



Using 16:1 Serialization to Achieve Signal Rates Exceeding 20 Gbps

By Robert Schwanke, Senior Applications Engineer, Inphi Corporation

High-speed serial signaling has become a mainstay technology in a wide range of video, computing, telecommunications, data communications, satellite, military, and aerospace applications. Whether developers are using industry standards or proprietary implementations, serial signals operating at ultra-high-speed data rates up to 20 Gbps and even as high as 50 Gbps provide the most cost-effective approach for achieving the ever-increasing data rates these applications demand.

Combining multiple channels of signals into a single serialized signal enables the robust and reliable transmission of data across densely packed PCBs, over chassis backplanes, through fiber-optic cables, and through free space via satellite, optical or microwave transmission. This is achieved using a parallel-to-serial (P/S) converter, or serializer, that converts an 8- or 16-bit wide parallel data bus into a single, serial-bit signal. The resulting signal enables a significant increase in serial link capacity and overall system throughput.

To reach ultra-high-speed rates, signals are combined in two stages. The low-speed, input signals are typically generated by an ASIC or FPGA. Four of these low-speed signals are combined in each first stage serializer to create a mid-speed signal. In the second stage serializer, the four mid-speed signals are combined to create a single high-speed signal.

The reality is that high-speed signal design becomes significantly more challenging as signal frequency increases. Designers of yesterday's 1 GHz designs cannot approach today's designs with the same assumptions. The effects of inferior transmission line design, connector reflections, and slow edge rates serve to increase jitter and introduce ambiguity between the different timing relationships across a system. In order to generate reliable multi-Gbps signals using serializers, developers must carefully consider several key issues, including managing clock phase margin, determining what type of I/O buffer technology to employ, selecting the most appropriate device process and packaging technologies, deciding between narrowband or broadband operation, and maximizing architectural flexibility.

Margin for Error

Clock phase margin (CPM) plays an important role in developing a reliable communications channel. Clock phase margin equals the data bit (clock) period minus the combined setup and hold time. The lower the clock phase margin, the more of the overall timing margin of the system that is consumed by the device and the more developers will have to concentrate on recovering timing margin in other parts of the system (i.e. by optimizing jitter and precisely routing signal lines to achieve optimal timing).

The effect of CPM on timing margin can be illustrated by comparing a 4:1 standard digital multiplexer (mux), the equivalent of a single-pole, quad-throw, or 1-pole 4-way, switch, to a 4:1 serializer with a high CPM. The timing margin of the 4:1 mux is affected by several timing parameters: input low-speed data jitter, input-to-input low-speed data skew, select signal jitter and optimum data to select line phase margin. The 4:1 Mux/switch selects between the four low-speed data input signals, using two select lines, in such a way that the data on the selected input passes through the digital mux directly to the data output, thus propagating the same output data rate, timing and jitter as the individually selected signal. Because the input data is not latched within the digital MUX, it must be maintained throughout the desired bit time. This also means that the select lines require a very accurate, external, control circuit.

For example, the four input signals a 4:1 MUX brings in may match in frequency because they come from the same ASIC and are generated by the same clock, but, they may still be out of phase with each other due to ASIC output skew and transmission line trace lengths. Because a MUX requires each data input to be valid for the period of time during which that input is selected (greater than 1/4X the input data rate, equivalent to less than 270 degrees CPM), developers must make sure that each input must be optimally phase-aligned to the proper select state to be valid for the proper data sequence. As a consequence, developers need to manage multiple timing relationships since the four input data lines to the mux must be in phase with the proper select state of the two select lines for the mux.

In contrast, a serializer is designed to latch the four low-speed data inputs and output them sequentially using a single high-speed clock input. This alleviates timing constraints for a high CPM serializer because the required setup and hold time is low which allows the line-to-line skew and jitter of the four input data lines to use up a relatively large



portion of the low-speed bit time. Also, the single high-speed clock is divided down internally to generate the necessary select signals for the on-chip multiplexers. This architectural chip feature serves to bring most of the critical timing onto the chip where it is better controlled. Thus, the only critical timing relationship that is external to the chip is that of the set up and hold time of the serializer which, for 20 Gbps serializers, is on the order of 7 ps (or 347 degrees CPM for the second stage serializer), total, and therefore is relatively easy to line up. This eliminates the need for developers to worry about precisely skewing lines to bring the data and clock into phase. Furthermore, note that the set up and hold time is independent of the input and output signal rates, resulting in even greater clock phase margin (357 degrees CPM or better for the low-speed, first stage serializers).

Differential and Process Technology Options

For high-speed applications, differential signals are far more robust than single-ended signals because they offer common mode noise rejection. There are several differential I/O buffer technologies, however, that offer different advantages and disadvantages at higher speeds.

Low Voltage Differential Signaling (LVDS) is one of the lowest power options available to developers. However it suffers from small output amplitude swings which reduce the signal-to-noise ratio (SNR) since noise stays the same regardless of the output amplitude. Emitter-Coupled Logic (ECL) is another alternative but it requires undesirable off-chip termination; developers want to connect chips together as directly as possible while avoiding impedance matching and imbalanced rise/fall times issues. A third option, Current Mode Logic (CML), achieves an excellent balance of power consumption and a high SNR that is well-suited for ultra-high, multi-Gbps signal rates.

Regardless of the signaling technology used, signal integrity can be improved in several ways. By bringing signal termination onto the chip, the best impedance matching for the least effort and cost is possible. On-chip power supply decoupling protects the signal from noise in the power supply. Having a short ground current return path further reduces noise and jitter and must be implemented both in the serializer and board design.

New process technologies, such as Gallium Arsenide (GaAs), sub-micron Silicon CMOS, Silicon Germanium (SiGe) and Indium Phosphide (InP) are available to facilitate ultra-high-speed signal rates. Gallium Arsenide (GaAs), Silicon Germanium (SiGe), and more recently, Silicon CMOS have served well as the foundation for serializers ranging from



1 to 10 Gbps. Sub-micron Silicon CMOS offers the advantage of low power but its typically narrowband PLLs/DLLs, higher jitter and high development (mask set) cost makes it infeasible for all applications other than the high volume demand of standardized high-speed serializers. While Silicon Germanium typically has the benefit of access to low power CMOS gates for extra features, it gets its speed from the higher powered, bipolar transistors which also tends to have more jitter than GaAs and InP technologies. To achieve very low jitter at data rates of 20 Gbps and as high as 50 Gbps, technologies such as Indium Phosphide (InP), which is the fastest semiconductor technology in production today, are required. Because InP dies can be developed with fewer mask sets than silicon-based technologies, it is not only the highest performance but also the lowest cost and most robust technology available.

With current state-of-art packaging technology, 25 Gbps appears to be the point at which surface mount packaging technology begins to present significant problems. While these issues will more than likely be resolved within the next decade, today's serializers that exceed 25 Gbps require more robust packaging, such as a channelized metal package using coaxial connectors which provide better impedance matching. For this reason, 16:1 serializer implementations often make use of four slower serializers (i.e., 20 Gbps or less) for the first stage and a single higher-speed serializer supporting 40-50 Gbps as the second stage. Of course, if the final serialized signal is less than 20-25 Gbps, the second stage can utilize a slower serializer as well.

Architectural Flexibility

Narrowband refers to a serializer that accepts input signals only within a very narrow frequency range. For applications using standardized serial signaling, a narrowband device might be sufficient as the signal rate is relatively fixed. Many high-speed applications, however, may find it difficult to keep within the tight constraints of a narrowband signal. For these applications, a broadband device that can operate over a continuous frequency spectrum is more appropriate. For example, a narrowband 10 Gbps serializer may have an optimal frequency range from 9.9 to 11.5 Gbps. A corresponding broadband serializer might support from 100 Mbps to 20 Gbps.

The primary difference between narrowband and broadband serializers is the clock source. Narrowband devices often use clock multipliers, prone to added jitter, to create the appropriate high-speed clock frequency while broadband devices utilize a low-phase noise



external clock, extending developers full control over the signal rate at all times. Furthermore, broadband serializers are architecturally less complex than narrowband serializers.

In addition to serializing high-speed signals, serializers are also extremely useful for developing pulse pattern generation test equipment for high-speed applications. For many new high-speed communications standards, developers must begin design before device samples or corresponding test equipment is available. For proprietary applications, there may never be appropriate off-the-shelf test equipment available.

Broadband, multi-Gbps serializers enable the generation of ultra-high speed pulse patterns at any data rate using multiple channels of existing pulse pattern generators and combining them into an even higher rate, serialized signal. Serializers enable developers to carry over in-house testing systems to next-generation architectures. For example, tools developed to test PCIe Gen 1 can be leveraged to test PCIe Gen 2 systems by increasing the signal rate by up to 4X using a single 4:1 serializer. In this way, developers can create arbitrary digital pulse patterns for early device testing without abandoning the tools they already have.

Serializers also present an alternative to using cutting-edge FPGA technology in applications where reliability is important. Many developers are understandably shy about driving FPGA transceivers close to their specified limits, for a number of reasons. Consider upgrading an existing architecture to HyperTransport 3. In order to operate at the specified rate of 3.125 GHz [confirm], developers would have to upgrade any underlying FPGA architecture as well. Such a modification is not trivial, as the FPGA may introduce significant architectural differences.

Rather than move to a new FPGA architecture, developers have the option of staying with the same FPGA running at a lower signaling rate which is then combined using a serializer to achieve the desired output signal. For example, rather than working with a single 3.125 GHz signal and struggling with all the high-frequency issues that come with it, developers can work with four ~800 MHz signals serialized into the final 3.125 GHz signal.

This approach also has the advantage of enabling developers to work with lower frequency FPGAs which are more cost effective per gate to use than cutting-edge higher frequency FPGAs. Another advantage of working with serializers is that developers are able to achieve nonstandard signal rates, such as 6.4 Gbps, which may be optimal for particular



applications. The final architecture is also more flexible and doesn't lock developers to a particular implementation or vendor.

The value of such flexibility cannot be overestimated. Broadband serializers enable developers to test devices over a wide range of potential real-world signal rates. For example, while a system might support up to 3.125 Gbps, it may be advantageous to support a low-speed mode at 400 Mbps. Such a system has the advantage of being able to support legacy signal rates by using training or handshaking mechanisms to negotiate from lower- to higher-speed modes.

Serialization technology enables developers to take high-speed signaling to the next level through multistage architectures which provide the means for stepping up signals. While maintaining signal integrity requires careful consideration and design, developers who understand the key issues behind achieving multi-Gbps signals through serialization give their systems an outstanding competitive edge.

Inphi Corporation (www.inphi-corp.com) has developed two ultra-high-speed, multi-Gbps serializers. The 2080MX broadband serializer covers the operating range of 100 Mbps to 20 Gbps and is packaged in an inexpensive, 32-pin, 5 X 5 mm, plastic QFN package. The 5080MX broadband serializer covers the operating range of 1 Gbps to 50 Gbps and is packaged in a 17-pin, 1 in. sq., GPPO-connectorized, metal package. Both devices operate from a +3.3V power supply, utilize superior, differential, CML signaling and can also be purchased in die form.
