

Going Serial in Gigabit Ethernet Designs

edn.com/going-serial-in-gigabit-ethernet-designs/

EDN

June 20, 2002

With Gigabit Ethernet technology gaining steam in the communications sector, system designers must begin rethinking the choice of interface between media access controller (MAC), physical layer (PHY), and switching components. In a typical Ethernet design, designers would link these components using a wide, parallel bus. While good in Fast Ethernet designs, these parallel architectures fail to keep up with the speed demands of Gigabit Ethernet and emerging 10 Gigabit Ethernet designs.

What's needed is a move toward high-speed serial interfaces. Through the serial GMII (SGMII) and the XAUI interfaces, designers are able to address the interface problems that they currently face, while shortening the design cycle and saving on costs.

In this article, we'll explore the impact of implementing GMII interface in a Gigabit Ethernet design. We'll also explore a new version of the SGMII interface that drops overall pin count to four pins. Additionally, we'll look at the XAUI interface and its impact on emerging 10 Gigabit Ethernet architectures.

The Current Situation

The vast majority of Gigabit Ethernet-class switching chips shipped today use the GMII interface technology to link the MAC and PHY components. This is a parallel format, which uses 25 separate traces running simultaneously at 125 MHz between every MAC and PHY pair. In existing applications, the GMII technology has proven adequate to this task, but in upcoming designs it's too cumbersome.

While the first Gigabit Ethernet switching chips supported only a single data port, current designs are much more dense. Four- and eight-port devices are the norm, and 12-port chips are becoming more common. In the near future, the industry is expected to move to 24-port components.

A 12-port design contains 12 sets of MAC and PHY pairs. With GMII technology, this requires 300 traces and 300 separate pins on the switch device just for the interface; other components within the device will require additional pins. This can be an excessive number of pins to place on a single semiconductor. Not only does it place a high burden upon the engineer during the design phase, it also makes each chip more expensive to produce; a serious drawback for components targeted at the very cost-sensitive switching market segment (**Figure 1**).

Figure 1: Designers looking to develop 12-port Gigabit Ethernet chip designs could face up to 300 board traces when interfacing components.

There are other problems that must also be considered. The 300 pins on the chip must be plugged into 300 individual lines on a PCB. That much wiring takes up space, and extra space on the PCB means extra cost. As the amount of nets on the PCB increases, additional layers are required, which significantly increases costs; in general, adding two layers to the PCB will increase the cost of the PCB by 25 percent.

Furthermore, the nature of a parallel connection requires the data packets traveling across all of the traces to depart from one component simultaneously, and to also arrive at the other component at the same time. To satisfy this difficult synchronization requirement, designers must attempt to make each trace the exact same length.

With 300 separate traces weaving through the rest of the design's circuitry, it is difficult to plot enough paths with the same length, and traces of differing length create channel skew in the signal. While it is possible to deskew the signal within the components, this is an extra task for the chip design engineers to solve and requires additional performance within the device.

This is not the end of the design challenges that engineers face. With data packets traveling across 300 individual traces, signal integrity issues become much more significant than in previous designs that had far fewer traces.

Try Fewer Pins

All of the challenges above arise from the simple fact that more highly integrated designs, which support more data ports, require more pins and traces. The integration trend is certainly not going to reverse itself, which means that all of these problems just become more significant over time. But there is a simple way to correct these problems: use fewer pins.

That is exactly the approach that system designers have begun to implement, with new interface formats. One of the earlier attempts to solve this problem was the reduced GMII (RGMII) format.

The RGMII parallel interface is similar to the original GMII interface, but allows designers to decrease the number of pins by increasing the frequency at which data travels across each pin. Instead of 25 traces connecting each MAC and PHY pair, the RGMII format requires 12 traces. Halving the number of traces is achieved by increasing the frequency, from the 125 MHz used in GMII to an effective rate of 250 MHz in RGMII.

This approach offers some advantages, but does not solve all of the problems. Fewer pins reduce production costs, and allow designs to be more competitively priced. This requires fewer lines on the PCB, which will reduce routing congestion and allow chips to be placed

closer together on the board, and therefore reduces overall system costs. And in addition, fewer traces make the task of routing them all somewhat easier, which in turn makes it easier to synchronize all the signals, so data leaves and arrives in synch upon parallel traces.

However, increasing the frequency of each signal results in a definite trade off. The increase is necessary in order to maintain the overall data throughput rates within the device, but running more energy through each trace makes it more likely that each trace will create interference issues upon other traces. Because RMI uses the same LVTTTL signaling technology as GMII, there are no significant changes in the signaling, but the integrity issues are a greater challenge because of the higher frequencies involved.

This problem can be solved in two ways. Engineers can add more space between the traces, although this requires more space for the overall design, and therefore increases the production costs and mitigates some of the advantages of migrating to the RGMII format in the first place. Or, engineers can spend more time fine-tuning the overall design in order to eliminate EMI problems, but adding to the design cycle also leads to a more expensive product.

So, while the RGMII technology is a step forward because it requires fewer pins, it does not go all the way towards solving the interface problems faced by today's engineers. In fact, the increased signal integrity problems that result from higher frequency traffic almost erase the advantages gained from decreasing number of pins.

Entering the Serial World

The next step is to shift from a parallel interface to a serial interface. Cisco has proposed such a technology with a serial GMII (SGMII) format that uses just eight pins and a double data rate technology that enables signals to travel at an effective rate of 1.2Gbaud MHz.

In this format, the eight pins are allocated to four channels, with one pair each for the receive (Rx) data and clock signals, and the transmit (Tx) data and clock signals. Even more significantly, this SGMII format switches the signaling format to LVDS, which greatly improves the signal integrity.

There are several advantages that come from shifting from parallel to serial technologies. While parallel connections allow for very fast data rates over short distances, serial links permit longer connections, and eliminate the need for synchronized signaling. The result is that designers have an easier time planning the route for each trace, which shortens the design cycle and save on costs.

The Cisco version of SGMII also further reduces the number of pins, and this in turn leads to all the cost benefits seen in the shift from GMII to RGMII. There is an additional benefit that comes into play with the migration from a parallel to a serial interface technology. While the

frequency with this SGMII format has increased from the frequency used in RGMII, designers have to solve fewer interference issues because the signals are no longer traveling in synch across the traces and because of the change to LVDS technology.

Moving to Four Pins

However, it is still possible to improve upon the SGMII serial interface technology proposed by Cisco. The Cisco version uses eight pins, but a newer version of the SGMII interface has been developed that cuts that number in half by embedding the clock signal within the data channel. Let's explore this format further.

The proposed SGMII format uses just four pins that are arranged to support two data streams, Rx and Tx, each with a single pair of pins. One pin in each pair is dedicated to the signals moving from the MAC to the PHY, and one pin dedicated to traffic moving in the opposite direction, from the PHY to the MAC. The signals all move at 1.25 GHz.

This format also uses LVDS signaling, so the signal integrity issues are comparable to the Cisco format. However, by reducing the number of pins to just four, this SGMII technology has reached the theoretical minimum number of pins for this type of interface.

The proposed SGMII format has been successfully implemented in devices that also contain integrated serializer/deserializer (serdes) components. This offers the additional benefit of allowing for direct connections between the switching chip and the optical fiber modules because the data no longer needs to be shipped off-chip to a serdes to be translated into the serial format that is required for the optical modules and other sub-systems within the network **(Figure 2)** .

Figure 2: An SGMII interface has been developed that requires only four pins and an LVDS signaling format.

Now, compare this SGMII format with the existing GMII standard. For a 12-port switching chip, the GMII device requires 300 pins; in the same chip, an SGMII device uses just 48 pins. This is a much more manageable number for engineers, who will have a much easier time routing just 48 traces. Furthermore, because the SGMII design uses serial technology, this routing task is even easier because the traces do not need to be the exact same length. And while the SGMII format uses much faster signaling, the risk of signal integrity issues is lower because the data need not travel in synch with data moving across other traces, and because it uses the superior LVDS technology.

Moving to 10E

The interface formats described above apply primarily to switching chip designs in the Gigabit Ethernet markets. However, there is a similar technology that has already been developed for devices in the 10 Gigabit Ethernet class.

The current generation of 10 Gigabit Ethernet components uses XGMII, another parallel interface designed for faster speeds. However, there is already a specification defined for a serial interface that can function at the 10 Gigabit Ethernet level. This technology is called 10 Gigabit Ethernet Attachment Unit Interface, and is generally abbreviated as XAUI.

All of the advantages of switching from a wide, parallel interface to a serial connection at the standard Ethernet speeds also apply to the 10 Gigabit Ethernet designs. Parallel interfaces still require more pins, more traces, more complicated signaling schemes and are more likely to produce interference and signal integrity issues. And in systems moving data at Gigabit speeds, even minor interference problems can lead to the loss of a significant amount of data **(Figure 3)** .

Figure 3: While many developers are turning to the XAUI serial interface, a host of developers are still evaluating the XGMII parallel interface for 10 Gigabit Ethernet switching designs.

The XAUI format uses four serial channels, compared to two channels with SGMII, and the higher system speeds require each channel to utilize four pins, instead of two. That leads to a total of 16 pins connecting the MAC and PHY components of a 10 Gigabit Ethernet chip, with data rates on each trace of 3.125 GHz.

Because of the built-in deskew capabilities found in standard-based XAUI cores, it is not necessary for all of the nets to be the exact same length. In fact, since the data traffic will likely be coming directly from optic fiber connections that could be miles in length, the length of the nets for XAUI-based systems can vary by as much as a foot **(Figure 4)** .

Figure 4: The XAUI interface some equipped with a built-in deskew capability, thus easing the implementation of this interface in a system architecture.

Standardizing the Interface

While the XAUI interface has been adopted as a standard in the communications sector, the SGMII format has not yet been approved as an industry-wide standard. However, with several companies successfully implementing SGMII in current product designs, designers can clearly see that this high-speed serial connection works. This success goes a long way towards proving that it is the next logical step in communications designs.

History indicates that it may well become a de facto market standard. In the standard Ethernet and Fast Ethernet segment, the original interface used was MII. In order to reduce the number of pins, the industry moved to RMII, and later adopted a serial version, SMII.

This is the identical trend now emerging in the Gigabit Ethernet and 10 Gigabit Ethernet segments. All of the advantages that come from adopting a serial interface add up to a faster design cycle and lower costs. Moreover, fewer pins and fewer lines on the PCB mean smaller chips, smaller boards and a lower bill of materials. These are all of the factors that create happy designers, and even more importantly, satisfied customers.

About the Author

Martin Lund is a product line manager at Broadcom. He holds a technical degree from Copenhagen Technical Academy in Denmark and can be reached at mlund@broadcom.com.