White Paper

Compromises of Using a 10-Gbps Transceiver at Other Data Rates

Introduction

Many applications and designs are adopting clock data recovery-based (CDR) transceivers for interconnect data transfer. Designs and protocols are standardized on data rates between 1 and 3.2 Gbps for the current data node. Protocols such as PCI Express, RapidIO, and Gigabit Ethernet all currently use these data ranges. As data requirements continue to rise, designers are turning their attention to next generation architecture, looking for solutions capable of supporting higher data rates, but over legacy systems and backplanes. This challenge is more difficult because of properties of FR4 printed circuit boards at high frequency, such as skin effect and dielectric abortion, which can cause severe degradation of signal quality.

A new generation of products is emerging that addresses the surge in the 1- to 3- Gbps market and enables the move to the next generation transceiver data node. Many standards bodies, including the Peripheral Component Interconnect Special Interest Group (PCISIG) and the Optical Internetworking Forum (OIF), believe the next logical node step for transceivers to be between 5 and 6 Gbps. However, some applications are starting to implement 10-Gbps interconnectivity (currently, more so for line applications than backplane). Considering the complexity of transceiver design, the challenge facing transceiver manufacturers is how to cater to all applications.

Established transceiver manufacturers favor dual transceiver product philosophy to support all applications. The manufacturers provide a generic transceiver to support applications between 500 Mbps and 6.375 Gbps and a more dedicated transceiver to support a data range with data rates between 10 to 11 Gbps.

This paper describes why a data rate transceiver that is too wide comes with system compromises and can introduce factors that may not be acceptable for some applications operating at lower data rates, where power and cost are generally most important. It also explains why good jitter performance at high data rates does not necessarily translate to good performance at lower data rates; specifically, data rate margin does not equal good signal integrity.

The point is that there are no easy compromises when selecting the relative transceiver data range. A transceiver optimized for 10 Gbps cannot deliver low power and low jitter at a low-cost at 3.125 Gbps. Rule-of-thumb assumptions that good performance at 10 Gbps provides extra margin at 2.5 Gbps are incorrect.

Transceiver Area

FPGAs must provide solutions for wide ranging applications and their requirements. That is, devices must be feature-rich and provide for many eventualities or offer flexibility in the architecture to cope with changes in industry. As an extension of this, the transceiver must provide the same capabilities. A transceiver designed to operate across a wide bandwidth must therefore provide for the requirements for protocols across its data range, otherwise it becomes redundant for certain applications and is not viable as a product.

A transceiver designed to operate up to high data rates must support applications ranging from 155 Mbps to 11.1 Gbps. This encompasses SDH/SONET with its exacting jitter requirements, standard protocols such as Gigabit Ethernet and PCI Express with 8B/10B encoding and subtle signaling requirements, and 10-Gbps Ethernet with its requirement for 64B/66B encoding techniques. The cost of providing this level of support is increased die area. Transceivers operating at the extreme data rates are particularly susceptible to this because the higher data rate protocol features become more complex, leading to the use of more transistors and a need for higher overall operational system clock.

The following are typical features required to support high data rate applications:

- Complex PLL implementations for robust jitter performance (see "PLL Considerations")
- New encoding schemes to support higher data rates (for example, 64B/66B)
- New schemes for signal integrity (if the device is to stand any chance of operating on standard PCB fabric)
- Faster bus interfacing into FPGA fabric

Increased die area has some negative effects:

- Power dissipation
- Device cost

Power

An important consideration when selecting a transceiver is the power dissipation. FPGAs with embedded transceivers are generally used in environments where power and heat dissipation is imperative. The transceivers are often located close to backplane interconnects or chassis panelling, where forced-air cooling is difficult to manage. It is therefore prudent for transceiver manufacturers to generate products that are suitable for a wide range of applications and select a technology and feature set that does not increase power.

PLLs have a big impact on power. The transceivers within the Altera® Stratix® GX device family take up to 30% of the overall power budget of the transceiver. As the data rate widens, the PLL architecture becomes considerably more complex, leading to the use of more logic within the PLL because the PLL must be capable of providing good jitter quality across the entire data rate (see "PLL Considerations"). The additional circuitry is similar to having multiple PLLs within the transceiver architecture. This leads to a significant increase in die area and, therefore, power consumption.

Because the new protocols require greater layers of integration, the new generations of protocols based on higher data rates also have an impact on the digital blocks within the transceiver. Standards bodies are still scoping specifications for future backplane standards, which are likely to support 11.1-Gbps data interfaces. At these data rates, it might be necessary to change encoding schemes from the standard 8B/10B, used in many of today's standards, to schemes such as 64B/66B.

While 8B/10B encoding provides a good solution, it also adds a 25% overhead into the data stream, because every 8-bit character is encoded to 10 bits. At higher data rates, a number of protocols, for example 10-Gbps Ethernet, are moving to 64B/66B encoding schemes. Although they have the same properties as 8B/10B encoding schemes, they require significantly less data overhead. Unfortunately, 64B/66B encoding is relatively complex, so if it is implemented within the transceiver it can take up significant die area within the transceiver itself, thereby adding a further power requirement.

A further consideration with power is the data rate itself. Because

Power = Capacitance × Voltage2 × Frequency

any increasing system frequency has a direct effect on system power. Operating the system at a slower data- or edge-switching rate helps to reduce the total power budget. For example, with Stratix GX devices

1 transceiver operating at 1.25 Gbps = 120 mW

1 transceiver operating at 3.125 Gbps = 200 mW

Based on V_{OD} = 400 mV Pre-emphasis = 0

The same principles follow as data rates increase further. In order to facilitate the higher data rates, designers must design transceivers with the requisite specifications at those higher data rates. From a transmitter PLL perspective, that translates to voltage controlled oscillator (VCO) design and sufficient buffer strength to drive the clock distribution network. In addition, the CDR must be programmable to be wide ranging, depending on the architecture. From the transmitter perspective, the deterministic jitter is partially determined by the intrinsic parasitic and the amount of drive power. However, at the lower data rates, this power consumption becomes a penalty. High-speed systems exhibit sufficient deterministic jitter that present signal integrity issues not only from inter-symbol interference but several other sources internal to the transceiver. Both on- and off-chip considerations to meet faster data rates mean increased power requirements and sometimes additional area consumption that does not offer the customer the most power and

area-efficient solution. Finally, designers should remember that transceiver developments generally move to lower feature sizes (technology) as data rates increase. As a result, it is not obvious that power consumption penalties exist when migrating to higher data rates.

Most applications today utilize transceiver operation between 1 and 3 Gbps. Future road maps suggest many of the next-generation interconnectivity standards will require operation between 5 and 6 Gbps. It therefore follows that most next-generation applications will be captured by a transceiver operating between 622 Mbps and 6.375 Gbps.

Most applications will not require 10-Gbps transceivers until beyond 2008. This is partly because of the slow emergence and high cost of infrastructure, including connectors, development tools, and test equipment and a lack of clarity on backplane standards and transceiver design attributes. Applications using this data rate concentrate on line-side applications that typically use SHD/SONET or 10-Gbps Ethernet-based protocols, which on average require only a single channel within a transceiver. Currently, this type of application is better addressed in a dedicated 10-Gbps transceiver external to the FPGA because of the following:

- It is easier to manage the exhaustive jitter requirements of the protocol specifications.
- The majority of customers that don't require 10 Gbps do not need to make the power and cost trade offs.
- The overall system performance of the FPGA does not become hampered by the bandwidth f_{MAX} requirements of a single transceiver bus interface into the FPGA.

Embedded transceivers within the FPGA can provide a more complete solution at lower data rates, while removing the complexity required to support higher data rates and protocols. This is important, because most of the complexity will be redundant for most applications. Altera's next-generation FPGAs with embedded transceivers will follow this approach.

The transceiver design will be targeted at applications operating between 155 Mbps and 6.375 Gbps. This allows transceiver architects to concentrate on producing a device with exceptional jitter performance across the entire data range. It also ensures that power consumption is manageable. When comparing the Altera solution against a competitor's 10-Gbps solution using actual test chip results, there is a clear difference between the device architected for 6.375 Gbps and the device architected for 10 Gbps. Table 1 shows the power dissipation for Stratix GX, Stratix II GX, and a competitor's solution.

Table 1. Power Dissipation Comparison

Device	3.125 Gbps	6.375 Gbps
Stratix GX	200 mW	-
Stratix II GX *(1)*	125 mW	300 mW
V4 FX *(2)*	400 mW	670 mW

Notes to Table 1:

(1) Numbers anticipated from test chip results.

(2) Numbers taken from device user guide.

On a single channel, the difference of 250 mW between the Altera 6.375-Gbps solution and the competitor's solution is significantly large; multiplied by 20 channels, this is as much as 5 W @ 3.125 Gbps. This is a considerable power penalty for a user who does not need the extra features of the transceiver.

PLL Considerations

Generally, transceivers are designed to transmit and receive data across an imperfect link. The transceiver architect must design a device capable of working at various operating speeds across various operating conditions. Techniques such as pre-emphasis and equalization can be added to the transceiver to help overcome transmission line losses, which are a major cause of inter-symbol-interference (ISI) or deterministic jitter. However, the PLL must also be specifically architected to manage transmit jitter and random jitter components. This management is increasingly difficult across wider data ranges, and although possible, can lead to an increase in complexity and the die area of the PLL.

Transmit PLL

Achievable bit error rate (BER) largely depends on the quality of the transmitted data. Two major influences on data quality are the deterministic jitter components discussed earlier and the random jitter components. The random jitter components seen at the near end of the link (at the transmitter) are primarily associated with the transmitter PLL. This random jitter component can be controlled for a specific data rate by designing the transmitter PLL jitter generation to be minimized for that specific data rate. This is managed by carefully designing the filter components of the PLL for the given range. The filter bandwidth is limited, so if the PLL is pushed to operate over a wide area, the upper and lower data rates in the range see more jitter

Figure 1 shows the characteristics of a PLL when the VCO is optimized to operate at 6.375 Gbps. The normalized jitter generation becomes worse as the data rate is increased above the optimal data rate because a higher percentage of the signal's unit interval is made up of noise. This behavior is also apparent in the wide data range VCO design's phase noise, which also contributes to the overall jitter.

Figure 1. PLL Random Jitter Versus Data Rate

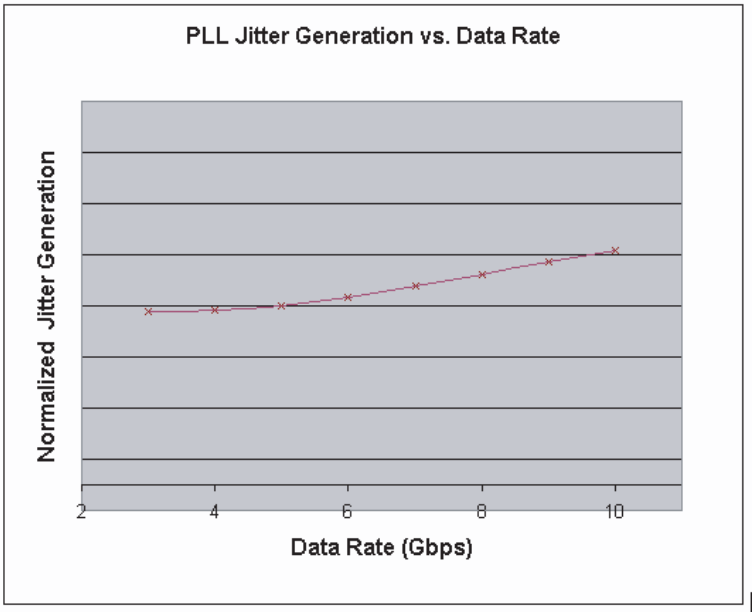

This data range can be widened in the PLL design by using multiple bands (or filters) within the PLL across a number of "specific" narrow frequencies or data rates. The narrow bands are configured inside the PLL to provide the full data bandwidth of the PLL. The bands are applied by changing divider ratios inside the PLL or by providing additional VCOs, both of which add to the die area of the transceiver.

The banding method is beneficial to the transceiver architecture and allows for a widening of the data rate. However, as the level of banding increases to cover the transceiver range, it can be more efficient to use a multiple PLL architecture. With each PLL covering different data rates, this method reduces the complexity of a single PLL architecture and allows additional flexibility in the type of PLL architecture used.

A side effect of both methods is that the jitter performance of the transceiver can vary across its data range. A transceiver operating and characterized successfully at 10 Gbps uses different PLLs or bandings at, for example, 2 Gbps, so jitter performance will differ. Therefore, results seen at 10 Gbps will not necessarily correlate to a better margin at 2 Gbps because the components of the PLL will be different. Jitter performance can be better managed with a dual PLL structure because it allows different oscillators to be used at various frequencies.

In summary, the same PLL is used at 6 Gbps and 10 Gbps, which can be susceptible to higher jitter when out of optimal speed. Or, two different PLLs can be used, which means performance at one data rate does not correlate to another data rate. In either case, good performance at a higher data rate does not relate to better jitter margin at the lower data rate because the PLL characteristics may be different.

Receiver PLL

The bit error rate also depends on the receiver's ability to recover the data in the CDR under jittered conditions. The CDR cannot track jitter beyond its bandwidth well, so the bandwidth must be sufficiently large enough to meet the most tasking or highest data-rate application's jitter tolerance mask. Unlike conventional PLLs, a property of CDR transceivers is that as data rate decreases, jitter tolerance dips. When noise is present at the input of a standard PLL or a CDR PLL, they both behave in the same way, that is, within the loop bandwidth they follow the noise and outside of the bandwidth they do not. For a standard PLL, the relation between the jitter transfer peaking and tolerance dip is correct as shown in Figure 2. For the CDR, the relation becomes much more complex because the single most important factor at high frequencies is the ability to follow high input jitter slope. However, PLL limitations apply and there is a limit to the slope or bandwidth that the CDR can follow. This translates to large phase differences between the recovered clock and the data, creating bit errors if the jitter tolerance mask is not met for a specific standard. Ideally, there is a margin above the mask to allow for any deterioration of the recovered signal caused across the transmission line.

The CDR for any given data rate experiences differences in the number of transitions in the data stream leading to this phenomenon and, incidentally, less stability in the loop. To make the CDR capable of accepting lower data rates it is again necessary to add more bands to the PLL.

Figure 2 shows the receiver jitter tolerance penalty for operating a high data rate CDR at lower frequencies. Frequency response has been normalized to allow both high and low data rate curves to be shown together.

Figure 2. Receiver Jitter Tolerance Penalty for Operating the High Data Rate CDR at Lower Frequencies

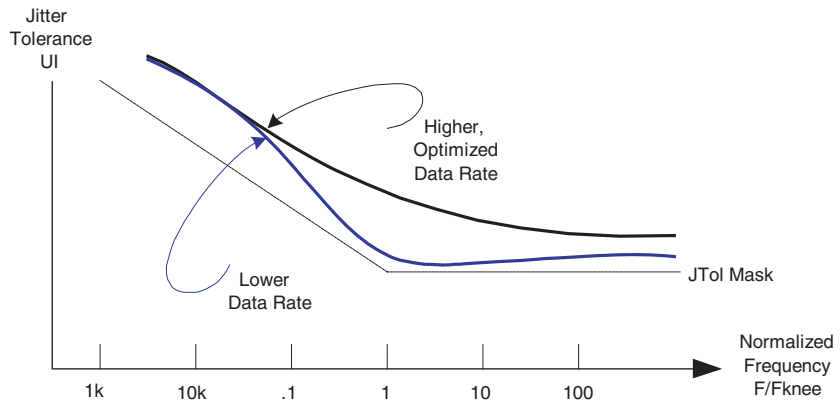

Transceiver Costs

As this paper discussed, supporting wide high-speed data rates requires extra flexibility and functionality within the transceiver so that it is as suitable for operation at 622 Mbps as it is at 10 Gbps. These requirements can significantly increase the area of the transceiver, causing it to take up more silicon die area. Even with shrinking process technology, second- and third-generation transceiver blocks are bigger in die area than their predecessors. The increase in die area directly impacts the cost of the device to the customer.

At 10 Gbps, many of these blocks are necessary, such as when the device is used in a 10 Gbps line-module application. However, a customer using a transceiver for a single-port PCI Express application requires a low-cost solution and may be unable to accept the additional cost of the additional features.

Conclusion

Many applications are now switching to transceiver technology. The mainstream requirement still remains at approximately the 3-Gbps data point. Although attention is turning to next-generation technology, currently most applications are addressed by a transceiver operating at up to 6.375 Gbps. Situations where higher speeds are required can be addressed in specialist devices designed to operate at a specific data rate. Providing a transceiver at a higher data rate does not guarantee delivery of better jitter performance. Jitter is a difficult problem to manage in transceiver design and it is not solved by higher transceiver speeds. The PLL must be architected to perform well across the entire data range, but stretching the range increases the PLL design challenges.

Transceivers within FPGAs can be architected to operate beyond 6.375 Gbps, but the impact on the transceiver complexity results in an increase in die area, which has a major impact on cost and, more importantly, power, even when the majority of applications do not need this extra performance. Altera's solution is to deliver a low-power solution suited to the broad base of applications while at the same time delivering a low jitter solution.

101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
http://www.altera.com