CHAPTER 2

# Why Do We Need Gigabit Serial I/O?

*A review of gigabit serial I/O design advantages*

## Design Concerns

The average design engineer is in a quandary. He would like to stick with tried-and-true solutions because they offer predictability and dependability. But he must also strive for performance improvements in parameters such as data flow, pin count, electromagnetic interference (EMI), cost, and backplane efficiency. Should he consider gigabit serial input/output (I/O)?

## Gigabit Serial I/O Advantages

What is the chief advantage of gigabit serial I/O? Speed. For getting data on and off of chips, boards, or boxes, nothing beats a high-speed serial link. With wire speeds from 1 to 12 Gb/s and payloads from 0.8 to 10Gb, that is a *lot* of data transfer. And with fewer pins, no massive simultaneous switching output (SSO) problems, lower EMI, and lower cost, high-speed serial is the clear choice. Multi-gigabit transceivers (MGTs) are the way to go when we need to move lots of data fast. Let's examine some of the advantages of gigabit serial I/O.

> **MGT:** Multi-Gigabit Transceiver—Another name for multi-gigabit Serializer/Deserializer (SERDES). Receives parallel data and allows transportation of high bandwidth data over a serial link.

### Maximum Data Flow

Some large programmable logic devices have 20 or more 10-Gb serial transceivers for a total bandwidth of 200 Gb/s in and out. While that is an extreme, let's look at an example application that shows us how serial I/O speed can help system architects, board designers, and logic designers.

Figure 2-1 shows a block diagram of a high-definition video mixer.

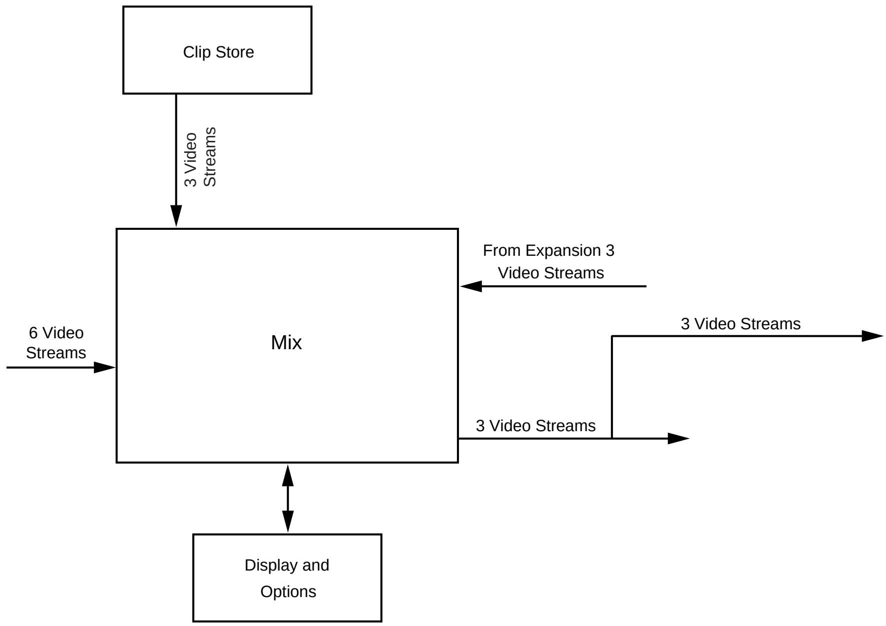


FIGURE 2-1: **High-Definition Video Mixer**

Each high-definition video stream needs 1.5 Gb/s when transported in a baseband or uncompressed format. One scenario for building this system includes discrete deserializer and serializer chips for the serial video streams, and parallel interfaces for the expansion bus and clip store. The other scenario uses gigabit transceivers inside the logic part to decode and encode the serial streams. The faster serial stream acts as the interface to the expansion connector and clip store.

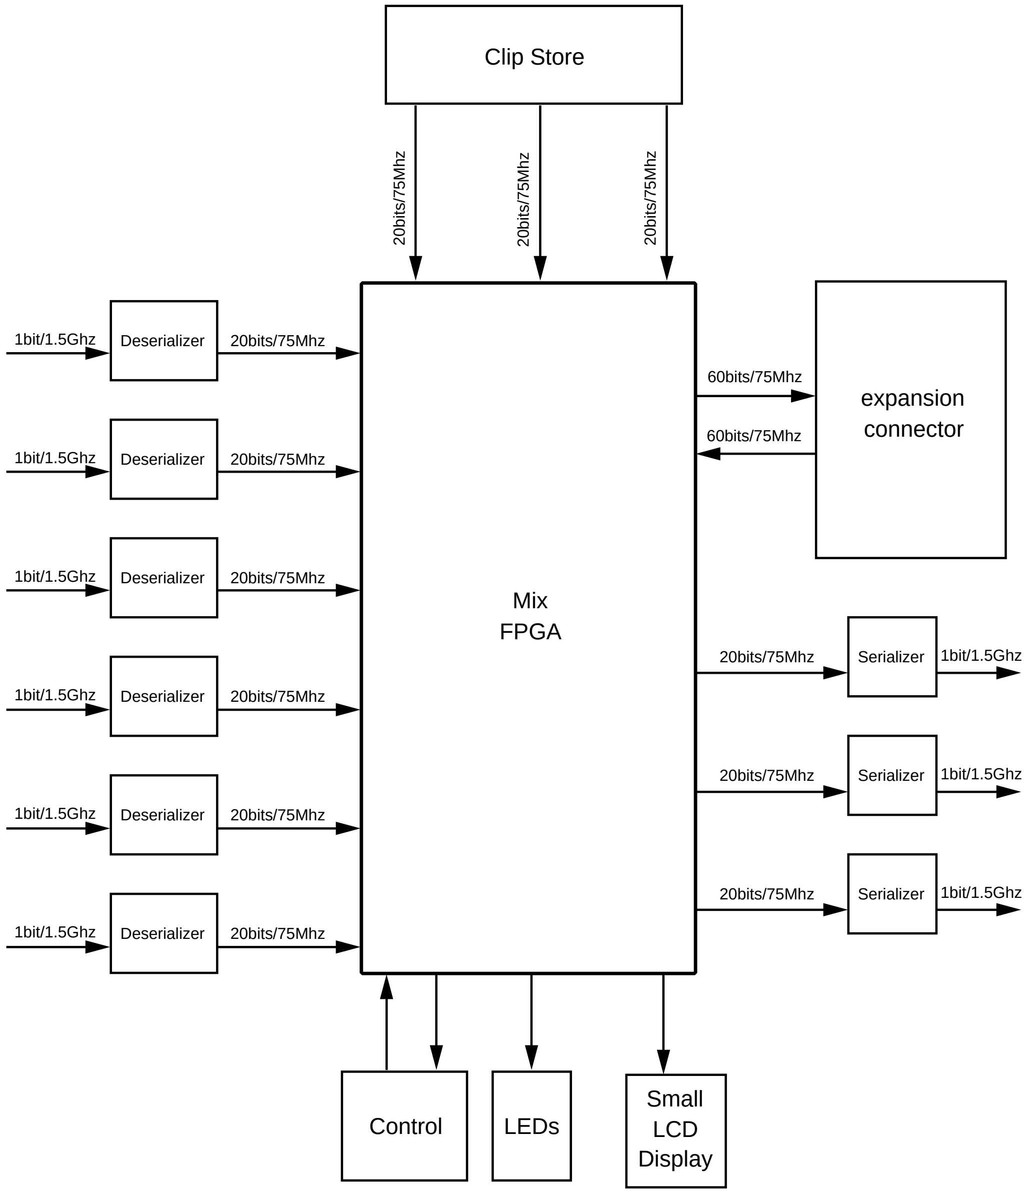


FIGURE 2-2: **Using Deserializer/Serializer Chips (Parallel)**

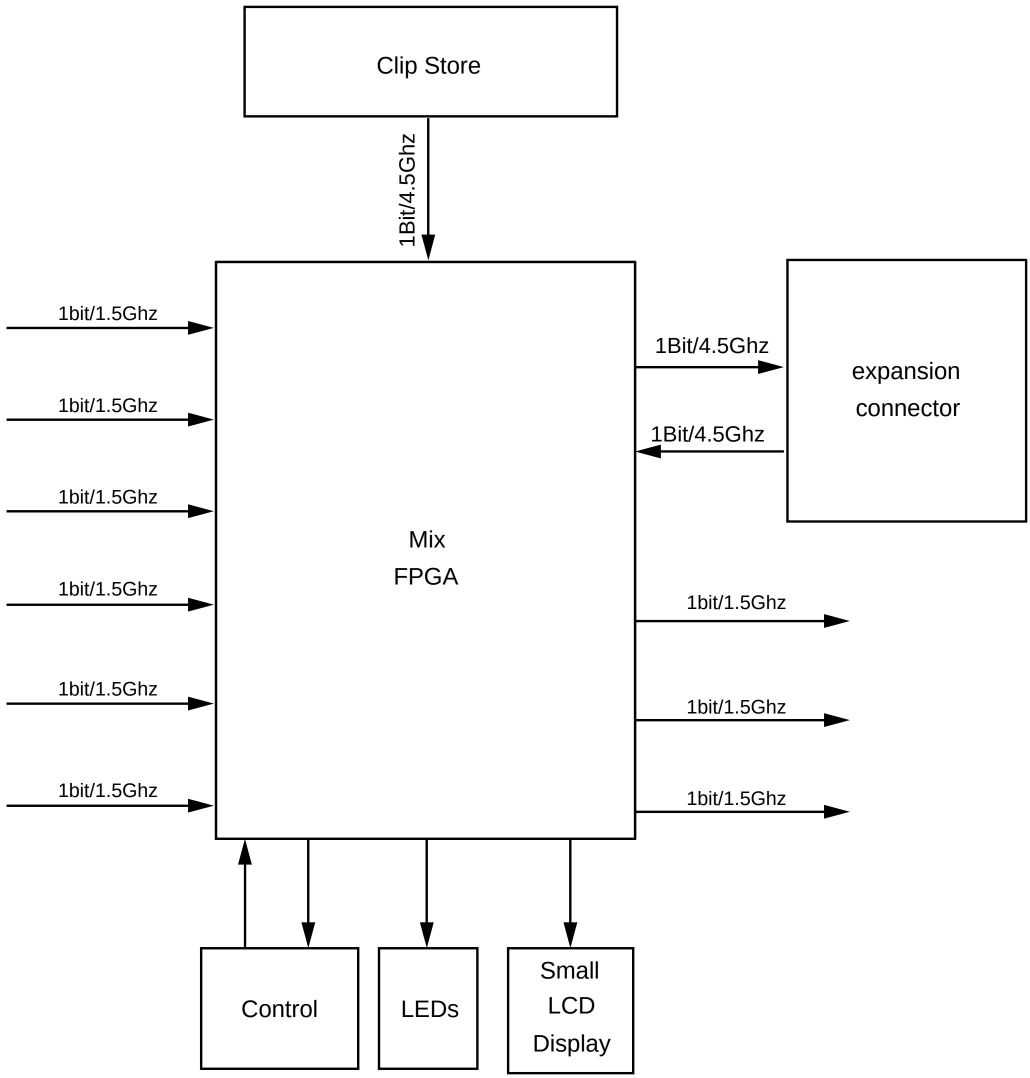


FIGURE 2-3: **Using Gigabit Transceivers (Serial)**

## Pin Count

Pin count is the first problem encountered when trying to move a lot of data in and out of a chip or a board. The number of input and output pins is always limited. Although pin count tends to increase over time, it is never enough to keep up.

The pins needed for our two possible scenarios are given in Table 2-1.

TABLE 2-1: **Pin Counts: Serial vs. Parallel**

| | Direction | Parallel | Serial |
|---|---|---|---|
| Inputs 1-6 | IN | 120 | 12 |
| Clip store | IN | 60 | 2 |
| Expansion inputs | IN | 60 | 2 |
| Expansion outputs | OUT | 60 | 2 |
| Outputs 1-3 | OUT | 60 | 6 |
| Control /status In | IN | 48 | 48 |
| Control /status Out | OUT | 52 | 52 |
| LEDs | OUT | 12 | 12 |
| LCD driver | OUT | 48 | 48 |
| **Totals** | | **520** | **184** |

To be fair, there are some pin issues for which we are not accounting. For example, some MGTs need more power and ground pins than a pair of slower pins. And a parallel interface may require special reference pins. But this example is close enough for a comparison.

Board design time and costs can go up dramatically when a large number of pins are used. Connector pin count is also extremely important for connector/cable selection and feasibility. And using all available ball grid array (BGA) pins might not be convenient.

### Simultaneous Switching Outputs

A designer should consider SSO when using single-ended parallel buses. However, some of those outputs are going to toggle at the same time. When too many switch simultaneously, ground bounce creates a lot of noise.

A designer could also employ differential signal processing on all I/O to get rid of the SSO problems, but that doubles the pin count. And if the data flow needs are more modest, the designer could use a parallel interface with a usable pin count.

## EMI

Experience has shown that as clocks get faster, emissions testing gets more difficult. Hence, gigabit design may seem nearly impossible. But a high-speed serial link will usually exhibit less radiated emissions than a large bus that moves at a slower rate. This is because functioning gigabit links require excellent signal integrity. As one expert put it, “Radiated emission problems are really just signal integrity problems.”

## Cost

Using MGTs will often result in lower overall system costs. With a smaller, cheaper package, the connectors can have fewer pins and the board design may be simpler as well. In the video mixer application, the parallel solution had nine more ICs (integrated circuits) than the serial solution. In this example, the cost of the serial solution is hundreds of dollars less than the parallel solution.

### Predefined Protocols

Another benefit of using MGTs is the availability of predefined protocols and interface standards. From Aurora to XAUI, designs already exist for many different needs.

## What are the Disadvantages?

Before we think that gigabit serial I/O sounds too good to be true, let's look at the downsides. In our designs, we must first we must pay close attention to signal integrity issues. For example, one vendor reported a 90% failure rate on their first attempt with high-speed, multi-gigabit serial designs for a particular application. To improve the odds, we might need to perform analog simulations and use new, more complex bypassing schemes. In fact, we may even need to simulate and model the bypassing scheme.

We can also expect to pay more for impedance-controlled PC (printed circuit) boards, high-speed connectors, and cables. We will have to deal with complications and smaller time bases in digital simulations. And when taking advantage of a predefined protocol, we must plan time for integration and extra gates or Central Processing Unit (CPU) cycles for protocol overhead.

## Where Will Gigabit I/O Be Used?

Initially, gigabit SERDES was confined to the telecommunications industry and to a few niche markets such as broadcast video. Today, MGT applications appear in every section of the electronics industry — military, medical, networking, video, communications, etc. They are also being used on printed circuit board (PCB) assemblies through backplanes and between chassis. MGTs are critical to the future of electronics. Here is a sample of the industry standards that use multi-gigabit SERDES:

- FiberChannel (FC)
- PCI Express
- RapidIO Serial
- Advanced Switching Interface
- Serial ATA
- 1-Gb Ethernet
- 10-Gb Ethernet (XAUI)
- Infiniband 1X, 4X, 12X

### Chip-to-Chip

SERDES was initially used to talk box-to-box. But it exploded into the marketplace because of how nicely it handles chip-to-chip communication on the same circuit board. Chip-to-chip communication had previously been almost exclusively a parallel domain. The amount of logic needed to serialize and deserialize far outweighed any savings that come from pin count reduction.

But with deep sub-micron geometry, an incredible amount of logic can be achieved in a very small amount of silicon. SERDES can be included on parts for a very low silicon cost. Add to that the ever-increasing need for I/O bandwidth, and SERDES quickly becomes the logical choice for moving any significant amount of data chip-to-chip. Consider the following benefits of SERDES chip-to-chip communication:

- **Pin Count:** Smaller, cheaper packages.
- **Pin Count:** Fewer layers on PCB assemblies.

- **Smaller Packages:** Smaller, cheaper boards and more compact designs.
- **SSO:** Fewer pins and differential signaling eliminate the SSO problem.
- **Power:** Usually a high-speed serial link will use less power than a parallel link. This is especially true of some of the actively biased/terminated high-speed parallel standards like high-speed transistor logic (HSTL).
- **Control Lines included:** Often a parallel interface needs a few lines for control and enable in addition to the data lines. Serial links have enabling and control capabilities built into most protocols.

## Board-to-Board/Backplanes

Although they were once the best available, parallel architectures are at their limits. Most parallel bus protocols have evolved to the point where adding data bits is physically impractical because of pin counts on connectors. Clock skew, data skew, rise and fall times, and jitter limit the ability to increase clock frequency. Doubling the data rate can help, but it often requires moving to differential signaling, and that drastically increases pin count. Also, controlling the cross-talk issues on parallel buses is difficult.

New serial backplanes are somewhat different than parallel backplanes. They typically have dedicated serial links from each node to every other node. Figure 2-4 illustrates the basic architecture of an old parallel bus and a new serial bus.

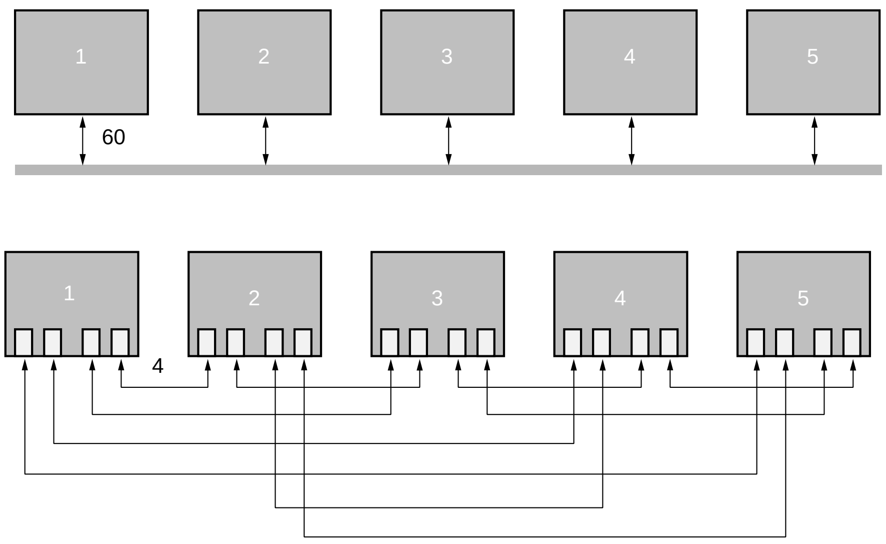


FIGURE 2-4: **Old Parallel Bus vs. New Serial Bus**

Serial bus architectures have a lot to offer. The pin count of a serial bus is a function of the number of nodes. For most practical node numbers, a serial architecture has fewer pins than the old parallel architectures.

Perhaps the most important difference between the two is the bandwidth access method. In the parallel architectures, one node can transmit to one or many nodes. But while that node is transmitting, all other nodes are blocked. All nodes share the available bandwidth.

In serial buses, each node has a dedicated link to every other node. So one node can talk to one or to all nodes while another node is talking. In fact, all nodes can talk to all other nodes at the same time. Of course, the nodes will have to have First In First Out (FIFO) buffering and storage so they can process all the information being received.

Serial bus structure advantages include:

- Higher bandwidth
- Reduced pin count
- Detected bandwidth node-to-node (no need to share)
- Solutions are built into the SERDES
- Easily supports protocols

### Box-to-Box

While SERDES got their start connecting boxes, many designers do not consider multi-gigabit serial links for box-to-box communication. A common misconception is that box-to-box communication cannot be fast without using fiber optics. However, there are many links that go box-to-box for short distances over copper cabling systems. One standard that uses these links is Infiniband. The Infiniband spec allows for 1, 4, or 12 channels of serial data at 2.5 Gb/s per stream. The standard has been commercially available for a number of years and includes cables, connectors, and a protocol that are all well defined and tested.

FIGURE 2-5: **Box-to-Box Connection**

## The Future of Multi-gigabit Designs

At first glance, multi-gigabit communication seems to impose unacceptable restrictions. Serial designers must contend with signal integrity, smaller time bases, and possibly the need for extra gates and additional CPU cycles. However, multi-gigabit advantages in box-to-box and chip-to-chip communication far outweigh the perceived shortcomings. For example, high speed, fewer pins, lower EMI, and lower cost make it the ideal choice in many communication designs. These advantages will ensure its continued use in communication applications far into the future.