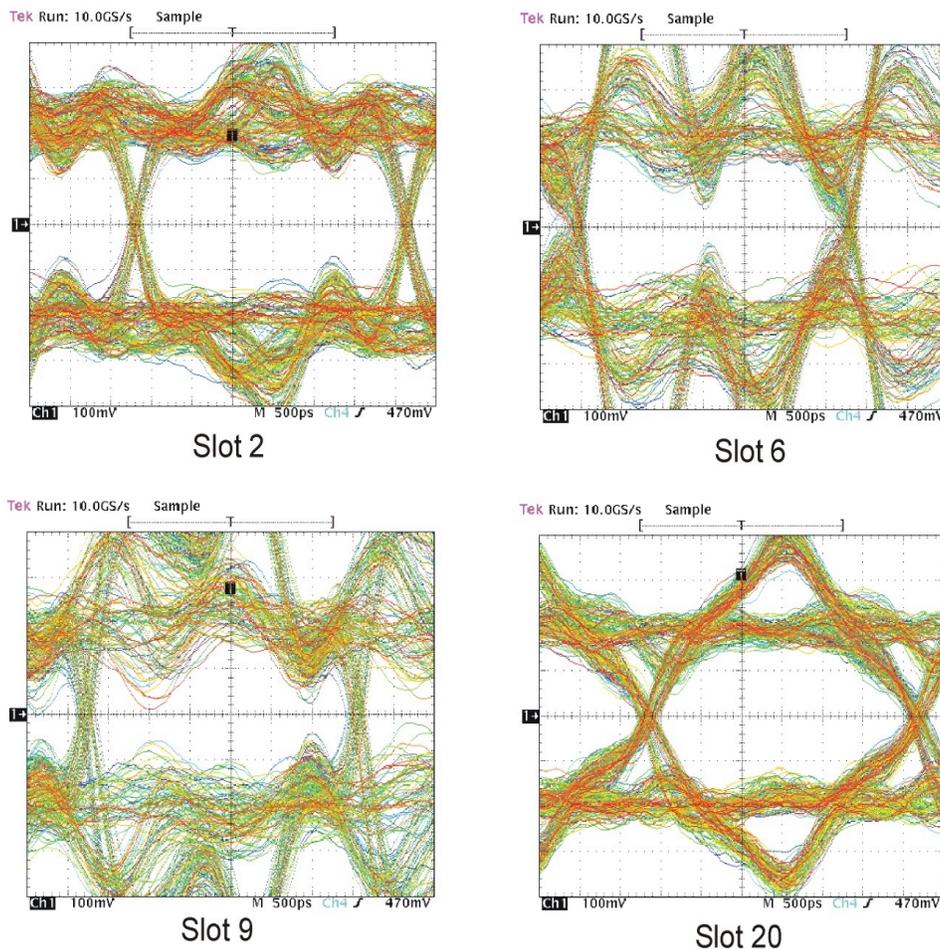

Figure 10. Eye Pattern at 300 Mbps With Driver at Slot 7, Measured at Slots 2, 6, 9, and 20

3.2 Bit Error Rate (BER)

BER measurements were taken with two, four, eight, and 21 daughter cards installed in the backplane. For the partial loads (four and eight), the arrangement of the daughter cards was varied, either by grouping all the daughter cards close together in a single cluster, by arranging them in several tight clusters (e.g., for eight cards, having two groups of four, four groups of two, and so on), or by distributing them evenly along the entire backplane. For each configuration, the worst case situation relative to the positions of the drivers and receivers was evaluated, and the maximum signaling rate was determined. The criterion for the maximum signaling rate was a BER of transmitting at least 10^{12} bits error free. Error rates shown as 1E-12 in the following tables indicate the error-free transmission of at least 10^{12} bits. Some tests were run up to 10^{13} bits or more. However, because of time restrictions, not all measurements could be performed for that long.

As illustrated in Table 4, the worst performance occurred with a partial loading of eight daughter cards, and it is obvious that the specific placement of the daughter-cards along the backplane has a major effect on performance. When the daughter cards were distributed evenly along the backplane, a BER of 1E-12 was easily achieved. In addition, grouping all daughter-cards together created a large impedance discontinuity, which caused substantial disturbances (seen in the eye pattern). A trigger adjustment would likely help this situation, but the more obvious solution would simply be to size the chassis for the number of slots to be used, and keep this backplane fully (or almost fully) loaded.

Table 4.   Signaling Rate vs BER for Test Set 1 Measurements (Worst Case)

SIGNALING RATE (MBPS)	BIT ERROR RATE			
	POINT-TO-POINT	FOUR NODES	EIGHT NODES	FULLY LOADED
175	1E-12	1E-12	4.3E-12	1E-12
200	1E-12	1E-12	6.0E-12	1E-12
300	1E-12	2.2E-12	2.5E-05	2.0E-04

NOTE: A BER of 1E-12 indicates that at least 10^{12} bits were transmitted error free.

The same tests have been performed with the clock transmitted over the backplane as well as the data. (For details of the channel assignment, refer to Table 1.) As can be seen in Table 5, even better performance was achieved this way. Because of the high performance at 200 Mbps (Table 5), additional tests were conducted at 250 Mbps and the results of these are also shown in this table.

Table 5.   Signaling Rate vs BER for Test Set 2 Measurements (Worst Case)

SIGNALING RATE (MBPS)	BIT ERROR RATE			
	POINT-TO-POINT	FOUR NODES	EIGHT NODES	FULLY LOADED
175	1E-12	1E-12	1E-12	1E-12
200	1E-12	1E-12	1E-12	1E-12
250	1E-12	1E-12	4.7E-11	1.6E-11
300	1E-12	2.0E-12	2.0E-11	9.4E-11

NOTE: A BER of 1E-12 indicates that at least 10^{12} bits were transmitted error free.

The tests were repeated with a short-card extender installed between the chassis connector and the daughter-card connector, increasing the stub length to each receiver by approximately 12 mm or 1/2 inch. For a fully-loaded backplane, the signaling rate was decreased by approximately 25% to achieve an error-free transmission of 10^{12} bits, and the same derating was necessary for the partial load, too. This applied to both test set configurations. As expected, stub length had less effect on the point-to-point configuration.

These tests indicate that LVDM typically performs extremely well, even under less than ideal conditions. Significant improvement can be expected when the system is optimized for a particular loading configuration (i.e., only as many slots as actually needed, and matched termination), stub lengths are minimized, and the backplane and connector impedance is optimized for a characteristic impedance of 100 Ω.

3.3   Live Insertion (Hot Swapping)

Live insertion (also called hot swapping or hot insertion) is generally classified into three categories:

- Bus shuts down and resets: no damage to the system

- Bus is idle during insertion: no bus contentions

- Bus is active during insertion: no data loss during an insertion or removal

Texas Instruments devices provide a high impedance to the bus when disabled or when V_{CC} drops below 1.5 V dc. This avoids damage and bus contentions during power up and when adding daughter cards. The feature also minimizes data losses or errors because the high-impedance state of the devices means that other devices on the bus are not subjected to significant impedance changes as cards are installed or removed. However, there is no pre-charge feature of the bus pins, and each connector, trace, and device has a finite capacitance. The charging of this capacitance may cause some errors or failures during the insertion or removal of cards from the bus.

During these tests, daughter cards were installed and removed frequently while the bit error rate was monitored. No errors were observed during the full-load tests (21 daughter cards installed) but during the lightly- and partially-loaded tests (where the highest jitter occurred, as evidenced by jitter measurements using eye patterns), the worst case of three bit-errors per insertion was recorded when cards were installed.

The test chassis did not have card guides to align the daughter cards with the backplane connectors during insertion. During insertion, the daughter cards were intentionally tilted in an attempt to create bit errors during this operation. The daughter cards were designed for V_{CC} at one connector edge and Gnd at the other, so by tilting the daughter-cards, first contact would be made with either V_{CC} or Gnd. Possible reasons for potential errors might include (i) timing within the bit stream, (ii) capacitive loading, and (iii) a differential spike from contact timing. However, even when intentionally trying to generate errors, a maximum of only three bit-errors occurred. It would appear that the errors were the result of the specific timing of the insertion and capacitive loading rather than of the accuracy of connecting. Whenever errors occurred, their total number never exceeded three per insertion at transmission speeds of 200 Mbps and even as high as 300 Mbps. During daughter-card removal, no errors were ever observed. There was no measurable dependence of bit error rate on the precision with which the connections were made.

3.4 Bus Contentions

Conflicts were intentionally created to verify the performance (i.e., durability) of the devices under these conditions. Two drivers (on the same backplane channel) were asserted simultaneously and BER and eye patterns were monitored. One driver was driven by the generator of the BER, the other one by the HFS9009 generator (set to slightly different signaling rate and delay). As expected, the data were corrupted, but no damage to the system or any of the devices occurred. This test was run for more than 60 hours with no devices failing, and the system resumed normal operation after the fault was removed.

3.5 Power Consumption

Power consumption was monitored while powering the driving daughter card from a separate power supply and monitoring Icc with a precision Wavetek 1281 current meter. The number of receivers was increased from zero to 20, but the current varied only slightly. This effect was very small, but when using the Wavetek 1281, a 0.3% decrease in driver Icc was observed between the point-to-point configuration (one driver and one receiver) and a fully-loaded bus (20 receivers).

Figure 11 shows a comparison of the SN65LVDM176D between competing interfaces. It is quite obvious that the power consumption was almost independent from the switching frequency. The 50% increase was caused almost exclusively by the receivers and it resulted from the TTL output. For TIA/EIA-422 (one channel of AM26C31C and AM26C32C), the supply current increased significantly with frequency.
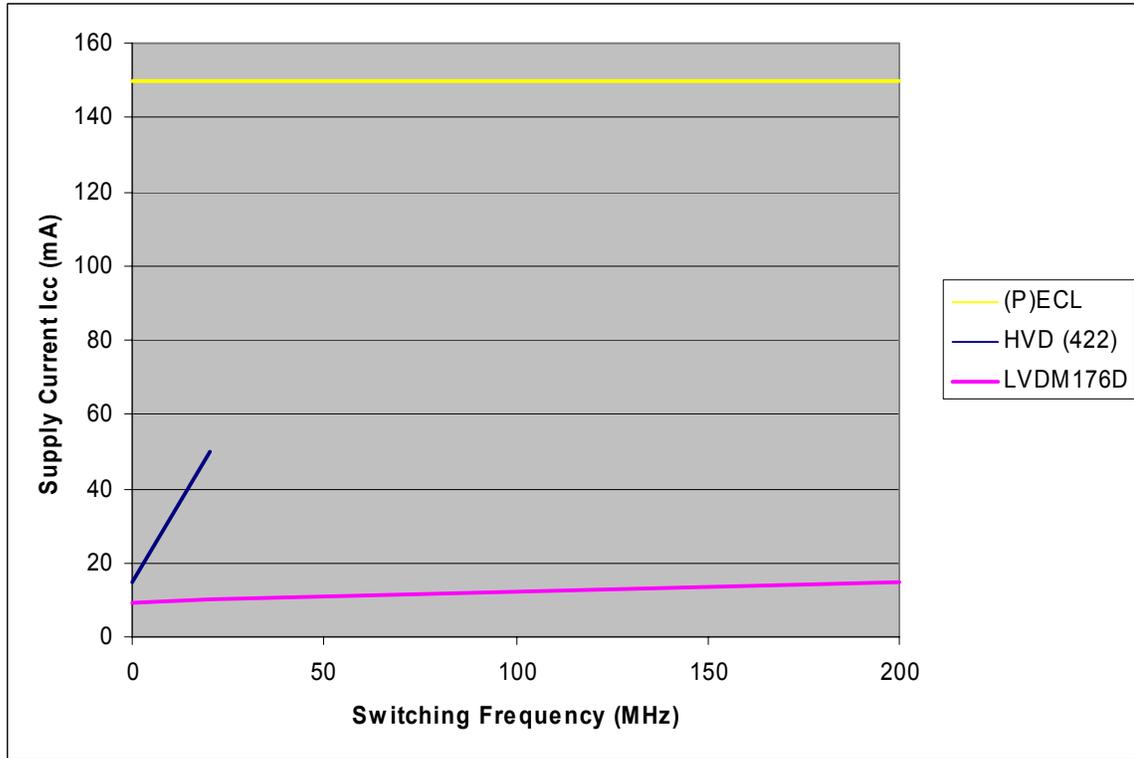
Figure 11.   Typical Supply Current vs Switching Frequency

3.6   Decoupling

For comparison, the boards without any V_{CC}-decoupling measures were used either (i) one at a time for driving, (ii) one at a time for receiving, or (iii) one for driving and the other for simultaneously receiving. As expected, when the driver had no decoupling of the power supply, the decrease in its performance was barely measurable. This insensitivity resulted from the balanced transmission mode in which little or no common mode current bounced between V_{CC} and Gnd. Further enhancement was achieved by the current-mode driver characteristic, because it always drew the same amount of current and only changed the current direction at the outputs.

When the receiver had no V_{CC}-decoupling measures, there was a significant degradation of its performance. The receiver had a voltage-mode output that had to maintain rise times of about 100 ps and swings of 2.5 V. It draws much more supply current during the transitions than at steady state. This additional current was supplied by the decoupling capacitors, located between V_{CC} and Gnd. Decoupling of drivers is recommended, therefore, only when it is required for receivers.

3.7 Shielding Between Backplane-Channels

The single, unconnected bus-line on the backplane between any two channels was either set to different voltages or it was left floating. For point-to-point interconnections, a fixed voltage on these lines gave improved performance from the standpoint of crosstalk. Voltages of V_{CC}, 1.2V and GND were applied at both ends of the backplane. For partial loading, the benefits of shielding were just about balanced by the extra capacitive loading. Fixing the voltage forced the signal to charge the parasitic capacitances that existed between the adjacent and differential lines. However, at full load, the drawback of additional capacitance outweighed the benefits, and the performance declined (the signal amplitude decreased with higher capacitive loading). With the maximum load of 21 nodes, the amplitude might even be insufficient to pass the threshold voltage.

The difference between the impact on crosstalk of Channel 3 (without any extra line between itself and Channel 4 under observation) and Channel 5 (with an additional line towards Channel 4) was negligible. Because of the additional pin-count and backplane area, the shielding provided insufficient benefits to justify itself.

3.8 Slew Rate Control

The board, with several different capacitors (ranging from 4.7 pF to 47 pF) soldered to the outputs of the driver, was used to evaluate the impact of slew-rate control. Eye patterns were observed on the channels with different loads. Figures 12 and 13 show plots taken at slot 18 of a fully-populated backplane running at 100 Mbps with the driver always located in the center of the backplane. Figure 12 shows the eye pattern with the driver outputs slowed by a 4.7-pF capacitor; Figure 13 shows the pattern when the driver was equipped with a 47-pF capacitor.
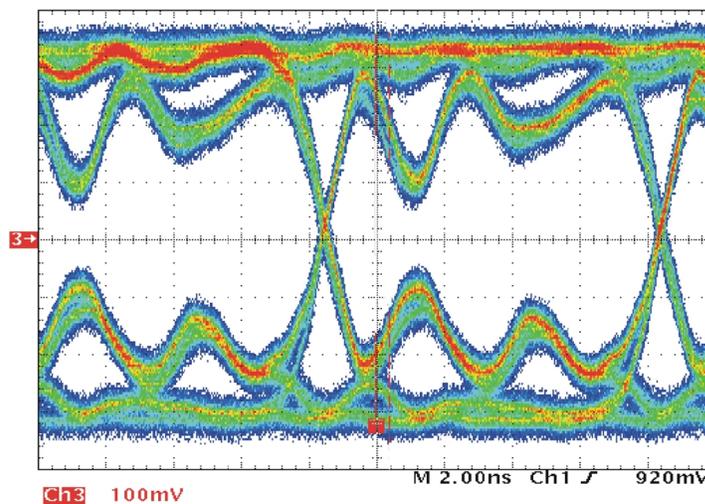

Figure 12.   Slew Rate Control With 4.7-pF Capacitor—Eye Pattern at 100 Mbps
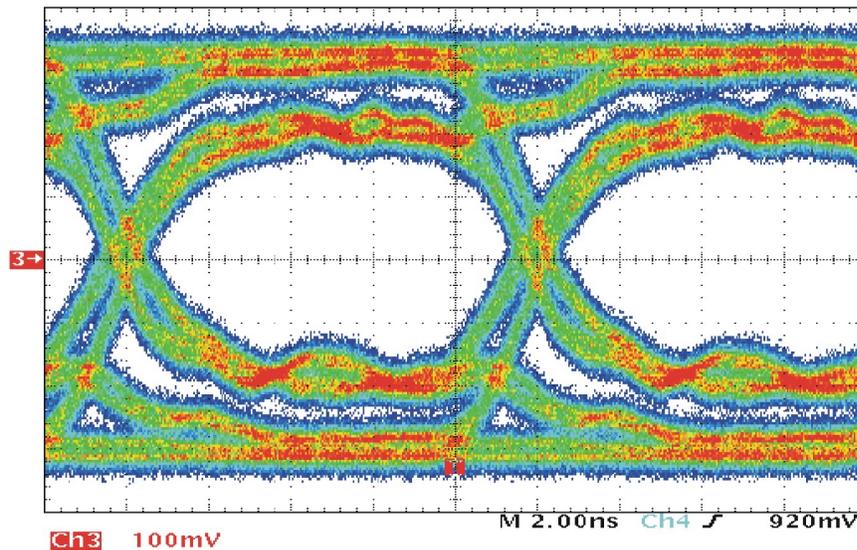
Figure 13. Slew Rate Control With 47-pF Capacitor—Eye pattern at 100 Mbps

In general, reflections caused by a misterminated line can be reduced by placing an appropriate output capacitor at the driver. The obvious cost is that this capacitor adds capacitive loading to the driver and reduces the slope of the output voltage transition.

In multipoint scenarios, every potential driver must be equipped with a capacitor. However, this degrades the output swing too much, because this distributed capacitance has an impact on the bus impedance. Therefore, the drawbacks caused by the slew rate control as implemented here outweigh the benefits. In the future, this problem will be overcome with slew-rate control integrated right into the driver, as in M-LVDS parts. Meanwhile, it is recommended that a compromise value of the termination-resistor be used if the load conditions change significantly.

4 Conclusion

These tests have shown that a VME backplane with an LVDM electrical layer can support signaling rates to 200 Mbps. Higher signaling rates can be achieved by keeping stub lengths as short as possible, and by optimizing for a specific configuration of loads (i.e., by limiting the slots to the number of intended daughter cards in order to avoid impedance imbalances and by terminating with the exact match). Hot swapping cards onto the bus may create bit errors but it does not cause bus failures. Power consumption and dissipation are not impacted by the number of loads on the bus and are lower than most other electrical layers at comparable speed. Decoupling measures are relatively unimportant for drivers but they are essential on the receiver side.

5   References

1. *Reducing Electronic Interference (EMI) With Low Voltage Differential Signaling (LVDS)*, Application Report, Texas Instruments Literature Number SLLA030b.

2. *LVDS Multidrop Connections*, Elliott Cole, Application Report, Texas Instruments Literature Number SLLA054.

3. *Data Transmission Seminar Reference Manual '99*, Frank Dehmelt, Texas Instruments Literature Number SLLD005.

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third–party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

Mailing Address:

Texas Instruments
Post Office Box 655303
Dallas, Texas 75265