

Overcoming Signal Integrity Challenges of 112G Connections

By Cadence

One of the big challenges with 112G SerDes (and, to a lesser extent, all SerDes) is handling signal integrity issues. In the worst case of a long-reach application, the signal starts at the transmitter on one chip, goes from the chip to the package, across a trace on a printed-circuit board (PCB), through a connector, then a cable or backplane, another connector, another PCB trace, another package, and arrives at the receiver. By the time it arrives, the signal is very distorted, making it a challenge to recover the clock and data-bits of the information being transferred.

This white paper looks at how to handle these signal integrity issues and ensure that data is faithfully transmitted with a very low bit error rate (BER). The focus will be on 112G long reach. Many similar considerations apply to shorter reach, and to 56G (and lower data rates), but in some sense, 112G long reach is the most challenging case. While this white paper does not cover the actual design of 112G SerDes silicon, it does consider how connectors are designed.

Contents

112G SerDes IP	2
Signal Integrity	2
3D Analysis	3
Designing Connectors.....	4
Putting Analysis and Layout Together in Practice.....	5
Signal Integrity for 112G.....	5



112G SerDes IP

All semiconductor IP is, to some extent, a “make versus buy” decision. In the case of a 112G SerDes, the answer is always going to be to buy, meaning to license IP. This is due to two especially big challenges. Firstly, only the most highly skilled and experienced design groups are even capable of attempting such a design. Second, any design that complex requires a test chip to validate the SerDes block, adding a huge amount of cost and schedule delay.

Cadence has a 112G SerDes IP design, with power, performance, and area (PPA) optimized for long reach and medium reach. It is designed in 7nm silicon targeted at advanced high-performance computing (HPC) data center systems on chip (SoCs). It supports lower speeds, too—in fact the entire Ethernet range of 10G/25G/50G/100G. For 112G and 56G, it is PAM4 (two bits per symbol, so three eyes), and for the lower speeds it is NRZ (one bit per symbol, so a single eye). It has >35dB insertion loss.

The 112G SerDes IP has fully autonomous startup and adaptation. That is, assuming that all the signal integrity issues that will be discussed later in this paper are addressed, the transmitter and receiver SerDes blocks will synchronize automatically without any special intervention by the rest of the SoC. This is achieved through equalization. The transmitter side is digital-to-analog converter (DAC) based, with four taps of transmitter finite impulse response (FIR) equalization. This adds pre-emphasis and de-emphasis at the transmitter to compensate for the channel.

The receiver side is digital signal processor (DSP) based with several stages of equalization. First, automatic gain control (AGC) boosts the input signal (but also the noise) and centers the input signal. Next is continuous time linear equalization (CTLE or sometimes just CTE), which attenuates low-frequency signal components and amplifies components of the signal around the Nyquist frequency. Then there is a decision feedback equalizer (DFE) and clock data recovery (CDR). This is where the signal is finally processed to recover the stream of data values and the associated clock ticks. The feedback aspect is that every few ticks (just how many is controlled by the designer), the parameters in use are updated based on previous decisions. Drift in time (a 112G signal doesn’t run at exactly 112G) and voltage (the value used to decide if a given voltage is 0 or 1) will vary. The DFE and CDR keep the variance of the signals under control—DFE for the signal-level discrimination, CDR for the distribution of where the clock edges are detected.

Nothing in the rest of this white paper depends on your using Cadence® IP for the SerDes. Indeed, in a typical design, even if you are designing an SoC, the chip you are communicating with may well come from a different supplier and be outside your control. In fact, if you are just doing design at the PCB level, you may have multiple components, all with 112G SerDes interfaces, none of which you control at the silicon level. However, the signal integrity issues are still important, and the design will not work if problems are not appropriately addressed.

Signal Integrity

Perhaps surprisingly, the channel between the transmitter and the receiver can be completely characterized by a single simulation of an impulse response. However, both the simulation and the description of the channel need to be done very accurately.




Figure 1: Clarity 3D Solver EM software optimizes the connector-PCB interface and models critical 3D interconnect from TX to RX

To do this simulation, a step function is put on the input to the transmitter output stage. Then the whole channel is simulated, consisting of that output stage, the channel itself, and the receive input stage. That single simulation contains all the information required to measure the distortion that the channel will introduce. Remember, the channel might be quite complex, with pins, traces, connectors, backplanes or cables, and so on.

It is not possible to use circuit simulation to analyze the transmission of the large number of bits required to get reliable eye diagrams, to calculate the BER, or even to simulate enough bits (hundreds of thousands of bits) for the transmitter and receiver to lock in.

Instead, the channel is modeled using the I/O Buffer Information Specification (IBIS), a nearly thirty-year-old standard for modeling the channel, with the Algorithmic Modeling Interface (AMI). Cadence was involved in developing both specifications from the beginning. They make it possible to simulate very large bit streams with fast and accurate simulation models. The first part of the simulation, before lock-in of the adaptive equalization, can be discarded. Later parts of the simulation can be used to check that the eyes are sufficiently open to comply with the spec, and to calculate the BER.

IBIS is a modeling standard that allows the driver and receiver to be modeled, with pull-up and pull-down transistors, transition times, clamping diodes, and the inductance, resistance, and capacitance of the package pin. The input model is the same but without the driver transistors. Traditional IBIS models were first used to replace SPICE models for much faster simulation and are text based.

AMI allows the text-based form of IBIS to be replaced with compiled code, concealing proprietary information (and running faster). The specifications, used together, are known as IBIS-AMI.

For 112G SerDes analysis, IBIS-AMI is the way to go. Assuming the design is using SerDes IP, then the IP developer should supply the IBIS-AMI models required. There will usually be more than one such model since the SoC being designed has to talk to another chip that may come from a commercial semiconductor vendor. Of course, sometimes the SoC being designed is being used repeatedly and communicating only with another instance of itself.

The group at Cadence that designed the 112G SerDes IP provides IBIS-AMI models, of course. They created them automatically using another Cadence tool, the Sigrity™ Advanced IBIS modeling tools, which allows production-grade models to be produced using a wizard.

During training of the models, to calculate parameters to initialize the equalizers, a backchannel is used that does not exist in reality in order to close the loop between transmitter and receiver. The parameters won't be exact when the system is powered up in reality, since they depend on process corner, temperature, and voltage, and perhaps on variability of other aspects of the design. But if the parameters start closely enough, then the adaptive equalization should lock into good values after a few tens or hundreds of thousands of bits.

3D Analysis

Until recently, a pseudo-3D solution would have been used to model all the interconnect between the transmitter and the receiver. But Cadence's Clarity™ 3D Solver is now the tool of choice. It has gold-standard accuracy, and it has been designed from the start to be ready for large data centers or the cloud, with world-class parallelization technology. As a result, it delivers up to 10X speedup or more compared to alternative approaches.

The analysis by the Clarity Solver uses finite element method analysis (FEM). This is done in two stages, both of them highly parallelized. The first does the adaptive meshing, breaking up the complete system interconnect into the small elements. The second phase then calculates the solution for different parts of the design across a range of frequencies. The entire solution is finally combined to create an S-parameter (frequency response) of the n-port network represented by the structure being analyzed.

Under the hood of the Clarity Solver, there is a massively parallelized matrix solver. This is a breakthrough algorithm and is part of Cadence's secret sauce in the system analysis area. In addition to near-linear scalability without any accuracy loss, it provides virtually unlimited capacity using a large number of low-capacity machines, eliminating the cost of huge machines that are either unavailable when needed or are sitting idle much of the time waiting for jobs to show up to use them. The whole infrastructure is dynamically deployed into the cloud (or a data center) and has fault-tolerant restart since, with huge numbers of machines, rare things happen regularly.

Sigrity technology can be used to analyze the power delivery network. Creating a successful high-performance SerDes-based design depends on having clean power and ground delivered to the chips. Modern chips run at low voltages, but that also means that the currents involved can be very high. Low voltage and high current make resistance losses, which cause voltage droop, especially critical.

Designing Connectors

So far, this white paper has focused on analyzing the signal integrity of the path, with an implicit assumption that the boards and connectors are “good”. The issues of what types of PCB materials should be used for high-performance designs are reasonably well understood. But the connectors themselves also must be designed while taking the signal integrity issues into account—some things simply cannot be fixed up by the equalizers in the SerDes transmitter and receiver. In particular, return loss results in less noise margin at the receiver since the part of the signal that never makes it to the receiver is obviously not useful during equalization there. Crosstalk, interference between separate signals, is something else that cannot be compensated by equalization since it happens at roughly the same rate as the signal (assuming the aggressor is also running at a similar data rate).

The design of connectors has gone through three phases, depending on the amount of compute power available in that phase and the sophistication of the algorithms for the finite element meshing and analysis.

Connector design is a complex problem, since connectors must perform well from a signal loss and crosstalk point of view, yet they are typically manufactured from injection-molded plastic and stamped metal. (This description excludes specialized and expensive microwave connectors.) What makes designing connectors for high-speed SerDes so challenging is that they must be economically manufactured in high volume at a low price—but 112GHz is microwave frequency without access to expensive and bulky microwave connectors.

In the distant past, connectors were basically designed by trial and error by designers with decades of experience and a feel for what would work. However, it could take well over a year, involving four or more prototype cycles where a trial connector would be built, assembled with a test board, and then measurements taken.




Figure 2: High-Speed Connector

Once software tools for analysis were available, this process could be shortened somewhat by creating the connector in a mechanical design system, and then analyzing the resulting structure. There were two limitations on this. Often, getting the connector to mesh correctly and thus be ready for analysis was tricky. And secondly, the amount of computer power required for the meshing and the analysis was very large. However, this could reduce the connector design cycle to six to nine months, needing only a couple of prototypes.

In this era, the connector would be analyzed separately from the reference board, and then the two sets of measurements combined, on the assumption that there was no electromagnetic interaction between the breakout region of the board, sometimes called “the final inch”, and the connector. At low signal frequencies, the interactions were second order and the errors resulting from ignoring them were small.




Figure 3: Connector-PCB interface

For 112G (and 56G) SerDes connections, the signal encoding is PAM4 and high frequency. The signal rate is 56G symbols per second (or 28G) but higher frequencies must also be transmitted cleanly to recover the data and clock given the encoding scheme. However, a lot of connectors are mechanical press-fit and electrically large and so are not very good. The pins are inevitably close to each other, with no (or limited) shielding. Crosstalk is an especially demanding problem—because it is so dynamic, none of the equalizers can do anything about it. Return loss, essentially energy in the signal that is reflected at the connector and never gets to the receiver, is another issue that the connector designer must minimize. Additionally, at these frequencies, the assumption that the connector and board can be analyzed separately and then “added” together is no longer good enough. There is too much interaction between the board and the connector that is not captured in the separate analysis and can no longer be ignored.

The modern approach to connector design is to use a full 3D analysis tool like the Clarity Solver to do as much of the design as possible “on the computer” instead of by building prototypes. The Clarity Solver, with its ability to take advantage of the huge levels of parallelism available in cloud data centers, can significantly reduce the required wall-clock time for analysis, without compromising accuracy, which is still the first requirement. Getting the wrong answer fast is a poor compromise.

Of course, eventually at least one prototype design will need to be manufactured to take real-world measurements, too. This can reduce the connector design cycle to under six months. The Clarity Solver makes it possible to design high-quality connectors and associated reference designs, and to predict accurately how they will behave after manufacture.

Putting Analysis and Layout Together in Practice

A designer putting a real 112G system together needs to operate in several domains at once. The system designer is usually selecting a suitable connector, not designing a custom connector as this would take too long and be too costly.

However, it is still a requirement to use the Clarity Solver to model the connector and the PCB together. This needs to be done at a physical level, since the connector needs to attach to the board, usually with pins going through vias of some sort on the PCB. But it also needs to be done electromagnetically to analyze any signal integrity issues.

The designer will want a board that has as few layers as possible and is built out of the lowest cost materials that meet the signal integrity constraints of 112G. Cadence’s Allegro® technology and the Clarity Solver are linked in a way that allows the optimized PCB structure in the Clarity Solver to be implemented in Allegro technology without having to redraw the PCB structure. This not only makes things easier for the design and signal integrity engineers, but also ensures that the design that is analyzed and the design that is drawn are, indeed, the same—another source of possible error is thus avoided.

Signal Integrity for 112G

The Cadence approach handles the signal and power integrity issues of 112G long-reach designs, calculating IR drop and EM issues accurately, including the ways that the different parts of the signal path interact. This analysis can be used at all stages of the design: the SerDes transmitter and receiver, the packages, the boards, and the connectors.

Cadence’s Intelligent System Design™ approach involves many design tools and technologies working together seamlessly—in this case, the 112G SerDes IP, Sigrity IBIS-AMI modeling, Clarity 3D Solver, Sigrity Signal and Power Integrity Analysis, and Allegro technology for board and package design and analysis. This combination of tools brings Cadence’s expertise in computational software to bear on designing successful systems running at the fastest current serial link speeds, even in what seem like hostile environments involving connectors, cables, and backplanes.

cadence®

Cadence is a pivotal leader in electronic design and computational expertise, using its Intelligent System Design strategy to turn design concepts into reality. Cadence customers are the world’s most creative and innovative companies, delivering extraordinary electronic products from chips to boards to systems for the most dynamic market applications. www.cadence.com

© 2020 Cadence Design Systems, Inc. All rights reserved worldwide. Cadence, the Cadence logo, and the other Cadence marks found at www.cadence.com/go/trademarks are trademarks or registered trademarks of Cadence Design Systems, Inc. All other trademarks are the property of their respective owners. 14246 04/20 SA/RA/PDF

