FAIRCHILD SEMICONDUCTOR™

March 2002
Revised March 2002

Section 3 - Backplane Architecture Backplane Designer's Guide

The primary criteria for backplane design are low cost, high speed, and high reliability. To attain these often-conflicting goals, it is necessary to maximize the performance of every aspect of the design, including the backplane configuration. A prime consideration is how data is transmitted between the driver and receiver(s), and data transmission depends of on three backplane characteristics:

1. The distance or space between the sending and receiving devices
2. The speed at which data must be sent to a receiving device
3. The capacitive loading the signal will see traveling between the driver and receiver(s)

When constructing a backplane, the designer has to make several choices with regard to connectivity, topology, and timing architecture. Connectivity refers to data transmission, and specifically to whether data is transmitted in parallel across multiple lines or in a serial stream. Topologies are the different methods for configuring a backplane layout: point-to-point, multidrop, and multipoint. Timing architecture refers to the choice between synchronous, source-synchronous, or asynchronous signal interfaces. Because all these options have strengths and weaknesses, the final design choices depend on the requirements the backplane must meet.

Section Reference

This section presents information covering:

- Connectivity of serial and parallel transmission designs
- Configuration of point-to-point, multidrop, and multipoint topologies
- Synchronous, source-synchronous, and asynchronous timing architectures

Section	Section Title	Contents
1	Introduction	Application demands, basic backplane considerations, and how to use this guide.
2	Backplane Protocols	Descriptions of different backplane bus protocols, including PCI- and VME-based protocols.
3	**Backplane Architecture**	**Topics relevant to backplane configuration, including parallel versus serial configuration and different configuration topologies and timing architectures**
4	Backplane Design Considerations	Issues relevant to backplane layout, including distributed capacitance, transmission line effect, stub length, termination, and throughput.
5	Backplane Signal Driving and Conditioning	Signal driving and conditioning, including power consumption, rise/fall time, propagation delay, flight time, device drive, pin conditioning, live insertion, and incident wave switching.
6	Noise, Cross-talk, Jitter, Skew and EMI	A review of the enemies of signal integrity and high frequency.
7	Transceiver Technologies	Detailed information about the following technologies: TTL-based (ABT, FCT, and LVT); ECL; and GTLP.
8	Mechanical Considerations	Information about mechanical considerations such as backplane chassis/cages and connectors.
9	Layout Considerations	Physical layout of the receiver and driver cards plugged into the backplane, primarily focusing on construction of the physical layer and the configuration of the devices that comprise the cards.

Connectivity

Connectivity refers to how drivers and receivers are connected. In some systems it is necessary to connect multiple backplanes together. This can be done from a card to a backplane or from a card in one backplane to a card in another backplane. Figure 1 shows these connections. The backplanes have a parallel bus type configuration with each card slotted onto its respective bus through a connector.

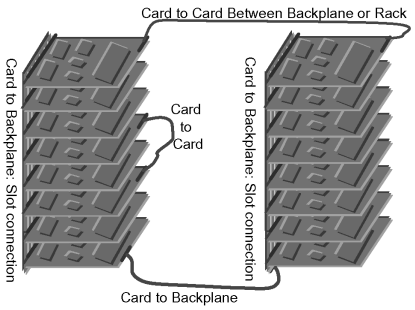

FIGURE 1. Backplane Connection Scenarios

In general, two types of transmission schemes are used in today's backplanes – serial and parallel. These can be differentiated as follows:

- In a parallel transmission scheme, the driver attached to the bus places bits of data in parallel signal lines onto the bus, and all the information is sent along the backplane in either a multidrop or multipoint configuration.
- In serial transmission, all data is sent point-to-point in a serial stream on a single data line. Defined, multiple channels may be used. For the serial technique to achieve the same data rates as the parallel technique, data must be sent at much higher speeds. Serial backplanes can be implemented using special serializer-deserilizer devices that convert a parallel data stream to a high-speed serial data stream and then convert it back to parallel data at the backplane output.

Each design type has the advantages and disadvantages summarized below. When designing a backplane, it is important to select a scheme based on the needs and constraints of the backplane.

Parallel Design

Advantages

- Individual lines can be configured for control signals enabling a fast reaction time and a high degree of system and device control.
- Data throughput can be increased with the addition of parallel data lines.
- High data throughput can be achieved without high data speeds.
- No time delay is required for serializing and deserializing data blocks.
- Parallel design is a common and widely used design type with a significant amount of industry design and implementation information available, including several standards.

Disadvantages

- Significant board and system space can be required, especially with wide or high card count designs.
- Implementation costs can be high because of multiple terminations and devices, wide connectors, and high board trace counts.
- Line-to-line signal skew must be matched and controlled.

Serial Design

Advantages

- Serial designs remove line-to-line signal skew issues in high performance backplanes
- Serial data offers an adjustable approach to data rates, allowing longer cable lengths.
- Fewer ground lines and signal lines are required to transmit the data from one point to another.
- Serial design offers significant space savings over equivalent parallel designs.

Disadvantages

- A time delay is incurred when serializing and deserializing data blocks.
- Serial devices are generally more expensive than parallel backplane drivers.
- Higher speeds require greater attention to signal path layout and impedance matching.

Configuration Topologies

Several signal connection configurations or topologies can be used to transmit data across wires. These include point-to-point, multidrop, and multipoint topologies. Each configuration addresses different issues in backplane design, and each has advantages and disadvantages.

The choice of topology is application-dependent. When selecting a configuration, the designer must be sure that it meets the performance parameters and functionality of the backplane. Each of the three configurations can be implemented with multiple I/O technologies that have been designed specifically for backplane use. Although there are notable variations among these technologies, the most significant difference is between single-ended and differential topologies.

Single-Ended versus Differential Topologies

Single-ended or unbalanced communication means one signal from a driver is carried on one wire or trace to a receiver. The receiver compares the signal to a reference level with respect to ground. In most backplanes, the reference is a standard logic level, although some technologies, such as Gunning Transceiver Logic Plus (GTLP), use a variable reference voltage (V_{REF}).

Differential or balanced communication refers to the signal and its complement. The driver sends this differential signal pair on two matched wires to a receiver, and the receiver compares the two signals with each other. This topology offers noise immunity advantages over most single-ended technologies. Common mode rejection is the ability of the receiver to reject input noise. Because the two differential signals are compared to each other and not to a ground or V_{CC} reference, noise or ground shifts that affect both inputs equally will be ignored by the receiver. Also,

Configuration Topologies (Continued)

most differential technology signal swings are offset from both ground and V_{CC} furthering input noise immunity. Single-ended setups are much simpler than differential topologies because they use half the signal lines, but noise can be a problem. For this reason, some systems benefit from differential signaling because of its superior noise immunity.

The following tables summarize the advantages and disadvantages of single-ended and differential topologies.

Single-Ended Topologies

Advantages	Disadvantages
Less expensive	Susceptibility to noise
Easy implementation	Lower data transfer rates per bit
Fewer components	
Many technologies available	
Higher drive capability available	

Differential Topologies

Advantages	Disadvantages
Common mode rejection	Devices are more expensive
Typically higher data rates	Board space, more components
Lower output swing requirements	Difficult to implement
Generally lower EMI emissions	

Point-to-Point Configuration

Point-to-point configuration is the transfer of data, unidirectional or bidirectional, between two nodes that share a common bus. Data rates, which can exceed 10 Gbit/s, are typically higher than those in a multipoint configuration because there is usually less loading on the bus. Differential technologies are often used in high-speed, point-to-point applications. Point-to-point configurations have the following benefits:

- Straightforward design
- Faster data transfer rates than multidrop configurations
- Each driver operating a single receiver, allowing more consistent receiver control

The major drawbacks of point-to-point configurations are also some of their strong points. The point-to-point connection scheme is not efficient because one driver is dedicated to driving one receiver. This type of connection is best suited for applications that have a low number of receivers.

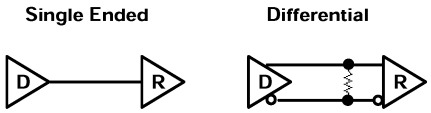

FIGURE 2. Point-to-Point Unidirectional Connections

Point-to-point configuration is implemented with one transmitter and one receiver per line as illustrated in Figure 2. This figure also shows the difference in the setup for single-ended and differential connections.

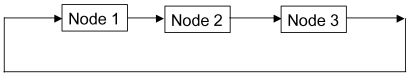

FIGURE 3. Ring Connection

A ring is essentially a point-to-point connection from node to node. A signal starts at one node; that node sends it to the next node; and the next one passes it on to another. Each node scans the signal to check if the message pertains to itself. If it does not, the node passes the signal on. There are a variety of protocols designed for this type of data transmission, but the physical layer is simply a point-to-point connection, as illustrated in Figure 3.

Multidrop Configuration

The multidrop configuration is implemented with one transmitter and many receivers per line. This configuration is not a common design choice because it allows just one device to interface with other cards in a backplane. In this implementation, the receiving devices are not configured to transmit data back to the transmitter or other receivers on the backplane.

Figure 4 shows a typical connection scheme for a multidrop configuration. Multidrop backplanes are generally configured “board-to-board,” with multiple daughter cards linked together through connectors and across printed circuit traces contained on a system motherboard.

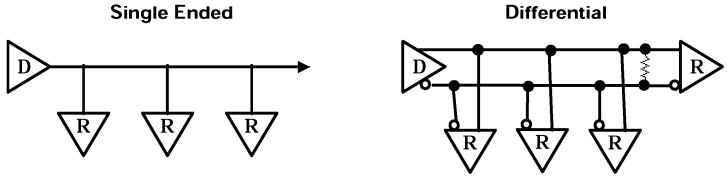

FIGURE 4. Multidrop Unidirectional Connections

Multidrop and multipoint backplane configurations represent worst-case situations from the perspective of signal integrity. Each drop or load consists of a series of potential impedance mismatches that can lead to poor signal integrity and line propagation problems. These mismatches can be caused by drivers, receivers, connectors, stubs, or a combination of all of these. In these configurations, the backplane termination(s) and signal traces must also be factored in. Elimination of impedance mismatches is key to achieving the highest signal integrity.

Configuration Topologies (Continued)

Multipoint Configuration

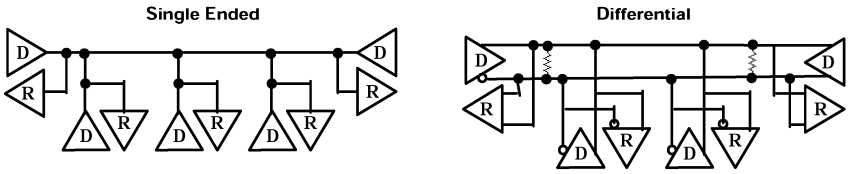

FIGURE 5. Multipoint Connection Using Multiple Transceivers

Multipoint refers to a common bus containing multiple nodes, as illustrated in Figure 5. Multipoint configuration facilitates the transfer of data between the various boards and devices in the system. This is the most efficient backplane configuration because devices connected to the backplane can interface with all others, both receiving and transmitting signals. This characteristic makes the multipoint connection a popular and practical choice for a backplane configuration.

The multipoint configuration is set up as a parallel transmission system, which means that data is simultaneously fed over multiple parallel lines. Under normal conditions, there are a number of bus receivers connected to these lines. Increasing the number of transmission lines increases the data throughput of the system.

Signaling is typically single-ended. This means that data rates, which are typically less than 200 Mbit/s, are usually slower due to noise and loading issues when compared to differential signaling.

One drawback associated with multipoint connections is the decrease in speed. This can be attributed to the large number of capacitive loads that the driver is signaling. The number of loads has a significant effect on maximum backplane speed. For example, with equivalent layout and devices, a 13-slot multipoint backplane has the ability to run at twice the speed of a 21-slot multipoint backplane.

Backplane Configuration Terminology

Terminology used in the industry makes reference to backplane types, including serial or parallel backplanes. In most instances, a reference to a serial backplane denotes a multi-channel, point-to-point backplane.

The transmission is serial with multiple channels, and each trace running from one point to another is considered a channel. In this configuration, a device will be located in the system that performs a switching function, which allows for the communication between many different points in that particular backplane. This device switches one data path to another depending on the address that it senses. When a serial backplane is implemented as a multi-channel design, it is beneficial to restrict the number of channels to a minimum for cost efficiency.

Parallel backplanes can be thought of as one or more single-channel multipoint backplanes. Parallel backplanes incorporate a greater number of data and control lines connected in parallel to all cards. This is usually designed as a single-channel implementation that includes several lines or traces in the single channel.

The desired performance is gained through the use of many signals transmitting in parallel. This is referred to as “data bit width,” and most backplanes are implemented in multiples of eight, with 64-bit widths representing the high end of parallel backplane design.

In most applications, parallel backplane frequencies are slow when compared to those on a point-to-point backplane. However, wide bit widths can often result in the same overall throughput.

The backplane does not necessarily consist of only a single channel. However, a single channel backplane configuration is the typical implementation.

Timing Architecture

Timing architecture refers to the choice between synchronous, source-synchronous or asynchronous signal interfaces, and these three interfaces have many variations. System and backplane signal timing architecture can play a significant role in throughput, cost, and ease of design.

Synchronous Clock

This is the standard architecture used in many low and medium data rate systems. The design is implemented utilizing a single clock source that provides timing for all components requiring a clock. This means that all devices throughout the system receive the clock at the same time. It is accomplished by laying out clock lines of equal length to each device, as is seen in Figure 6. Careful layout, system timing analysis, and minimization of clock skew are critical to this design.

If implemented carefully and properly, this architecture is a robust, effective platform. However, maximum system speed is limited by clock propagation time to the furthest device.

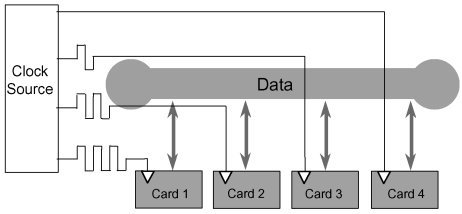

FIGURE 6. A system clock source uses matched trace lengths to generate edges that arrive at all synchronous system devices at the same time.

Source-Synchronous Clock

Source-synchronous architecture uses a clock generated at every driver to toggle the receiving device. This is implemented by laying out a clock line between every transceiver or driver receiver pair that passes data, as is seen in Figure 7. To ensure correct timing, the clock lines must be the same length and have the same loading as the data lines. A system master clock is normally used to re-time received data coming out of backplane boards.

Since timing is now “relative” from each driver to each receiver, much of the system latency is removed as compared to a synchronous system. Data can be passed between system devices without regard to clocking times of other components in the system. This relative timing allows for a significantly faster data throughput when compared to a synchronous backplane.

The benefits of this design are a significant increase in throughput and a potential for larger, practical systems and backplanes because of the elimination of latency and speed issues associated with synchronous designs. The drawback is a significant increase in clock lines – one set for each pair of data-linked devices. This becomes a trade-off entailing either a larger backplane or fewer data lines. However, the potentially significant higher throughput speeds and the capability of producing a system with more devices on the backplane can offset the reduction in data lines.

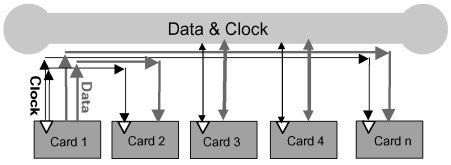

FIGURE 7. Source-synchronous clocks follow the same path as the data and are generated by the transmitting device. Source-synchronous clock architecture minimizes clock-to-data skew and system speed limitations due to clock latency.

Asynchronous Systems

Asynchronous system architecture does not use a clock. Instead of clock-synchronized data, a local “handshake” protocol between the driver and receiver is implemented. There are many different design approaches for this protocol. One example is a request signal from the driver to send data, followed by an acknowledge signal from the receiver when the data is processed.

Asynchronous design can be complex and is often unique for each design implementation. However, this architecture does have some unique benefits. For example, without a clock there are no clock timing or latency issues. This lack of clocking results in a potential increase in data-transfer speeds. In an asynchronous system, only the frequency and length of the data transmission limit data-transfer speeds. Clocked systems have a constantly running clock, regardless of whether data is being transferred or processed. An asynchronous design can yield a potential reduction in power usage due to the lack of a constantly running clock.

Asynchronous architecture can be complex to implement and verify, since traditional timing verification methods do not work. In addition, data streams coming from a backplane of this type still need to be re-timed for system processing. Despite these potential issues, asynchronous architecture has found favor with some designs, such as the VME64 backplane specification.

Summary

The backplane designer has multiple connectivity, topology, and timing architecture choices to make. All design choices derive from system requirements. End system performance compromises should be slanted to achieve optimum performance within the boundaries of the primary design criteria.

At this writing, most implementations are parallel, synchronous, multi-drop designs using single-ended technologies. These remain popular by achieving the primary criteria for backplane design: low cost, high speed, and high reliability. However, as data throughput requirements increase and are coupled with more complex system implementations, more sophisticated solutions are needed.

Serial designs are especially well suited for longer distances. Differential technologies enhance distance signaling, improve signal integrity, and reduce EMI. Source-synchronous clocking has the potential to improve parallel backplane throughput, and an asynchronous design approach holds promise for some implementations.

The end result will almost always be performance compromises of one type or another. In most cases, the watchwords will still be low cost, high speed, and high reliability.