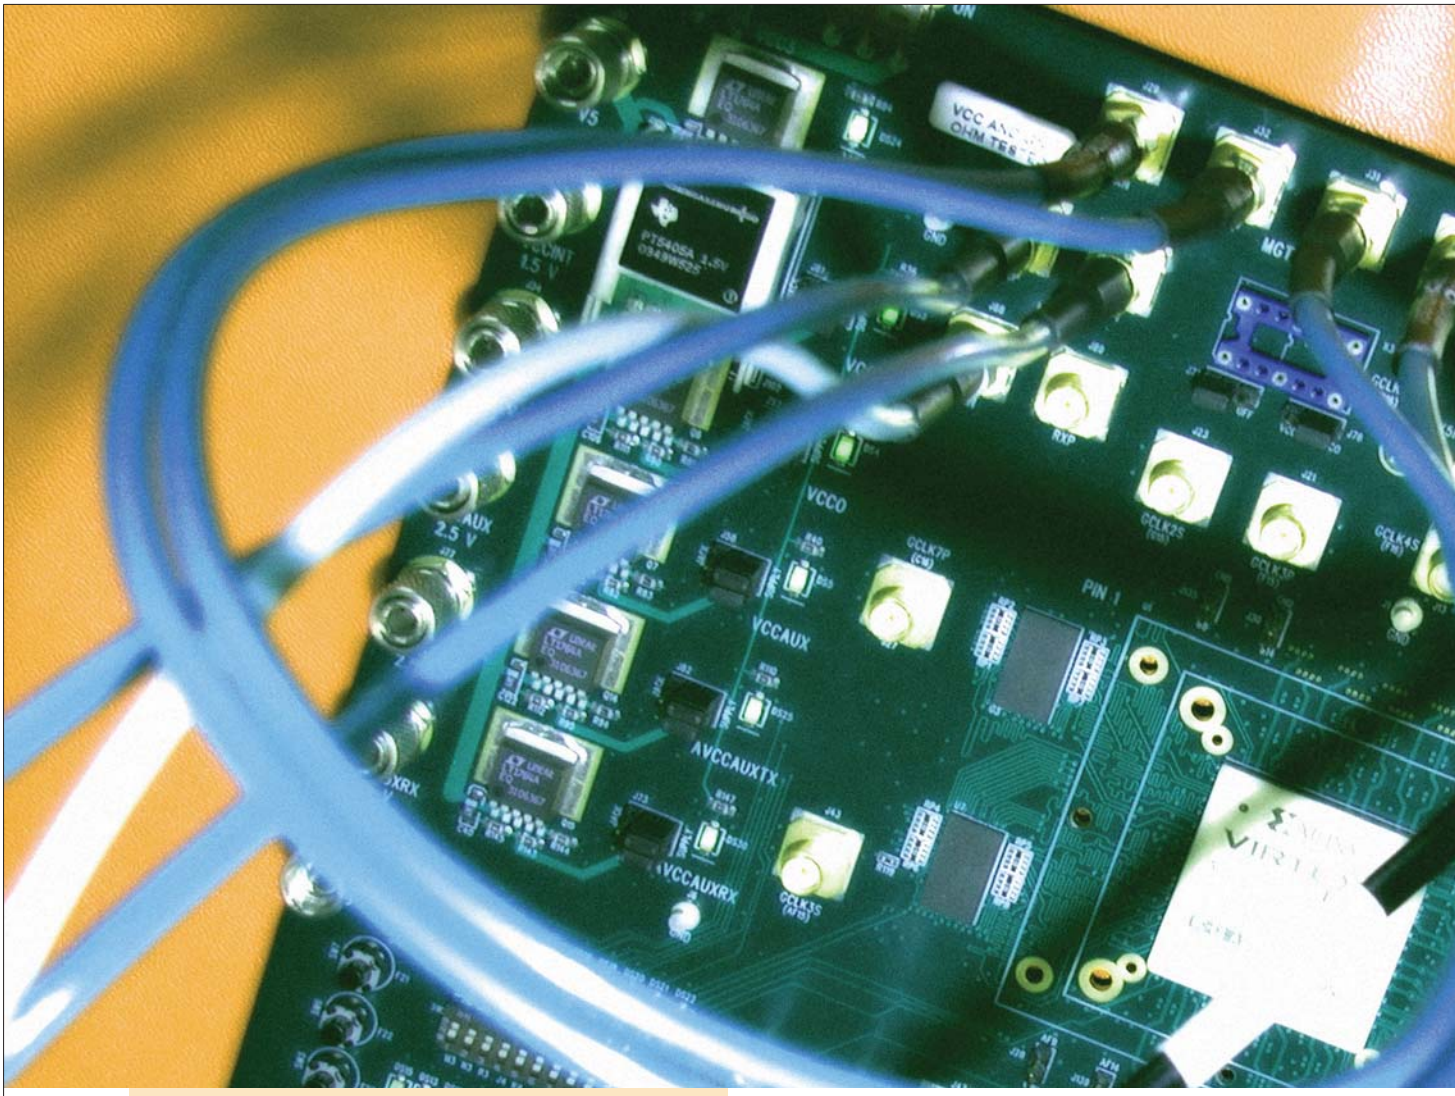

Serial I/O is putting greater constraints on board design, and bringing the disciplines of board engineering and layout closer together

by Chris Halford

Designers working with the latest field programmable gate array (FPGA) devices benefit from a large number of of high-speed I/Os at their disposal. Today's high-end FPGAs incorporate serial I/Os capable of more than 10Gbit/s, and the demand for new-generation products that require even higher increasing data bandwidths is making it difficult for the PCB designer to hide from the electrical effects.

Each year, close to four hundred printed circuit board (PCB) designs come through the company at which I work. We service a wide cross-section of clients from Formula 1 teams to consumer electronics giants. But, common to all the industrial sectors that we serve is a constant increase in the number and data-rate of high-speed I/Os used in designs.

Historically, PCB design has been carried out by draughtsmen working in the mechanical domain. The aim is to ensure a robust result that suffers minimal manufacturing issues. Electrical concerns have been limited to the current-carrying ability of traces and whether or not the layout connectivity reflects the netlist. The PCB has been treated simply as a substrate for connecting the components together. Today, however, the PCB is one of the most critical components in its own right.

In the past a casual approach to 'wiring' devices together has been possible due to the relationship between

Serial I/O layout shifts signal integrity design

signal transition times and the physical size of PCBs. With a 'slow' transition, the integrated-circuit (IC) bond wires, device packages and PCB routing can be considered to form a lumped system, throughout which a uniform voltage exists at all points. However, this situation is becoming rare as sub-nanosecond transition times become the norm.

Take a 400ps transition as an example. By the time the output reaches full swing, the start of the signal pulse will only have propagated approximately 6cm along a stripline trace. This means that virtually any interconnect being driven by this type of I/O must be treated as a distributed system: a transmission line, along which our signal propagates. In this type of system, traces must be of controlled impedance, properly terminated and designed with high-speed signalling in mind to avoid unwanted signal integrity (SI) artefacts.

Incredibly, the PCB technology we are using today has changed little over the past decades. We are still transmitting our signals across copper conductors embedded in epoxy fibreglass laminates (FR4). PCBs are still fabricated by etching away copper from both sides of cured FR4, clad with copper. These cores are then pressed together, sandwiched between layers of glass cloth, pre-impregnated with epoxy resin, commonly called pre-pregs. How well suited these raw ingredients are to high-speed I/O is indicated by two key parameters, the loss factor and the dielectric constant.

The material's loss factor, or tan-delta, is important in high-speed designs in which signals propagate across long links. All insulators heat up in the presence of electromagnetic waves, the same mechanism that heats up glass in a microwave oven will cause the heating of dielectric layers in a PCB. The tan-delta figure indicates the amount of energy that will be lost to the fibreglass through this heating mechanism.

Typical FR4 has a loss factor of approximately 0.02 at 1GHz. Improved resins in low-loss varieties can reduce this to 0.009 or so. Although these improved materials represent a significant additional cost, smart customers are arranging their cross-sections so as to only require the improved materials on one or two layers. Building a hybrid PCB in this way can be an extremely cost effective method of achieving a high-performance board that remains commercially competitive.

The other key factor, dielectric constant (E_r) of the fibreglass indicates the electric permittivity of the material, expressed as a factor of E_r in free space. This is a critical figure for calculating characteristic impedance. The figure needs to be as low as possible in controlled-impedance multilayer designs. Switching to a lower E_r material allows us to achieve higher impedances than →

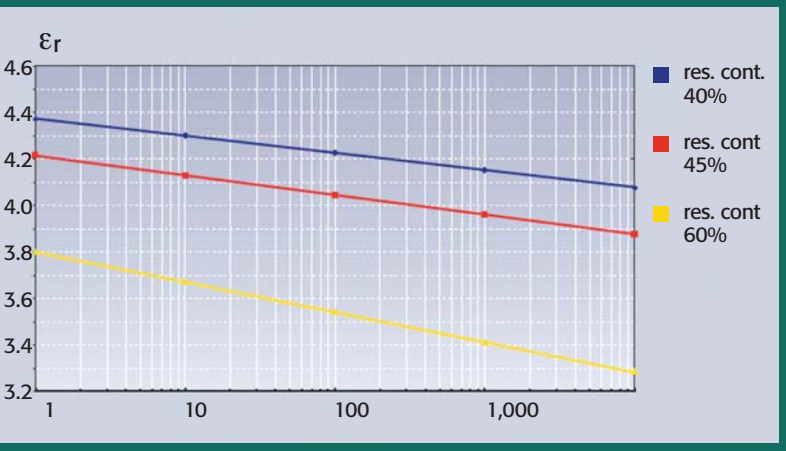

Fig 1 E_r of Isola 410 versus frequency for varying resin content

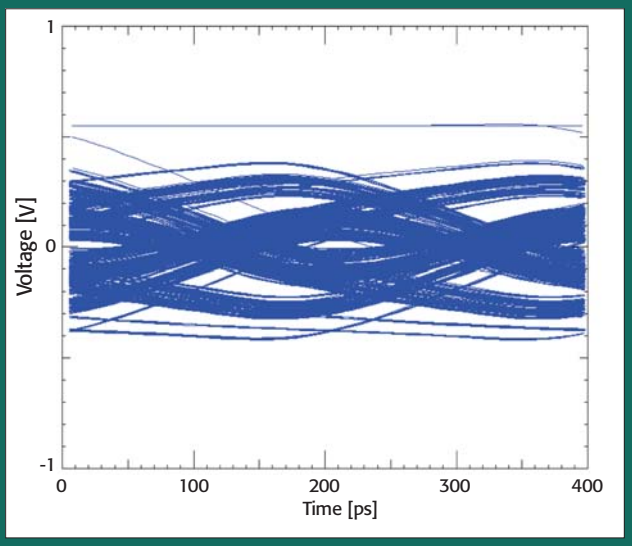

Fig 2 Receiver eye diagram for 24in, 5Gbit/s link with a poor PCB design

would otherwise be possible. This is particularly beneficial on internal PCB layers – striplines – in which the capacitive coupling of multiple plane layers makes it difficult to achieve sufficiently high impedances.

Similarly, a lower E_r will allow us to increase track width while maintaining the same impedance, a desirable change when trying to reduce signal attenuation caused by linear resistance. To predict impedance requires an extremely well-defined PCB cross section. Simply describing a PCB as a "12-layer construction in FR4" does not work anymore. The cross-section must include the exact materials being used to construct the laminate, and their electrical parameters at an appropriate operating frequency.

In a composite system such as FR4, the fibre and resin matrix have very different E_r values. Therefore, a varying ratio of glass to resin causes different material types to exhibit a wide variation in E_r, as can be seen in fig 1. Failing to properly specify materials during design can therefore lead to impedances being significantly out of specification.

The physical geometry of the trace structures themselves will also lead to local variation in E_r. For example, the gap between closely spaced traces will be almost completely devoid of fibres. But the area directly above the traces will contain comparatively little resin. Field solvers integrated into top-end layout tools take these variations into account and adjust the predicted impedance accordingly.

In addition to dielectric loss, another form of loss comes from the skin effect of the trace. As speed increases, a signal's current will tend to follow the path of least inductance. In the case of a PCB trace this means it will flow almost entirely in the outer skin of the conductor, as if it were a hollow box-section of copper. Furthermore, due to a proximity effect, the current will concentrate on the top and bottom surfaces that face the reference planes. The result of this is an increase in the linear resistance of the interconnect that can only be reduced by increasing the width of the traces. Although subtle, these effects appear at a surprisingly low frequency. By just 14MHz there will be no current flowing in the centre of a trace built on a 1oz copper that is 35µm thick. In long Gbit/s links this attenuation can be the difference between an acceptable error rate and the dreaded 'non-compliance'.

Although improved materials are extending the life of FR4, credit must also be given to the design of serial I/O drivers for supporting older materials. In order to propagate high-speed signals across legacy 'low data-rate' FR4 backplanes – a material once thought unsuitable for

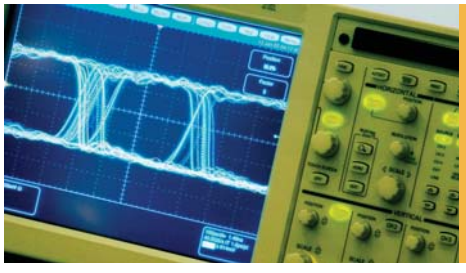

"Failing to properly specify materials during design can lead to impedances being significantly out of specification"

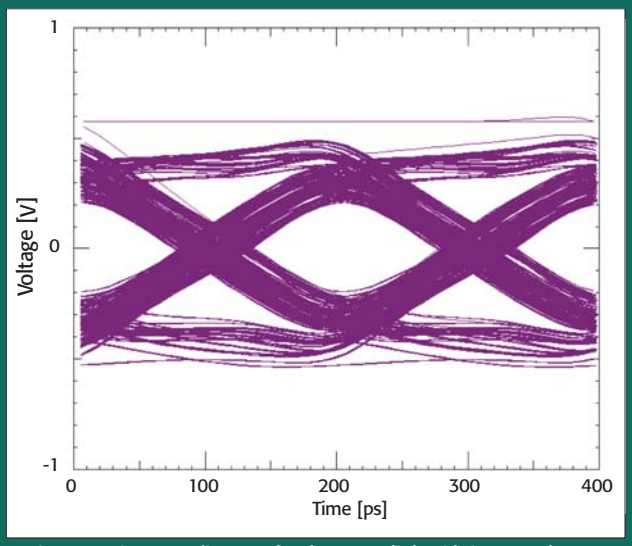

Fig 3 Receiver eye diagram for the same link with improved materials and design

Virtual assurance

Simulators are invaluable in today's PCB design flow. Regardless of the number of 'best practices' followed, the only way to reliably ensure a printed circuit board meets performance expectations is to run simulations. Engines such as Cadence Design Systems's TLSim, which include the effects of high-frequency signal loss, are essential tools in the quest for right-first-time designs. If a design does fail on its first iteration then the fixes required for high-speed I/O tend to be modifications that cannot be tested out very easily in hardware. Previous I/O types could be nursed into life by tacking on passive components in the lab, but newer high-speed I/Os may fail due to parameters that cannot be adjusted without a board re-spin, such as track width or dielectric material. This further emphasises the need to simulate, because having real prototype hardware in your hands may not get you much further forward.

Figs 2 and 3 show the difference that can be made by improving PCB layout in a high-speed design. The collapsed eye in the first diagram is opened up by improved interconnect design alone. A design flow that makes use of signal analysis allows such problems to be identified very early on, ensuring that the most suitable material is chosen, that the added cost of special laminates are only carried where necessary and that the design is right first time.

multi-gigahertz use – serial I/O drivers generally use pre-emphasis and equalisation. These techniques allow the frequency response of the PCB interconnect to be flattened, cancelling out the detrimental effect of the PCB. Pre-emphasis gives the first bit of a sequence of the same symbol a greater signal level, causing transitions to have a faster edge rate, effectively boosting the high frequency content of the signal. Equalisation in the receiver provides a similar high frequency bias by attenuating the low frequencies and boosting the signal, thus equalising the filtering effect of the interconnect.

Electrically, designing interconnect is damage limitation at best. Each PCB trace is a passive low-pass filter, the parasitics from every additional trace, via, connector and package will further reduce the frequency response, limiting the bandwidth that the interconnect can support.

The popularity, availability and data-rate of high-speed I/O continues to increase. However, offering these technologies up to the ball of a ball grid-array package alone does not ensure success. It takes a well-engineered PCB layout to distribute these signals across large, noisy PCB and backplane systems.

Links must be designed as a complete electrical system from die to die, with PCB layout a major part of the engineering effort, ruling out the traditional layout-bureau design-flow of 'netlist in, plot-files out'. With poor PCB design, the blistering performance of new generation I/Os is lost. Fortunately, the layout community is addressing this problem by integrating signal-analysis tools into its design-flows. Rather than draughtsmen, new recruits to layout tend to be design engineers and graduates with an interest in SI and system design, resulting in boards that cater for a client's electrical needs. It's not the game it was.

A successful design flow for high-speed I/O demands the use of constraints, field solvers and signal-integrity simulators to accurately control and predict the behaviour of signals prior to committing to manufacture. Using an external bureau for PCB layout raises design-flow concerns for many design authorities. Signal-integrity issues common on high-speed designs force engineers to take an increasing level of interest in the detail of a PCB layout, leading to a very inefficient design-flow in which the engineer watches over the shoulder of the layout designer.

As signal-integrity tools are very closely coupled to layout tools, it is more efficient for the same team to be responsible for both. Therefore, the choice is simple, bring PCB layout and analysis in-house or have your layout bureau handle the signal-integrity work. Bringing layout in-house requires significant investment in both maintained tools and trained users; a top-end PCB layout and signal integrity package will not leave much change from £150,000, a significant investment, especially for a tool that may only be used sporadically. For this reason the layout bureaux have started to extend their skill sets to become a closer part of the design team by deeply integrating signal analysis into the layout process. ■

Chris Halford is signal integrity engineer at Advanced Layout Solutions