A Signal Integrity Guide to HSD PCB Design

cadence®

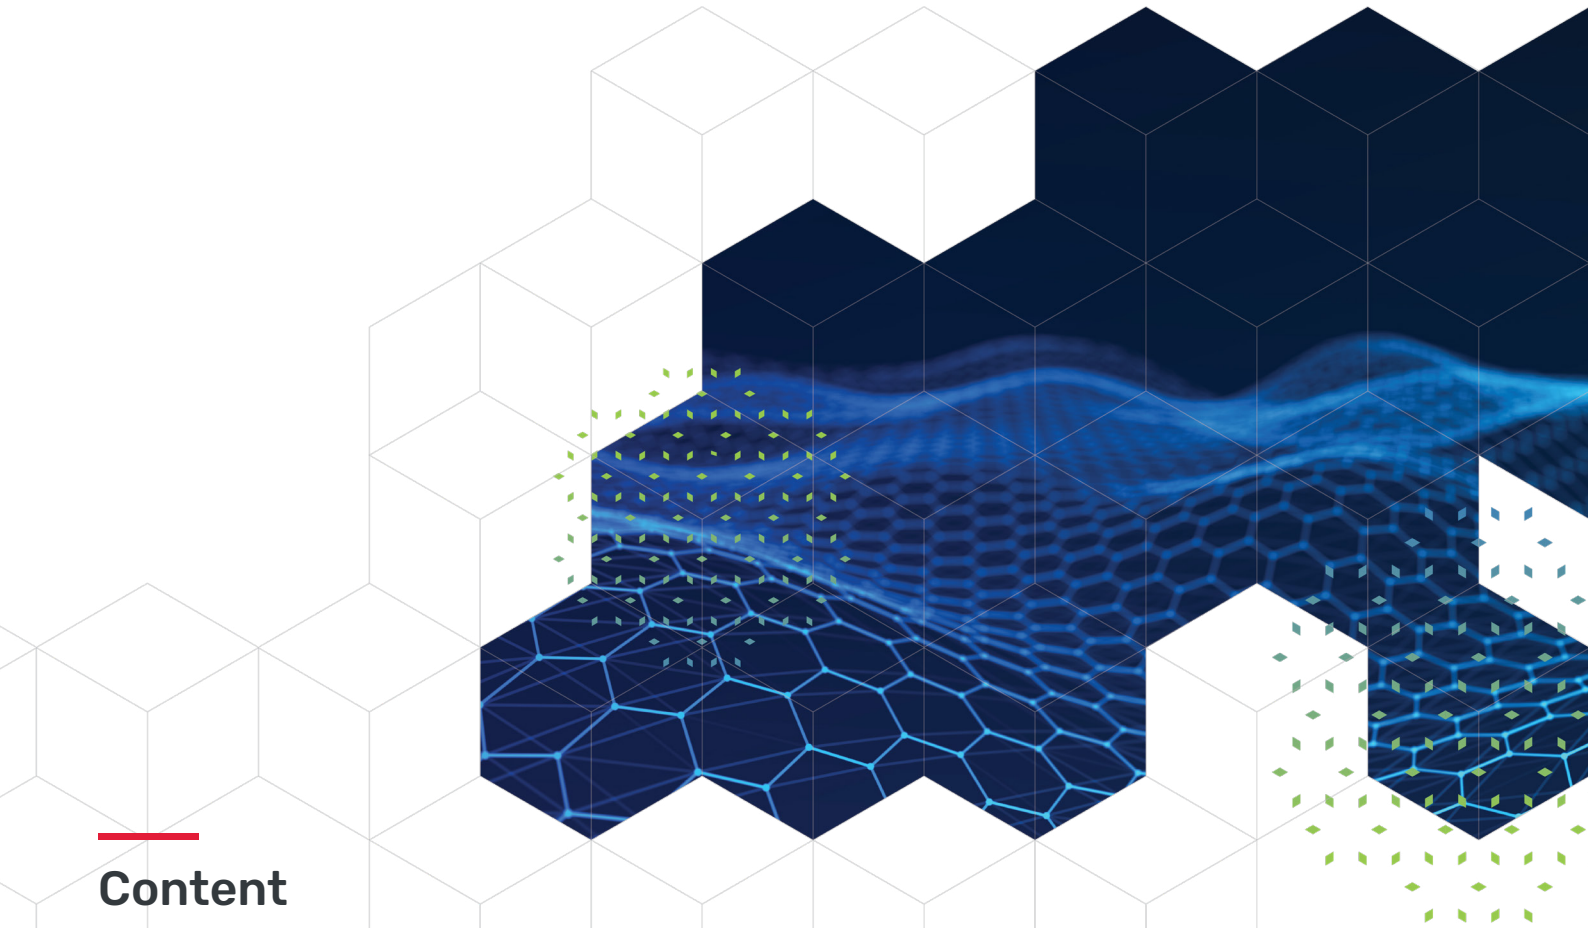

Content

HSD: Not quite RF design, but you still have to worry about harmonics 4

Characterizing the quality of your design 5

Differential signaling and transmission lines 6

Impedance mismatches 8

Meeting industry compliance specifications 9

Crosstalk 10

The Problem 10

The Solution 11

Clock and signal skew 12

The Problem 12

The Solution 13

Ground bounce or SSN 13

The Problem 13

The Solution 14

Conclusion 14

References 15

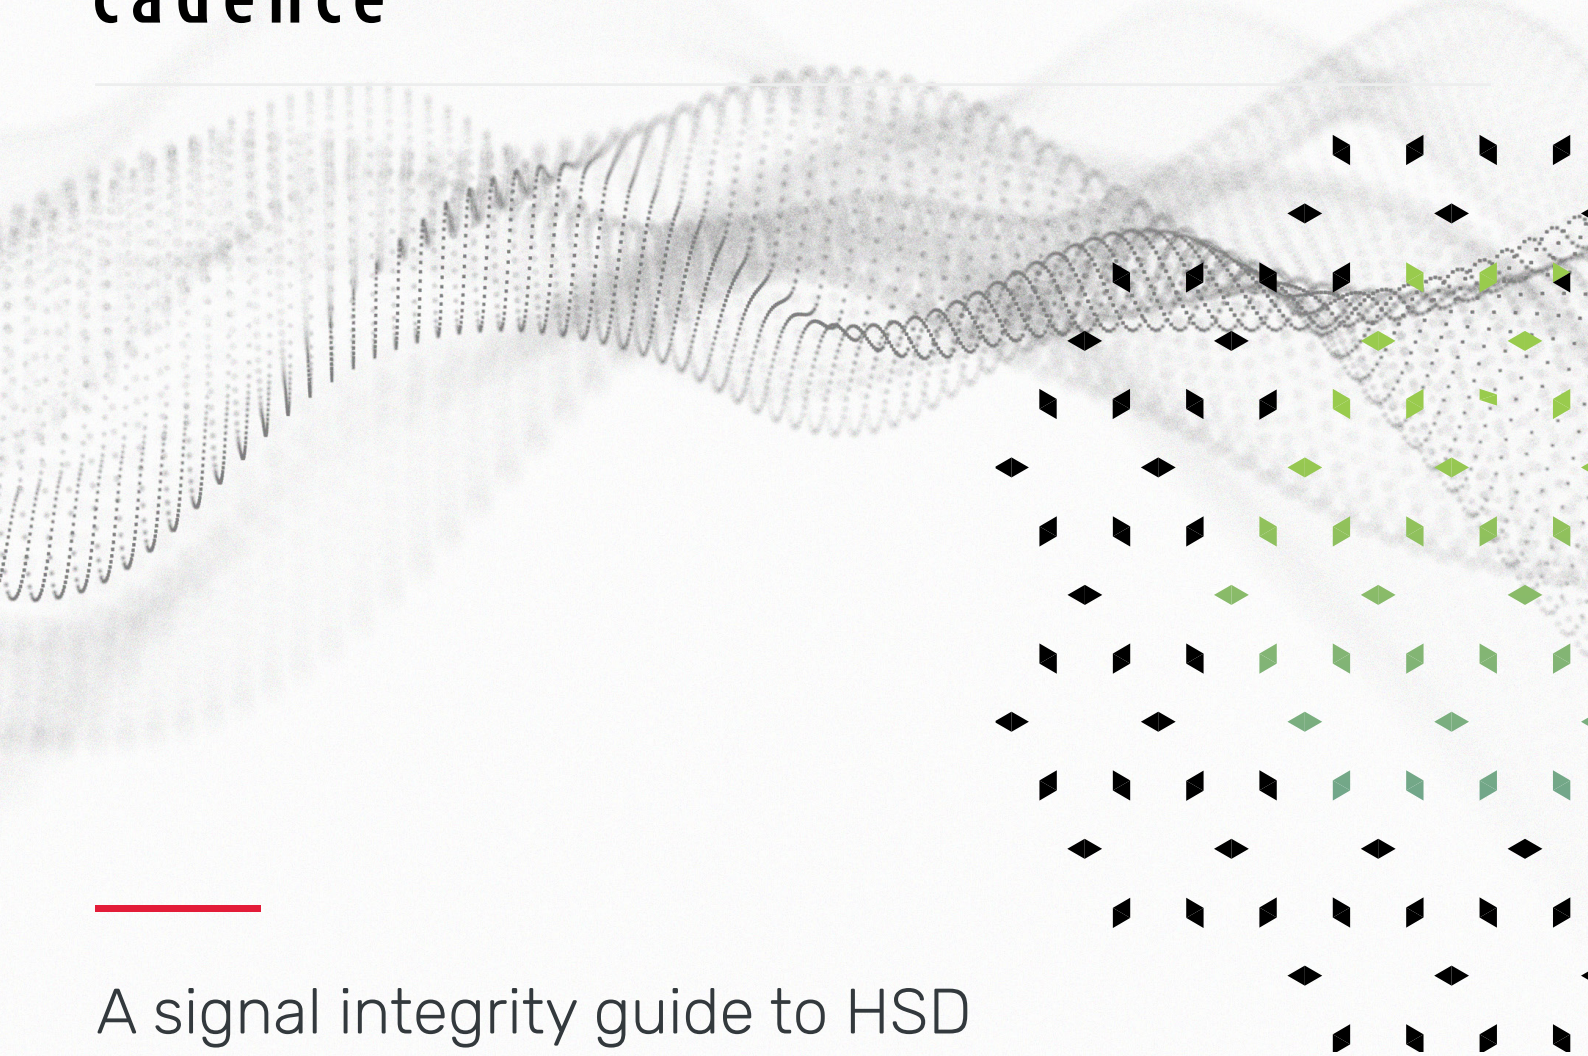

A signal integrity guide to HSD PCB design

The majority of PCBs are standard multi-layer rigid FR4 laminates, designed to reliably conduct signals along their copper traces with minimal signal integrity issues. Designers are constantly advancing the limits of data throughput on high-speed interfaces such as PCIe, DDR, HDMI, Gigabit Ethernet, and USB by employing high-performance serial modulation techniques like PAM4 and NRZ. As this occurs, the corresponding PCB systems must adopt multi-board high-density interconnect (HDI) architectures with integrated backplanes, which ensure structural stability and facilitate communication with multiple daughter boards. These high-speed digital (HSD) systems operate with a tight timing budget, meaning any signal distortions can compromise performance. Thus, at high speeds and corresponding high frequencies, PCB layout becomes increasingly critical.

Signal distortions are caused by signal integrity (SI) complications in your PCB, where SI refers to the quality of an electrical signal as it travels through the board. Among SI issues are: any signal reflections from impedance mismatches, crosstalk from parasitic coupling of nearby traces, or improper design of signal return paths that cause distortions. SI designs also involve a power-aware component—specifically, power integrity (PI) analysis, including challenges such as ground bounce. This ebook explores common SI challenges in HSD designs and offers guidance on how to avoid them.

HSD: Not quite RF design, but you still have to worry about harmonics

To achieve high signaling and symbol rates on the order of gigabits per second (Gbps), many HSD interfaces have signal spectrums that extend into microwave frequencies, with frequency components (harmonics) that can reach millimeter-wave ranges. For example, USB 3.0 can generate noise within the commonly used 2.4 GHz ISM band, potentially interfering with other wireless devices such as Bluetooth and Wi-Fi.

HSD interfaces use ultrafast clocks with high-frequency stability and additional parallel data lanes to transmit greater volumes of data per second, thereby achieving higher throughput. The below equation provides an approximate estimate of the spectral bandwidth required to capture the highest frequency components of an HSD signal:

$$BW = \frac{0.35}{\tau_R}$$

Where:

- BW is bandwidth
- τ_r is 10-90 rise time, or the time it takes for the signal to go from 10% to 90%.

This is the exact solution for the bandwidth of a single-pole low-pass filter (LPF) at its –3 dB point, or frequency point at which the amplitude of the transfer function rolls off by 3 dB [1]. Given that HSD systems are not LPF, this is a rule of thumb rather than a precise solution. Modern HSD interfaces have extremely tight timing constraints on the order of tens of picoseconds; for a PCIe 1.1 signal with a τ_r of 50 ps, the BW is ~7 GHz, for PCIe 2, this goes to 30 ps, requiring ~11.6 GHz.

Inadequate attention to the circuit's spectral behavior during PCB layout can create challenges, which is why early-stage simulation is essential for identifying SI problems before they result in expensive consequences. Conducting SI analysis as early as possible saves engineers from the time-intensive task of troubleshooting issues before producing a physical PCB or, worse, after the PCB is fabricated. Many engineering teams have faced costly design respins due to overlooked SI challenges. However, these expenses could be substantially reduced by integrating simulation during the layout phase, followed by comprehensive system-level signoff from both SI and PI engineers. This proactive, simulation-driven approach throughout the design process helps detect and resolve potential problems early, ultimately conserving valuable time and resources.

Characterizing the quality of your design

Timing margins are tight, and they can be impacted by many factors. The bit error rate (BER) measures bit synchronization, or the number of error bits that occur out of the total number of transmitted bits in a binary encoded stream, i.e., the signal. Every system has an intrinsic BER that is unavoidable— the goal is to ensure that the BER is not unnecessarily higher due to design issues.

Clock jitter worsens BER. Clock jitter, which refers to timing variations in signal edges from their ideal values, represents the instability of a clock signal. It can be classified into several types:

- **Random:** intrinsic to the system
- **Deterministic:** repetitive and narrowband
- **Periodic:** associated with the periodic frequencies of noise sources
- **Data-Dependent:** intermittent and dynamically changes its duty cycle, producing irregular clock edges

Layout issues that cause high signal attenuation or interference—such as long trace lengths, lossy traces, and suboptimal dielectric materials—exacerbate jitter and introduce additional problems by allowing high-frequency harmonics to affect the clock signal edges. (e.g., crosstalk).

While crosstalk generally creates a periodic jitter, intersymbol interference (ISI) can create a data-dependent jitter. ISI occurs when one signal in the bit stream interferes with later signals when read out by the receiver, causing errors in data transfer. ISI can be traced to poor layout design involving issues such as trace impedance mismatches, signal dispersion in the substrate material, and parasitic capacitances from components and other shunt elements.

Figure 1: *Oscilloscope eye diagram plot shows extreme jitter that includes ISI.*

Differential signaling and transmission lines

As mentioned previously, crosstalk causes a range of issues, including increased jitter, decreased BER of the signal, and degradation of the overall HSD signal performance. Before diving into crosstalk, it's important to understand the significance of trace routing and reference planes, and how they relate to the transmission line's impedance.

HSD transmits and receives data using multiple parallel lanes of differential signals, or two separate conductor paths carrying equal-magnitude currents of opposite polarity **(Figure 2).** The differential receiver then calculates the difference between the two, resulting in an interface with twice the voltage swing and improved noise immunity compared to a single-ended transmission. PAM4 specifications, for example, call for four to eight lanes of balanced differential signals to better reject common-mode interference—with each lane operating at 50 to 100 Gbps.

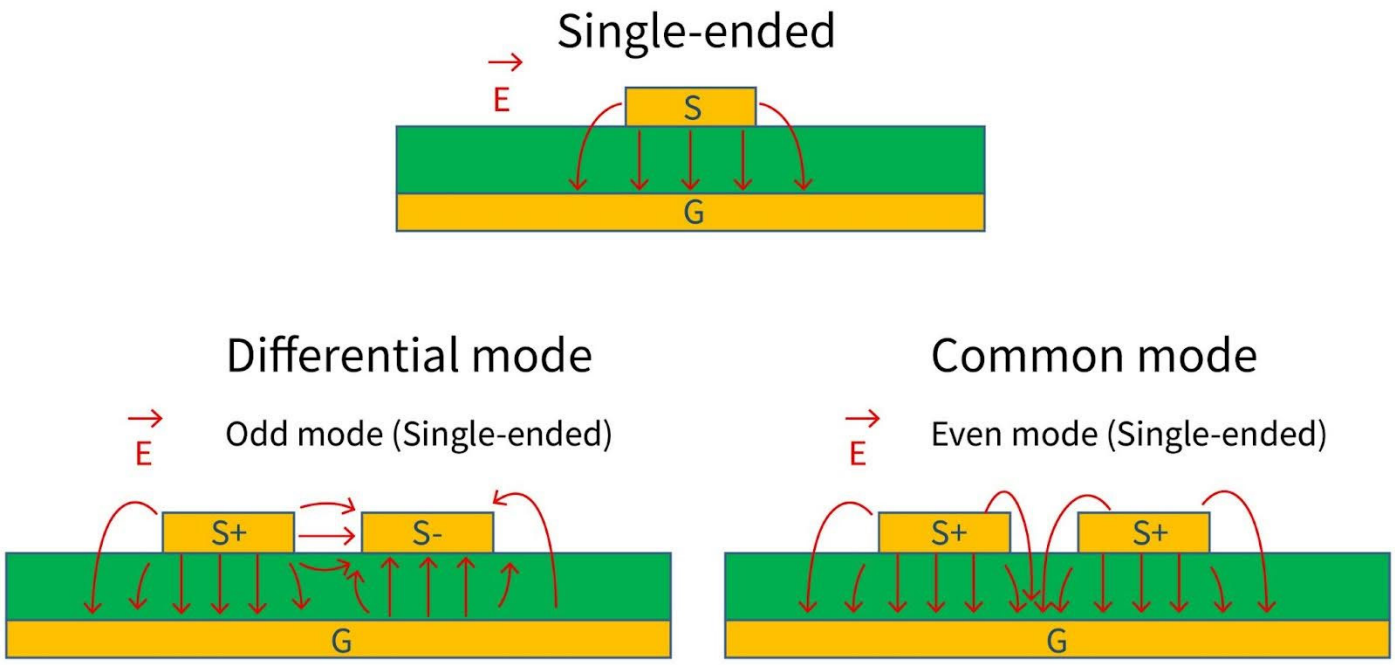

Figure 2: *Single-ended versus differential versus common-mode signaling. Source: Cadence*

Maintaining proper differential signaling and SI in HSD systems becomes more difficult when considering spectral performance. The path must maintain a consistent differential impedance along its route, or it will suffer from reflections and loss. Therefore, analyzing different transmission lines and their characteristic impedance equations is valuable **(Figure 3)**.

The signal's rise time and the length of the trace are related to whether it will behave like a transmission line:

$$L_{trace} \sim \tau_r * d_{prop}$$

Where:

- L_{trace} is trace length
- τ_r is the signal's rise time in picoseconds
- d_{prop} is the velocity in the medium, or propagation delay (ps/in).

While wavelength-based rules of thumb exist, their effectiveness is uncertain because digital signals encompass a wide spectrum of frequencies. Instead of calculating whether your signal will behave as a transmission line, it is safer to design a transmission line to the impedance requirements for HSD signals.

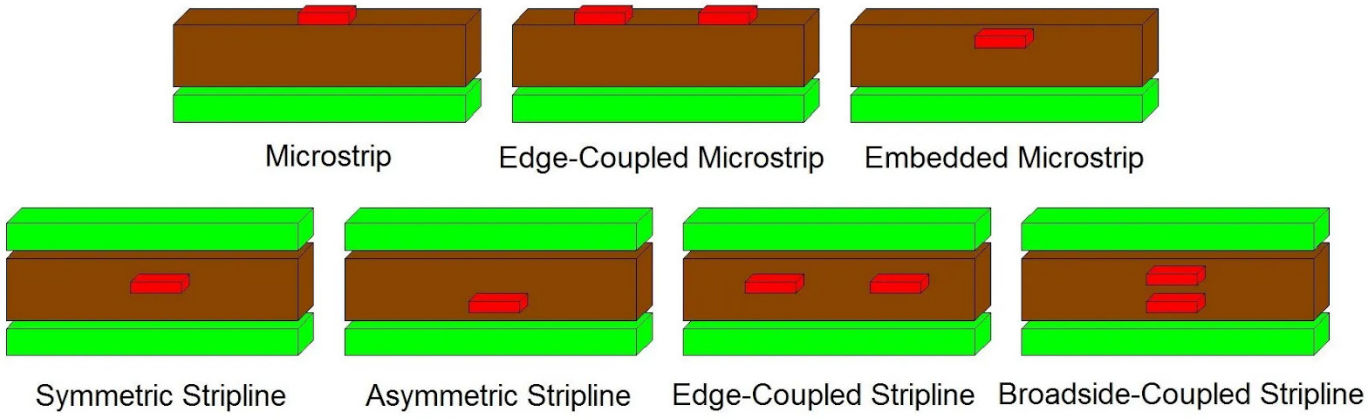

Figure 3: *The different transmission lines that can be routed in a PCB for either single-ended or differential signaling. Source: Cadence*

Differential signals are either of the microstrip or stripline variant, with two symmetric traces running in parallel and spaced equally. Any inconsistencies and asymmetries can potentially influence the otherwise convenient symmetries in the magnetic fields surrounding these traces, causing electromagnetic (EM) and SI issues. This is because the way that a differential pair is laid out is electromagnetically convenient, or fields that are equal and opposite. If you're familiar with the right-hand rule, this principle helps you visualize how the magnetic fields propagate in equal and opposite directions when the electric fields in the transmission lines exhibit opposite polarities or directions.

The differential impedance on a balanced differential pair of transmission lines is double the odd-mode impedance of each trace **(Figure 4)**:

$$Z_{odd} = \sqrt{\frac{L - L_m}{C + 2C_m}}$$

Where:

- Z_{odd} is the odd-mode impedance
- L is the self-loop inductance
- L_m is the mutual-loop inductance
- C is the self-capacitance
- C_m is the mutual capacitance

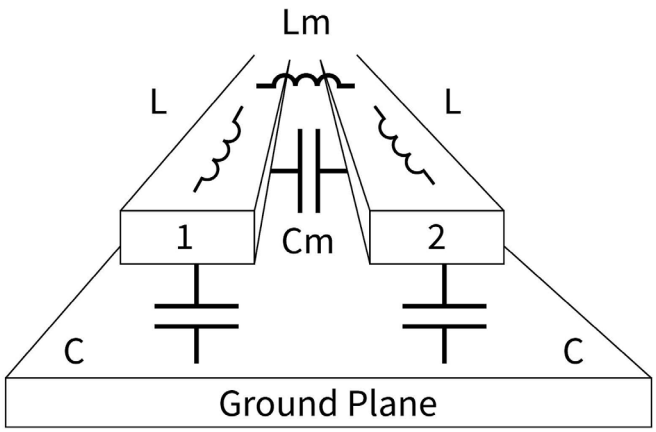

Figure 4: *Microstrip differential pair traces showing self-loop inductance, self-capacitance, mutual inductance, and mutual capacitance.*

Both the inductance (L and L_m) and capacitance (C and C_m) depend on the physical characteristics of the circuit. Inductance depends on factors such as trace width, trace thickness, layer thickness, and dielectric properties. Capacitance depends on the trace geometries concerning the return path, or ground plane. In other words, the characteristic impedance is determined exclusively by the physical geometry—such as the cross-sectional spacing between the line and its return path—and the dielectric properties, including its dielectric constant (relative permittivity). This is an important observation as it means the layout designer must adhere as closely to these physical geometries as possible to avoid EM and SI issues.

This also explains why modified epoxies and hydrocarbon-ceramic laminates are often used in HSD designs. Although they aren't on par with the PTFE substrate often used for high-frequency designs, they offer relatively low dielectric constant (Dk) and dissipation factor (Df) values. For modified epoxies, the high glass transition temperature (Tg) ensures that the structure remains stable and rigid even at elevated temperatures, thereby mitigating the unwanted timing delays that occur with the natural expansion of dielectric materials when heated.

Impedance mismatches

Having a basic understanding of the transmission line theory can help explain why SI issues occur in HSD designs. Any changes in the cross-sectional area of the trace would yield an adjustment in the characteristic impedance of the line, potentially causing reflections and distortion. Trace geometry should be kept as uniform as possible, with

little changes in trace width or sharp bends or turns that result in variations in impedance. Reference planes must be kept as uniform as possible because a lack of a consistent return path can cause impedance changes.

Vias can now be seen as potential "stubs" that can cause significant reflections due to impedance changes. Blind or buried vias, which connect an outer layer of the board to one of the inner layers without traversing all PCB layers, reduce stub lengths compared to through-hole vias. Microvias may be a more ideal alternative because of their reduced aspect ratio. The coupling between neighboring transmission lines can cause the actual impedance seen by a signal on a trace to differ from its characteristic value, leading to crosstalk.

High-speed traces on the PCB are often laid out with great care, while the remaining traces can be sent out for manufacturing without inspection. Ideally, a post-layout board impedance analysis should be performed to avoid any costly oversight.

Meeting industry compliance specifications

Both SI and PI are critical for meeting industry standards in high-speed digital systems. This applies at both the die level (e.g., DDR, UCIe) and the PCB level (e.g., PCIe, USB, 100G/400G Ethernet, RapidIO, InfiniBand), where reliable data transmission and power delivery are essential at the system level. As described earlier, each standard has its own data rate per lane, lane configuration, and total bandwidth **(Table 1)**. Each standard is optimized for a specific design focus, such as high bandwidth with minimal crosstalk, high bandwidth with low latency, or low latency and high reliability.

Table 1: *Specifications of various high-speed signal standards.*

Standard	Data Rate Per Lane	Lane Configurations	Total Bandwidth
PCIe Gen6	64 GT/s	Bidirectional x8, x16	256 - 1024 GT/s
Ethernet (100GBASE-KR4/ 400GBASE-KR8)	25/50 Gbps per lane	Unidirectional 4/8 lanes	100/400 Gbps
InfiniBand HDR/NDR	50/100 Gbps per lane	Bidirectional x4, x8, x16	200 - 800 Gbps
RapidIO Gen 3/4	10/25 Gbps per lane	x4, x8, x16	40 - 160 Gbps (Gen 3) 100 - 400 Gbps (Gen 4)
USB4 2.0	-	10 data wires: 2 full-duplex (2 data wire pairs), +1 half-duplex (1 data wire pair)	80 Gbps
DDR4	3200 MT/s (transfer rate over bus)	16 banks	25600 MB/s (bus bandwidth)
UCIe 1.0	4, 8, 12, 16, 24, 32 GT/s per pin	Unidirectional 16 to 64 lanes	-

All of these factors directly correspond to layout characteristics, such as differential impedance, lane width, pair spacing, and skew matching. For example, InfiniBand, RapidIO, and Ethernet call for a differential impedance of 100 ohms; USB is 90 ohms; and PCIe is 85 ohms. Back drilling, or blind/buried vias, might be mandatory to reduce via stub lengths. The differential impedance, in turn, influences the lane width and pair spacing, which then affects the high-speed routing strategy. Limitations on intra-pair skew also impacts routing, as they impose constraints on matched trace lengths for proper delay tuning.

Robust simulation tools are essential for ensuring compliance with industry standards across PCB design, SI and PI analysis, and timing verification. Effective PCB design software should provide a fully configurable set of design rules and constraints to specify the length and flight times of HSD signals. For system-level global analysis, it is crucial to use software capable of generating realistic simulation results that guide designers in meeting stringent industry requirements.

Crosstalk

The Problem

Crosstalk, or electromagnetic coupling, can occur with closely spaced traces that run parallel on the same board layer and even between side-by-side traces. The "aggressor" trace induces noise on the "victim" trace. It is caused by a parasitic mutual inductance (L_m) or a parasitic mutual capacitance (C_m) **(Figure 5)**.

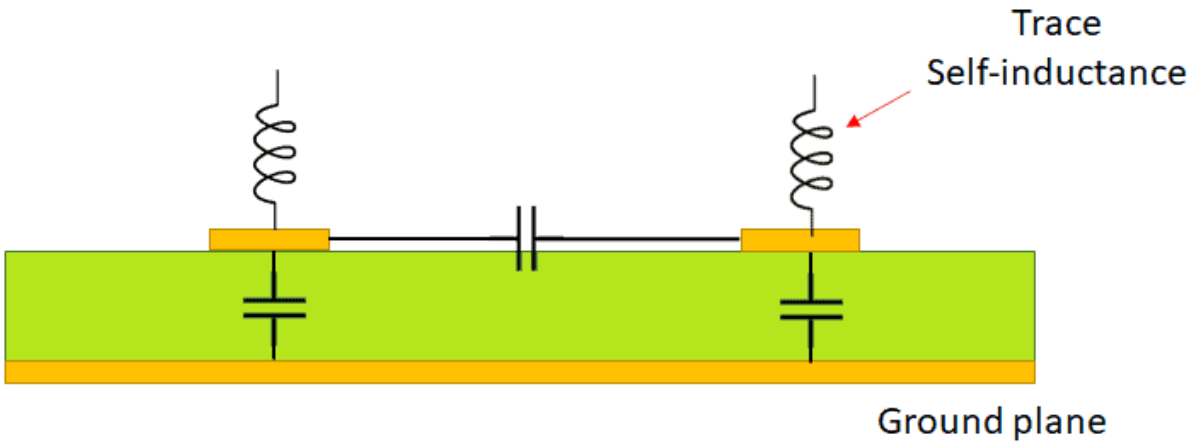

Figure 5: *The parasitics between two traces inevitably produce a mutual capacitance and inductance. If this is unwanted mutual coupling, i.e., parasitic, it couples noise into otherwise well-behaved traces as crosstalk.*

Capacitive crosstalk occurs due to a changing potential difference between two traces, while inductive crosstalk occurs when the aggressor signal changes levels (higher speed signals generate more inductive crosstalk). There are two major types of crosstalk: near-end crosstalk (NEXT) and far-end crosstalk (FEXT). NEXT occurs on the driver side of the victim trace, while FEXT occurs on the receiver side of the victim trace **(Figure 6)**.

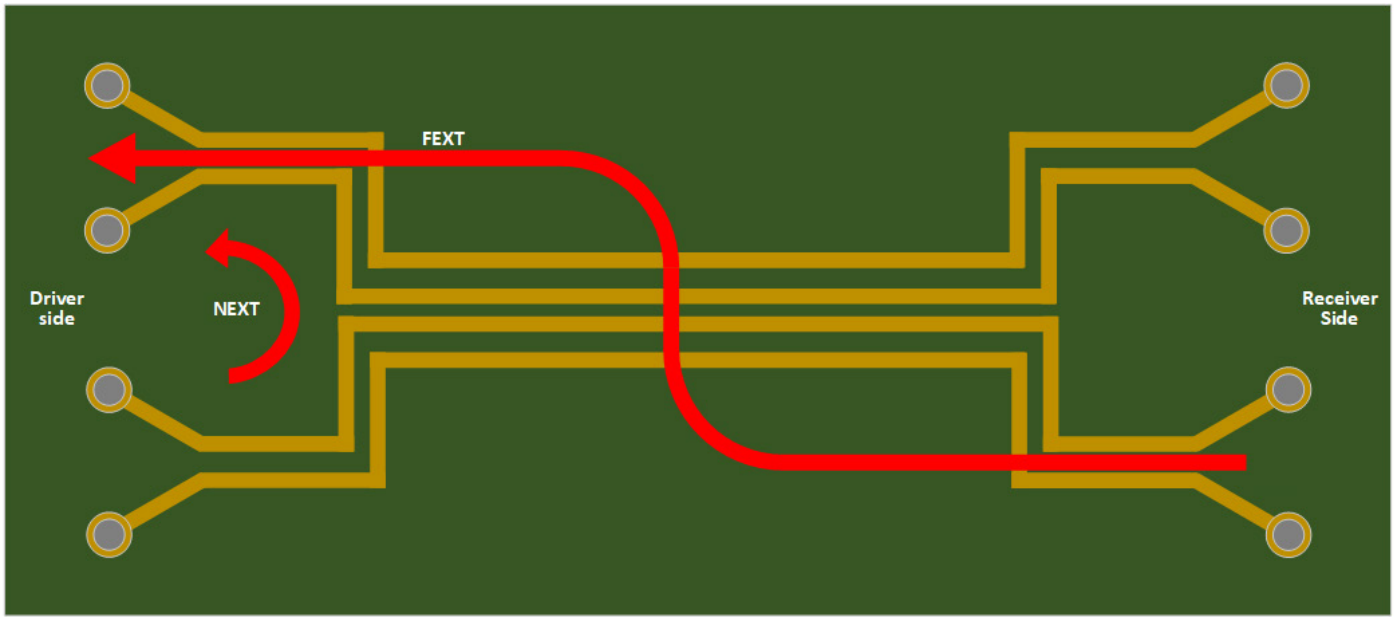

Figure 6: *NEXT and FEXT for aggressor traces on the victim traces for two sets of differential signals. Source: Cadence*

The Solution

Some degree of coupling is inevitable, but it must be minimized to avoid negatively impacting design performance (i.e., high BER, jitter, etc.). The following strategies can help mitigate these effects:

1. Increase the trace width and reduce the distance to the ground plane for interfering linesMaximize spacing between high-speed signal traces
2. Incorporate shielding and maintain solid ground references

Inductance can be a complex concept, but it is generally reduced by increasing trace width and decreasing the distance to the reference plane—both of which help lower self-loop and overall inductance. If this SI issue is identified after fabrication, resolving it may require a stack-up redesign to position the ground plane closer to the interfering traces. It's also important to remember that the term "plane" implies continuity; splitting the ground plane, especially under HSD signal routing, can severely degrade signal integrity and should be avoided.

When routing long parallel lines, or HSD signals, it is crucial to space them out to reduce both NEXT and FEXT. Cadence follows the general principle of spacing out traces three times their line width. Shielding structures such as additional ground pour or via fencing can be used to actively minimize interference between victim and aggressor routing. Human error can sometimes be unavoidable and mitigating crosstalk can become more challenging with increasing circuit complexity. Simulation tools can allow a designer to simulate the cross-section of coupled traces with field solvers to quantify and mitigate NEXT and FEXT.

Clock and signal skew

The Problem

In traditional parallel synchronous buses, one driver is connected to several receivers, where each branch carries data in parallel that needs to reach the load at the same time. Without precise length matching across these signal groups, a timing skew arises, which can result in increased bit error rates (BER) and latching errors—ultimately compromising data integrity.

Lane-to-lane skew in high-speed signaling is less impactful due to the embedded clock. However, between the differential pairs, which carry signals of opposite but equal polarities, there must be a level of skew matching between the D+ and D- legs of these lanes **(Figure 7)**. Given the high speeds and tight timing margins, timing differences in these differential lanes on the order of picoseconds can cause some of the differential-mode signals to become common-mode signals. This shift increases common-mode noise and EMI while causing other signal integrity issues.

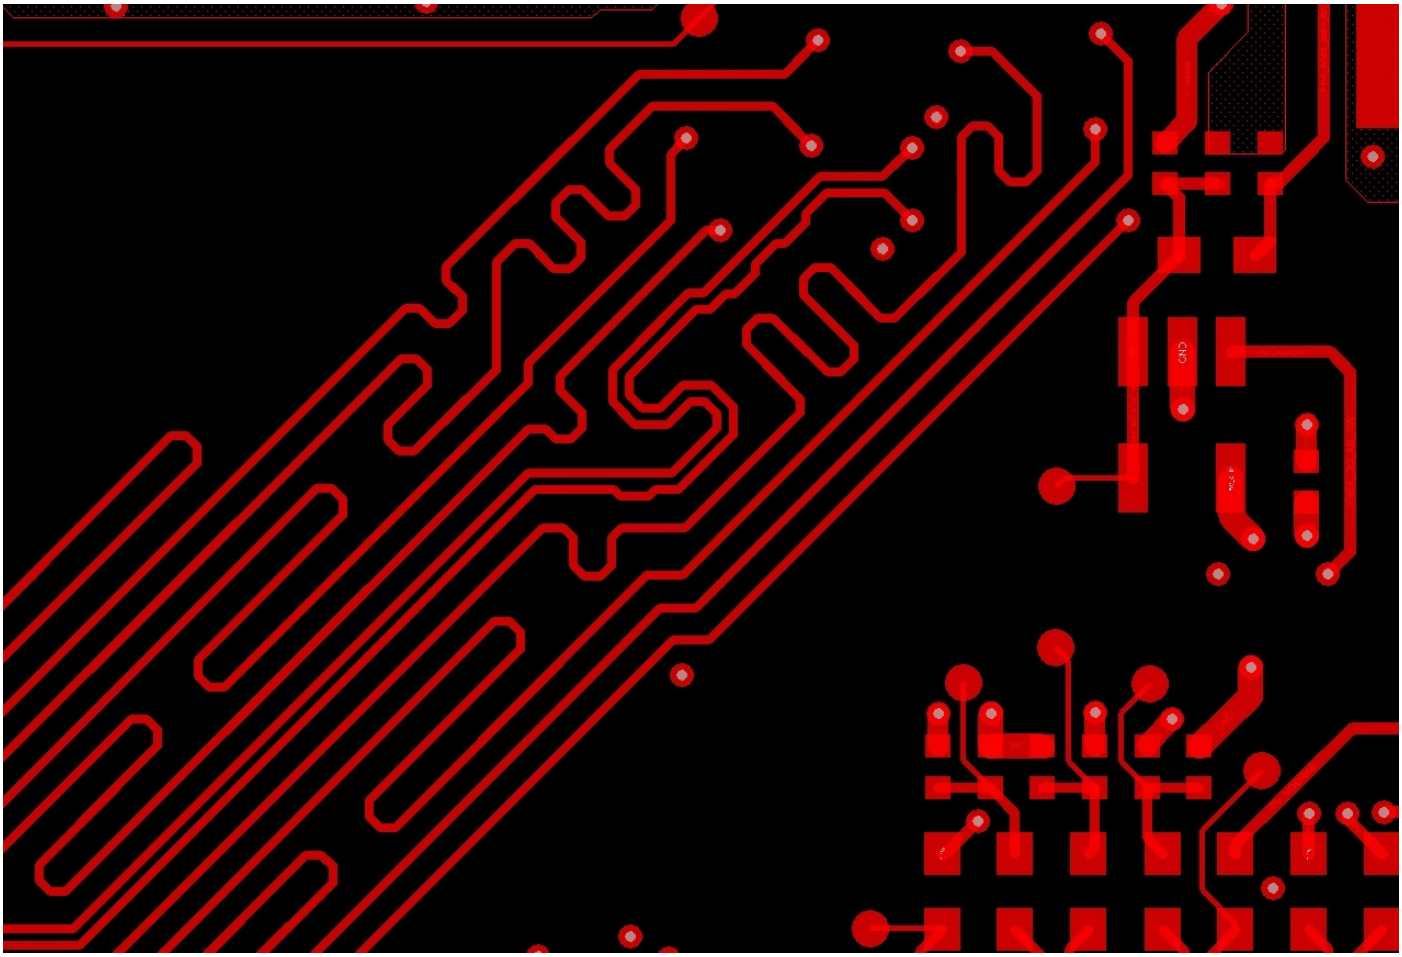

Figure 7: *Trace length matching for each differential signal within a lane is critical for SI. Source: Cadence*

The intra-pair timing, or phase skew, is based on each trace's "electrical" length. This is dependent on both the PCB laminate's relative permittivity (ε_r), or dielectric constant (Dk), and the mechanical length of the trace:

$$L_{ELEC} = L_{MECH}\sqrt{\varepsilon} \qquad (3)$$

The Solution

As stated earlier, HSD PCBs often use either modified epoxy or hydrocarbon-based ceramics. The modified epoxies, in particular, are reinforced with a woven e-glass with a low Dk. While the lower Dk ensures that the signals are less lossy, minute differences in the Dk because of the glass weave can lead to timing skew, which must be considered and accounted for in the skew budget. For instance, the minimum intra-pair timing skew for USB SuperSpeed is 15 ps, and this allows for a trace length mismatch of less than 100 mils on the standard FR4 laminate.

Start by routing the signal to the receiver pin farthest from the driver within the differential pair, then route the second trace to match the same mechanical length. Ideally, these pairs should be routed with minimal discontinuities—avoiding sharp corners or bends that can alter trace impedance by changing the effective trace width. Such impedance mismatches cause reflections that affect signal delay. To minimize these risks, establish a consistent routing strategy early in the design process.

Ground bounce or SSN

The Problem

Some problems don't occur because of the layout; rather, they occur because of the parasitics that are intrinsic to components mounted to the PCB. The power delivery network (PDN) is composed of a large number of fast-switching I/O drivers, each with its own bond wires, package pins, and on-chip interconnects (i.e., parasitic inductance). These drive I/Os switch from logic 0 to logic 1, charging up draws from the capacitance of the PDN. However, the parasitic inductance of switching ICs impedes current flow to logic 1 during fast edge transitions, causing current transients and voltage fluctuations in the power and ground lines, which can result in ringing on the supply rail.

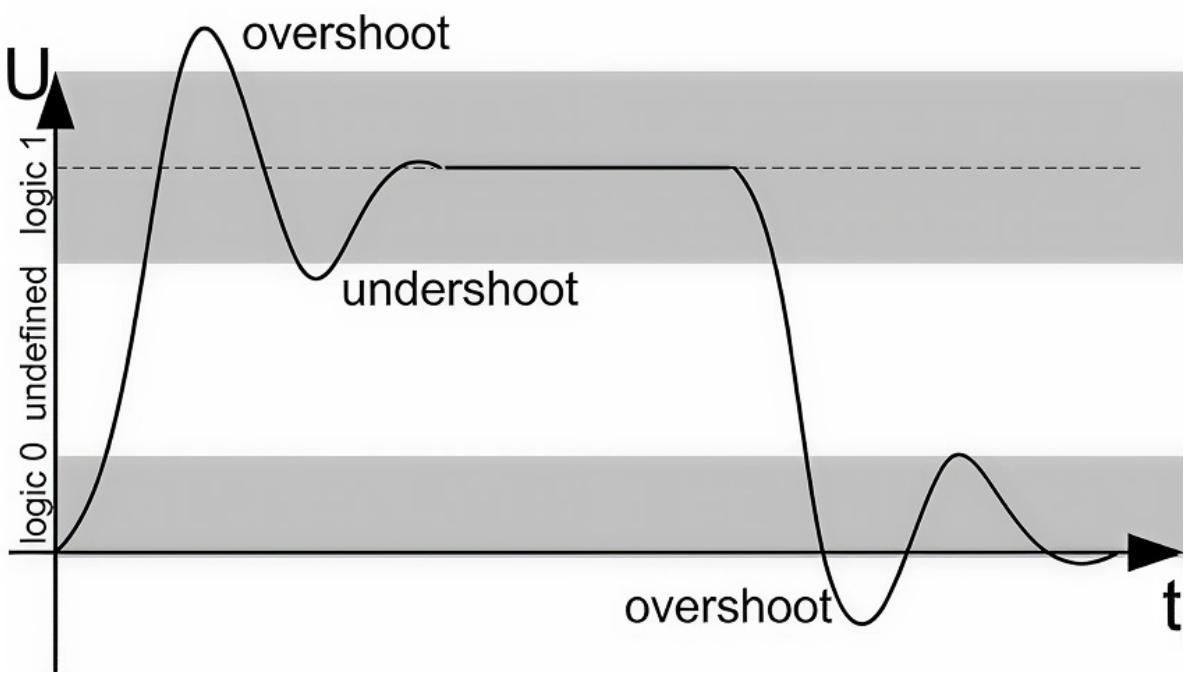

Figure 8: *Overshoot and undershoot are potential signal integrity problems that can arise in high-speed designs. Source: Cadence*

When returning to a low state voltage, it will "bounce" above that reference ground level, which leads to a noticeable change in ground plane potential near a switching IC. This is known as ground bounce, or simultaneous switching noise (SSN). This voltage level can bounce too high and be mistaken for a false high, resulting in double switching and disrupting overall performance. SSN leads to voltage and timing variations, or jitter, signal distortion, and a host of PI and SI issues. The plane and signal interactions and couplings due to SSN only occur after all of the routing is complete. As a result, the SI/PI engineer must simulate and perform an analysis.

The Solution

The most common solution to this problem is to add the correct number of decoupling capacitors of the correct values between the power and ground planes to make the plane less noisy. This capacitor optimization is a challenge because there are two uncertainties: the effective decoupling radius and the frequency ranges in which the decoupling capacitors are most effective. At this point, simulation software is generally required to minimize the power noise with current continuity so that the planes are close to an ideal reference.

Conclusion

HSD designs have raised the bar on data rate transfers with remarkable results: modern generations of these serial interfaces are many orders of magnitude faster than their previous iterations, with much larger payload sizes. However, these fast-switching networks with high clock rates and timing constraints make SI problems inevitable.

The Cadence Allegro PCB design suite can help you maintain the signal integrity of your high-speed design by applying these best practices:

- Impedance analysis tools to select nets of interest and note the amount of reflection in the signal based on a color-coded scheme.
- IR drop analysis to analyze the PI of a PDN and control ground bounce
- Return path analysis tools to analyze and simulate individual nets and check the integrity of their return path
- Coupling analysis to identify victim and aggressor nets in a table and resolve coupling problems without relying on an SI expert

The SI and PI engineers then need additional tools to ensure their design meets the basic design requirements. The Sigrity SI and PI analysis software enables more comprehensive SI/PI analysis.

References

1. https://www.signalintegrityjournal.com/blogs/12-fundamentals/post/853-back-to-basics-bandwidth-and-rise-time#comments

cadence®

Connect with us