Just a Matter of Time

Beyond Design

by Barry Olney, IN-CIRCUIT DESIGN PTY LTD / AUSTRALIA

Electromagnetic energy propagates at about half the speed of light within the dielectric of a multilayer PCB. This speed is inversely proportional to the square root of the dielectric constant (Dk) of the material. The lower the Dk, the faster the propagation of the wave. In the past, we ignored the board-level delay as it was relatively instantaneous waveform. But now that we have entered the realm of Gigabit/s design, an unaccounted 10 ps of delay can mean the difference between success and absolute failure of a high-speed design. Also, the trend is toward lower core voltages, which conserves power. However, reducing the core voltage also reduces the noise margin and impacts the system timing

Table 1: Overall DDR3-1066 timing budget allowances and resulting margin

Parameter	Setup (ps)	Hold (ps)
Open window from simulations	456	631
SDRAM setup and hold times from datasheets	25	100
Slew rate derating if >1V/ns	2.3	2.8
Timing offset with respect to VREF	13	11
SDRAM derating	88	50
Crosstalk	47	42
Controller error – skew	200	200
Clock error – jitter	30	30
PCB routing tolerance	10	10
Margin	**41**	**185**

Designing a memory interface is all about timing closure. Each signal's timing needs to be compared to the related clock or strobe signal in such a way that the data can be captured on both the rising and falling edge of the strobe—hence the term double data rate (DDR). The increase of data rates to 7800 MT/s for DDR5 has made the timing margin associated with each rising and falling edge even tighter.

Table 1 lists the various timing delays of a DDR3 memory interface running at 533 MHz. After allowing for the chip-level delay, setup and hold times, slew rate derating, clock skew, and jitter we are left with a total margin of just 41 ps on the setup time. Even at this relatively low clock frequency of 533 MHz, 10 ps is all the margin we have left for the board-level delay. Increase the crosstalk or jitter and we are looking at imminent system failure. So, the velocity of propagation of the electromagnetic wave of energy through the PCB is now very significant.

Most systems, whether at the chip or board level, operate synchronously; as such, voltage levels must rise or fall within a specified time or else the circuit will be out of sync and failures will occur. The timing budget tells us how much margin we have, or to put it another way, how much headroom we have before a failure occurs.

So, how do we go from a timing spec (Figure 1) to the actual flight time of the entire memory interface?

The signal and timing, relative to other signals, ride on an electromagnetic carrier wave at various speeds, depending on the surrounding dielectric materials. This energy transports the signal from the driver along the transmission line to the load and does not disrupt the original timing but rather adds the same delay to all the signals that travel the same path. So, one must keep all the relative signals on the same path (layer) or there will be discrepancies when the signals arrive at the load. Alter-

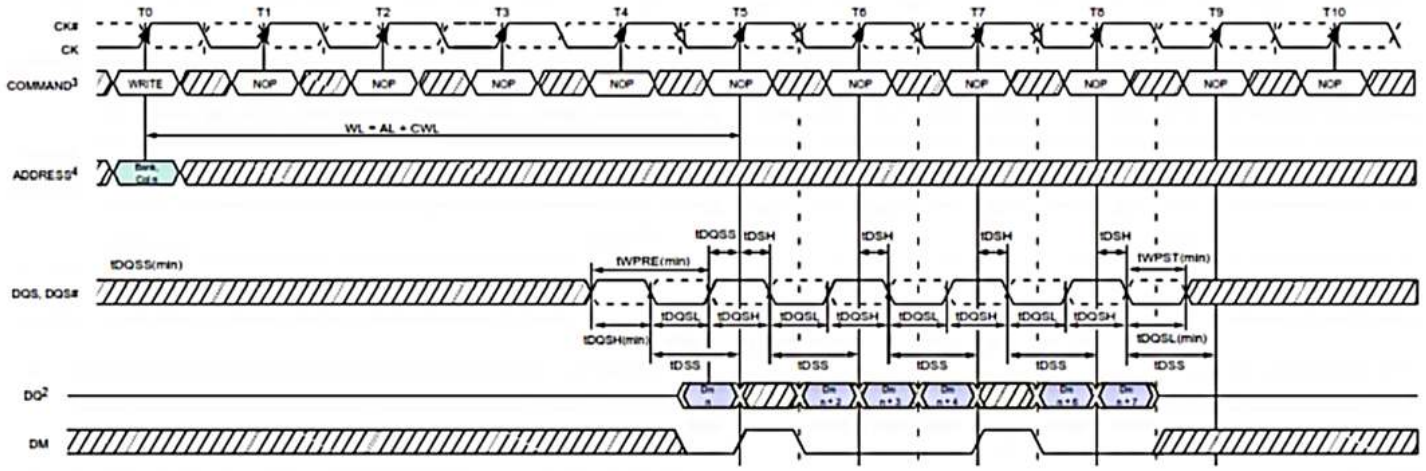

Figure 1: DDR3 synchronous timing. (Source: JEDEC Std 79-F3)

natively, one can compensate for a delayed layer by adding delay intentionally using serpentines to the signal trace in question.

Clocks are essential gatekeepers of the digital domain. Setting the pace for all that follows the clock can be a single trace or differential pairs that carry complementary signals. Because each bit of the data bus must arrive and become stable before the clock cycle, the clock signals establish setup time for accepting or extracting the data. The hold time ensures that the entire bus remains steady during the read-in or read-out of data after the clock.

Fortunately, synchronous buses, as typically used for parallel data signal transfer, benefit from an extraordinary immunity to crosstalk. Crosstalk only occurs when the signals are being switched and this crosstalk only has an impact within a small window around the moment of the clocking. The crosstalk must be specified during the setup (t_S) and hold (t_H) window at the receiver. During this interval, the crosstalk must never drive any valid signal across the receive threshold to the opposite logic state.

So, providing the receiver waits sufficiently long enough for the crosstalk to settle before sampling the data bus, the crosstalk has no impact on the signal quality at the receiver. If the crosstalk arrives during the signal transitions, then its only impact is jitter on the eye. However, this only applies to signals within the same group. Asynchronous and unrelated

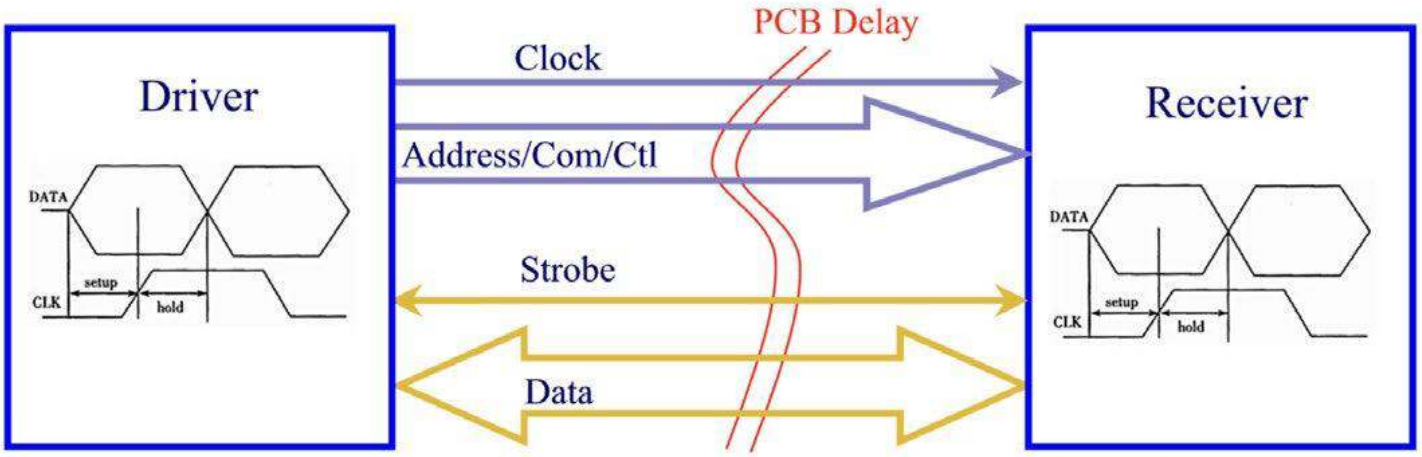

Figure 2: Source synchronous timing relationship across substrate.

signals, on the other hand, always remain sensitive to crosstalk.

From a PCB designer's perspective, we can only optimize what we can control and that is just the placement and routing. Ensuring the route tolerance is less than 10 ps is all we can do to stabilize the design. This may sound tight and difficult to manage but I do this routinely whether called for or not. Put in a little extra effort during layout and one can achieve perfect timing with every design.

A stripline is any trace sandwiched between reference planes on both sides. The electric fields of a stripline are totally contained between the two solid planes, so the speed of propagation for signals guided by the trace is entirely determined by the dielectric constant of the surrounding materials.

On the other hand, a microstrip is any trace fabricated on the surface layers of a PCB. A microstrip has dielectric material and a plane on one side and air on the other. An embedded microstrip is similar but is covered in a conformal coating such as solder mask or another dielectric material. In this case, the effective dielectric constant should be calculated by a field solver and represents a combination of the surrounding materials. There are also other variants of microstrip and stripline, such as build-up microstrip and dual asymmetric stripline.

The electric fields surrounding the microstrip exist partially within the dielectric material(s) and partially within the surrounding air. Since air has a dielectric constant of 1, which is always lower than that of FR-4 (typically 4.3), mixing a little air into the equation will speed up the signal propagation. Even if the trace widths are adjusted on each layer so that the impedance is identical, the propagation speed of the microstrip is always faster than the stripline, typically by 13–17%. The speed of propagation of digital signals is independent of trace geometry and impedance.

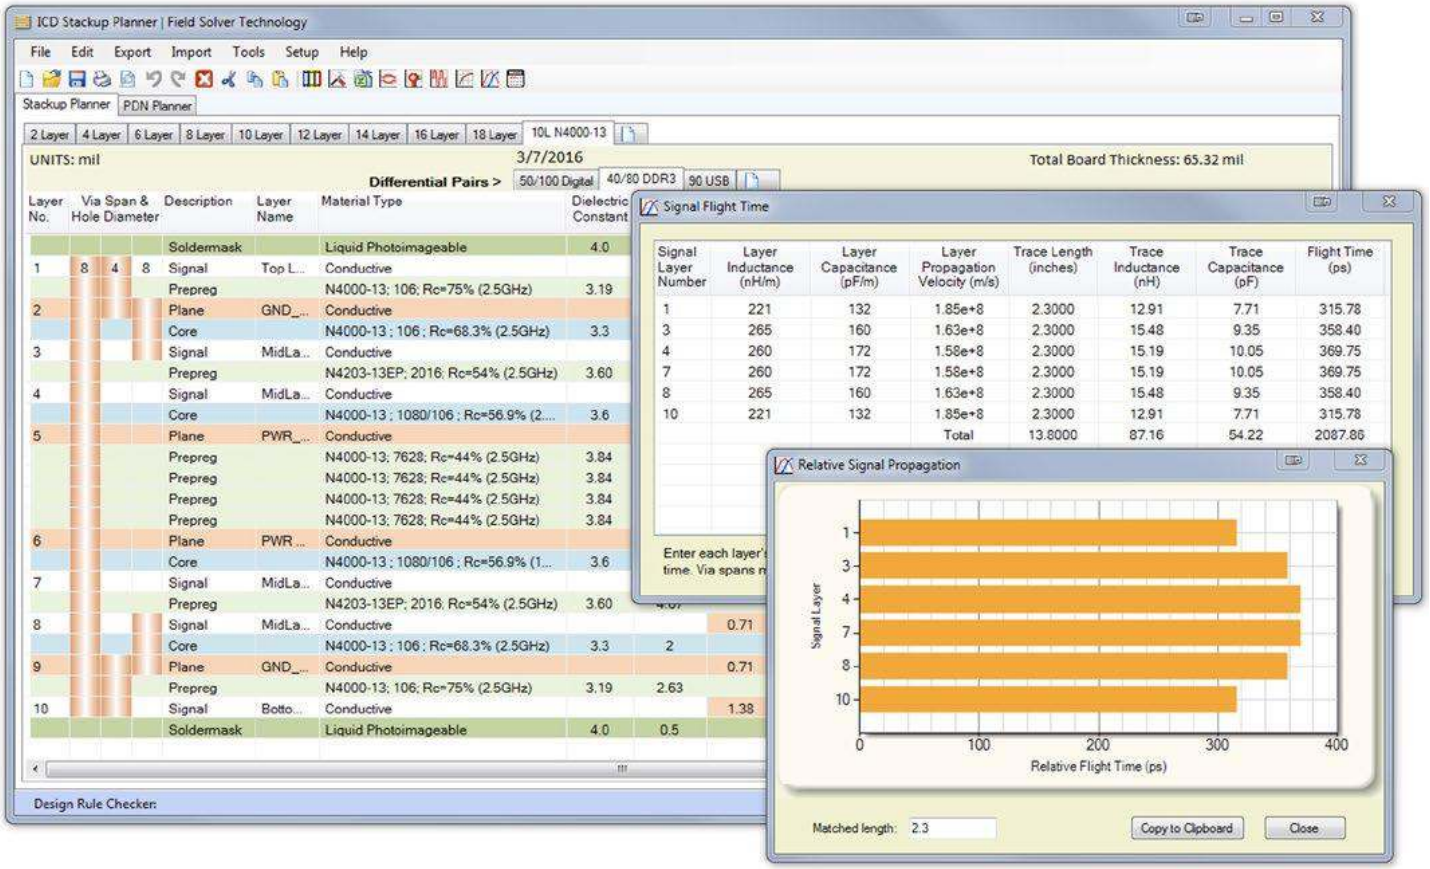

Figure 3: Relative signal propagation for each signal layer on a 10-layer DDR3 stackup. (Source: iCD Stackup Planner)

If you are aware of this issue, then the trace delays (Figure 3) can be matched to compensate for the flight time variance, so that at the nominal temperature all signals running on either microstrip or stripline will arrive at the receiver simultaneously. It is good design practice to route each memory bus on the same layers. That is, avoid routing on the surface layers other than fanout, and drop immediately to the internal stripline layers routing all data, address, associated clocks/strobes, and control signals on the same layer pair to avoid propagation delay variance.

Key Points

- Electromagnetic energy propagates at about half the speed of light within the dielectric of a multilayer PCB.
- The lower the Dk, the faster the propagation of the wave.
- Designing a memory interface is all about timing closure.
- The signal and the timing, relative to other signals, ride on an electromagnetic carrier wave at various speeds, depending on the surrounding dielectric materials.
- This energy transports the signal from the driver, along the transmission line to the load, and does not disrupt the original timing but rather adds the same delay to all the signals that travel the same path.
    - Clocks are essential gatekeepers of the digital domain.
    - Synchronous buses benefit from an extraordinary immunity to crosstalk. Crosstalk only occurs when the signals are being switched and this crosstalk only has an impact within a small window around the moment of the clocking.
    - Asynchronous and unrelated signals always remain sensitive to crosstalk.
- From a PCB designer's perspective, we can only optimize what we can control and that is just the placement and routing.
- Even if the trace widths are adjusted on each layer, so that the impedance is identical, the propagation speed of the microstrip is always faster than the stripline, typically by 13–17%.
- The speed of propagation of digital signals is independent of trace geometry and impedance. **DESIGN007**

Resources

• Beyond Design: "Crosstalk Margins," "DDR3/4 Fly-by vs T-Topology Routing," "Signal Flight Time Variance in Multilayer PCBs," by Barry Olney

• "Board-level timing analysis," Tech Design Forum Techniques.

• "High-speed PCB design timing analysis and simulation strategy," Engineering Technical, PCBway.

• "Understanding Simulation Analysis Parameters for DDR4 Bus Systems in SystemSI," Cadence.

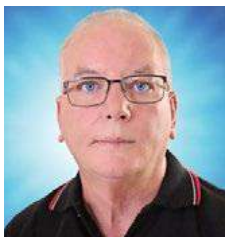

Barry Olney is managing director of In-Circuit Design Pty Ltd (iCD), Australia, a PCB design service bureau that specializes in board-level simulation. The company developed the iCD Design Integrity software incorporating the iCD Stackup, PDN, and CPW Planner. The software can be downloaded at www.icd.com.au. To read past columns, click here.

We DREAM Impedance!

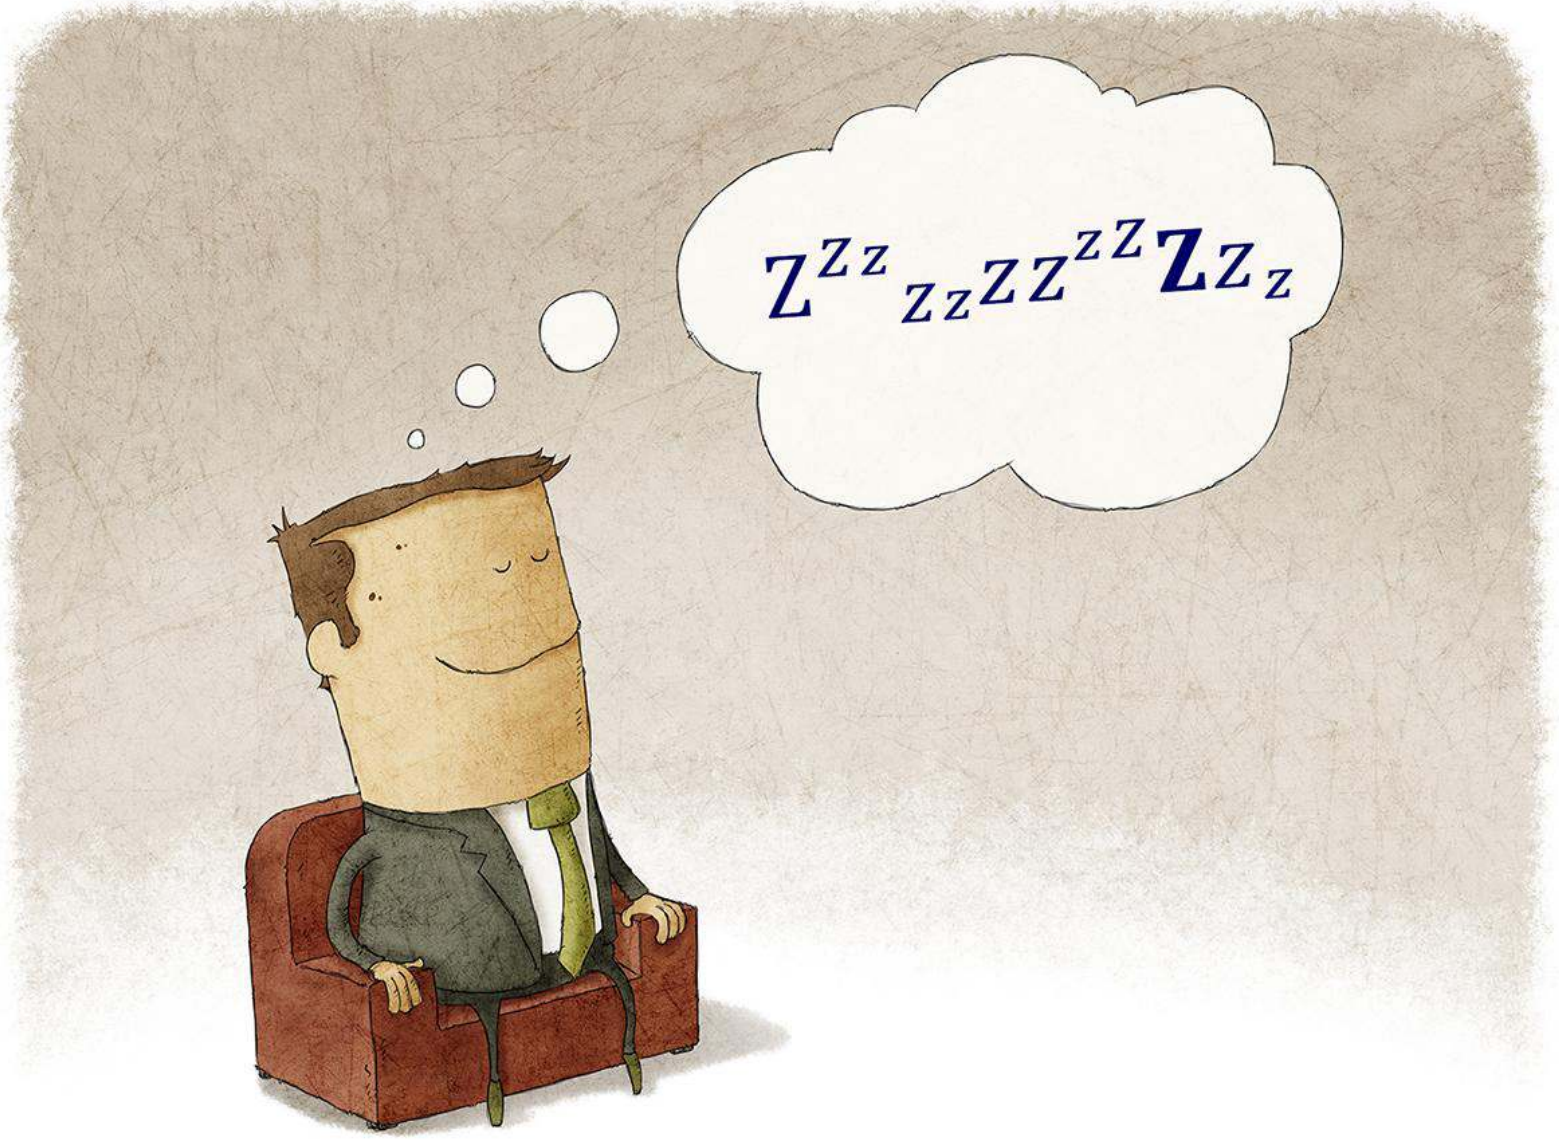

Did you know that two seemingly unrelated concepts are the foundation of a product's performance and reliability?

- Transmission line impedance and
- Power Distribution Network impedance

DISCOVER MORE

iCD software quickly and accurately analyzes impedance so you can sleep at night.

iCD Design Integrity: Intuitive software for high-speed PCB design.

"iCD Design Integrity software features a myriad of functionality specifically developed for PCB designers."
– Barry Olney

iCD
In-Circuit Design Pty Ltd

www.icd.com.au