

BOARD LEVEL SIMULATION SPECIALISTS

ICD Stackup Planner - offers engineers/PCB designers unprecedented simulation speed, ease of use and accuracy at an affordable price

- 2D (BEM) field solver precision
- Characteristic impedance, edge-coupled & broadside-coupled differential impedance
- Unique field solver computation of multiple differential technologies per stackup
- Heads-up impedance plots of signal and dielectric layers
- User defined dielectric materials library - over 16,250 materials up to 100GHz

ICD PDN Planner - analyze multiple power supplies to maintain low impedance over entire frequency range dramatically improving product performance

- Fast AC impedance analysis with plane resonance
- Definition of plane size/shape, dielectric constant & plane separation for each on-board power supply
- Extraction of plane data from the integrated Stackup Planner
- Definition of voltage regulator, bypass/decoupling capacitors, mounting loop inductance
- Frequency range up to 100GHz
- Extensive Capacitor Library – over 5,250 capacitors derived from SPICE models

DDR3/4 Fly-by vs. T-topology Routing

by Barry Olney

IN-CIRCUIT DESIGN PTY LTD AUSTRALIA

JEDEC introduced fly-by topology in the DDR3 specification for the differential clock, address, command and control signals. The advantage of fly-by topology is that it supports higher-frequency operation, reduces the quantity and length of stubs and consequently improves signal integrity and timing on heavily loaded signals. Fly-by topology also reduces simultaneous switching noise (SSN) by deliberately causing flight-time skew, between the address group and the point-to-point topology signals,

of the data groups. To account for this skew, the DDR3/4 controller supports write leveling. The controller must add the write leveling delays to each byte lane to maintain the strobe to clock requirement at the SDRAMs.

T-topology can be challenging to route, particularly double T-topology with four back-to-back SDRAMs as in Figure 1, but it can be advantageous when using multi-die packages. The fly-by topology used in Figure 3 is much easier to route but does not work well with high-capacitance loads, such as LPDDR3 DDP (double die package) and QDP (quad die package) devices. IC fabricators basically arrange dies in parallel (as in Figure 2) to increase package density which can also increase input capacitance by up to four times. Excessive ring-back is often present in the first few nodes of the daisy chain.

This is the reason why the T-topology was developed. However, if you are supporting only SDP (single die package) devices, then the fly-by

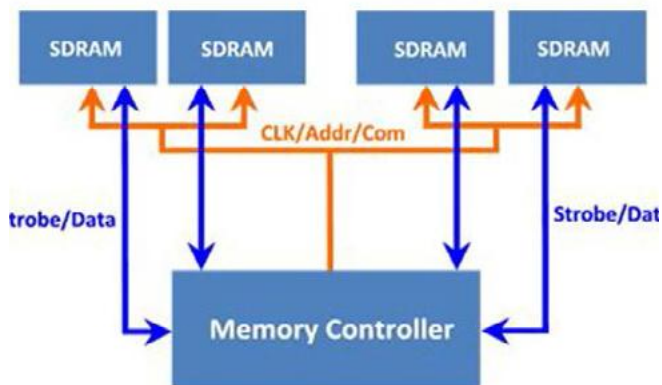



Figure 1: Double T-topology for clock/address/command/control routing.

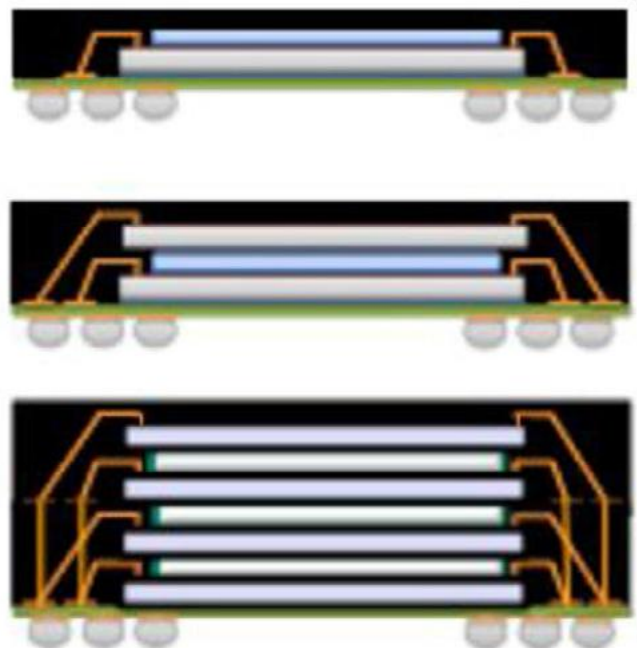


Figure 2: SDP and multi-die DDP and QDP memory devices.

is the most straightforward approach. It doesn't matter which topology you use, though—both fly-by and double T-topologies should work fine. If you are using a DDP device, then double-T topology works better than fly-by in terms of delivering a better system margin.

During a write cycle, using the fly-by topology, data strobe groups are launched at separate intervals to coincide with the clock arriving at memory components on the SODIMM or PCB, and must meet the timing parameter between the memory clock and DQS defined as t_{DQSS} of $\pm 0.25 t_{CK}$. The PCB design process can be simplified using the leveling feature of the DDR3/4. The fly-by, daisy chain topology increases the complexity of the controller design to achieve

leveling but fortunately, greatly improves performance and eases board layout for DDR3/4 designs.

It is not that you have to use fly-by write leveling, because it is a feature of DDR3 and DDR4, but rather that you have to use write leveling in order to allow fly-by routing. There is also no reason not to use the write leveling training for a T-topology in order to optimize the write strobe to clock timing. With this you can adjust slight differences in CA timings and avoid hard coding the skews that you normally have to manually take care of on the strobe to clock delay.

Fly-by topology is similar to daisy chain or multi-drop topology, but has very short stubs, to each memory device in the chain, to reduce the reflections. The double T-topology was used for DDR2 and had a downside in that the impedance discontinuities, due to branching along the traces, caused obvious margin losses. T-topology also tends to have overshoots, while the levels for fly-by are terminated and therefore do not reach the full swing voltage rails. Also, the length of the stubs has an effect on the maximum bandwidth of the transmission line. If you are employing high-frequency DDR4, then the bandwidth of the channel needs to be

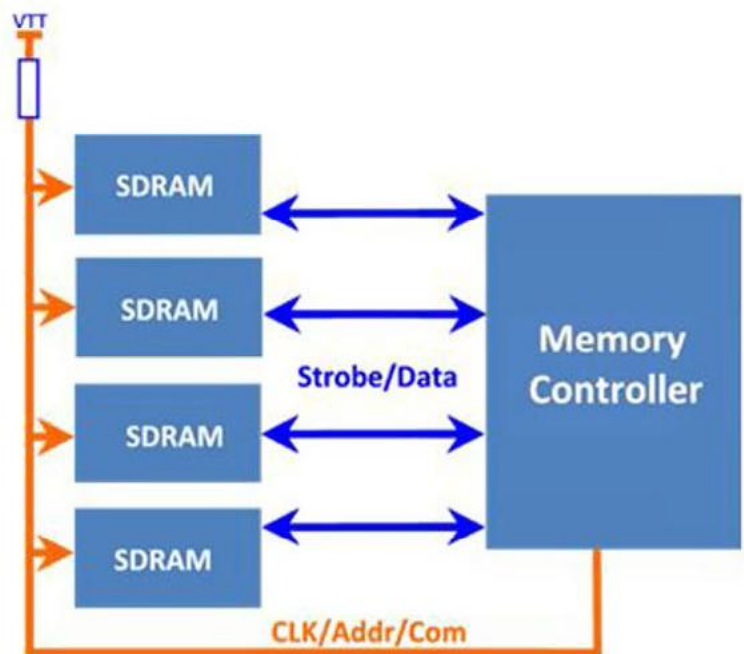



Figure 3: Fly-by topology for clock/address/command routing.

as high as possible. With conventional T-topology, the trace stub is lengthened with an increase in the number of memory device loads. In some cases, there can be as many as eight memory devices connected to the processor. The resonant frequency or bandwidth is inversely proportional to the stub length.

$$f_o = \frac{c}{4 \times (\text{stub length}) \times \sqrt{E_r}}$$

where f_o is the resonant frequency, c is the speed of light and E_r is the dielectric constant

The clock traces should be routed to a longer delay than the strobe traces per byte lane. This is necessary because:

1. The write leveling is capable of adjusting the clock to write data strobe alignment over a wide range, assuming the clock trace has a longer delay than the strobe traces.
2. The read leveling is capable of adjusting the read data eye to read the data strobe over a wide range. The adjustment is per byte, so board skew between the data and data mask signals should be minimized.

3. There is no automatic training for aligning command/address signals to the clock, but a fixed offset is programmable, in the processor, and can be used if necessary. Skew between the clock and address/control signals should be minimized.

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 4266MT/s for DDR4, has made the timing margin associated with each rising and falling edge much tighter. Even though a direct successor is not currently planned, sources speculate that the 5th generation DDR5 will use a serial interface to eliminate the issues associated with parallel busses. Serial busses are easier to scale up and have fewer connections, making PCB design less demanding.

It seems that every datasheet or reference standard you read on DDR design quotes different allowances for the timing budget. At a basic level, the differential clock is the reference signal for the address/control and command

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 CA	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

Table 1: Example of the overall DDR3-1066 timing budget allowances and resulting margin.

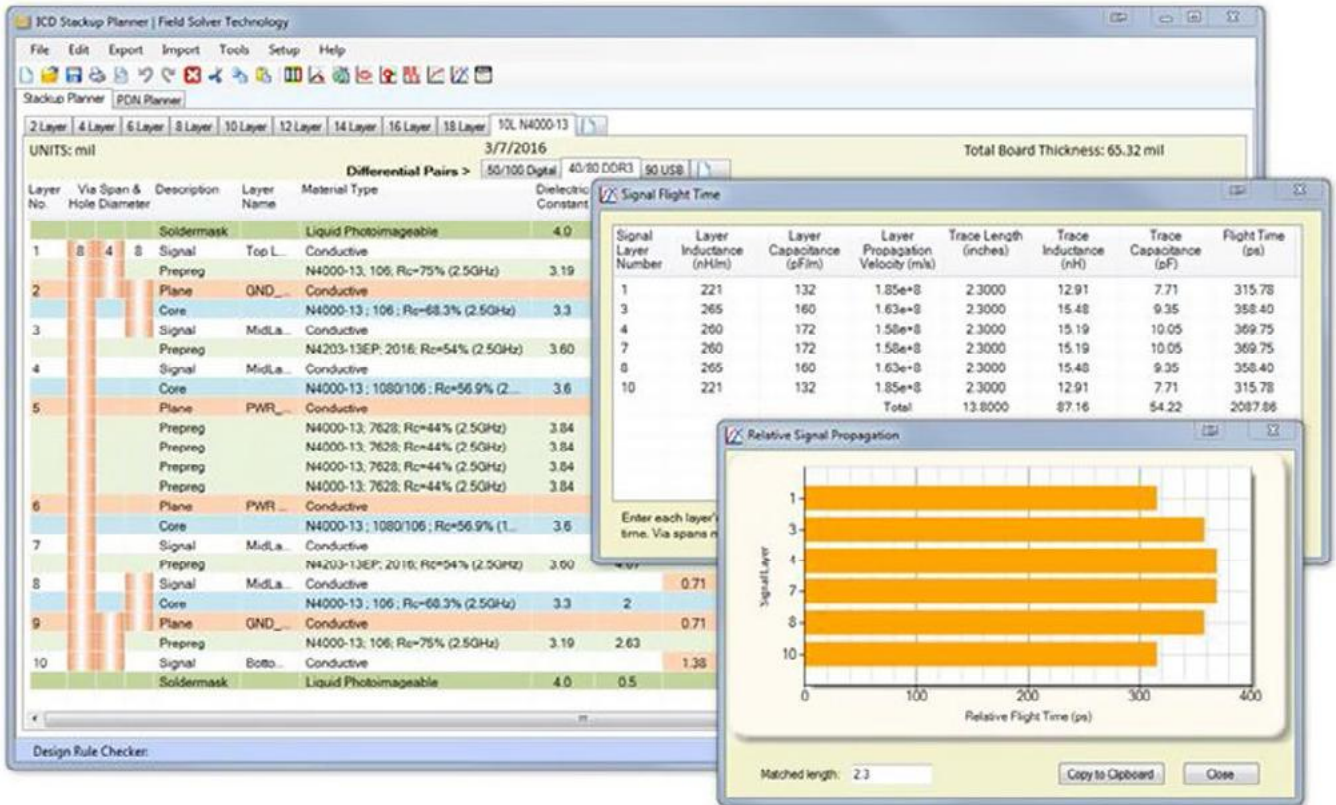


Figure 4: Relative signal propagation for each signal layer on a 10-layer DDR3 stackup.

signals. Whereas, the differential strobe is the clock for the data and data mask signals. The timing budget for the data byte lanes and the address group need to be determined and must be spread across the processor package, PCB interconnect and the SDRAM packages. The portion of the timing budget, consumed by the controller IC and SDRAM devices, is fixed and cannot be influenced by the PCB designer. The amount of timing budget remaining, after subtracting these fixed portions, is all that is left for the board interconnect—which is not much!

For a DDR3-1066 SDRAM for instance, data from the JEDEC, JESD79-3E DDR3 standard specifies 25ps for Setup and 100ps for Hold time as in Table 1. Ideally, one should use a simulation tool, such as HyperLynx, to measure the setup and hold times to ensure they are within the timing budget. However, if you do not have access to an analysis tool, work on 10ps delay, for the routing tolerance, and you can be assured that you are within the margin al-

lowing for any derating. That is, providing the transmission lines are matched to 40/80 ohms single-ended/differential impedance, the correct drive currents are being used and the waveforms are not distorted. Let's face it, it is not that difficult to route each signal to the exact propagation delay given that you have access to each layer's flight time. It is also worth noting that the margin limits can be increased, if the memory interface is not operating at the maximum frequency and/or if a fast memory device is used.

Now let's consider a typical 10-layer DDR3 stackup as in Figure 4. There are six routing layers and all the DDR3 signals are routed to 40/80 single ended/differential impedance and matched to 2.3 inches in length. In my 2014 column [Matched Length does not equal Matched Delay](#), I cited the difference between the propagation delay of the layers on a PCB. The most dramatic is that of microstrip (outer layers) compared to stripline (inner layers). In this case, the delta between

layers 1 and 4 is a massive 54ps—way outside the setup margin.

Whilst stripline layers 3 and 4 have an 11ps difference, even though they are routed to the same length. This is due to the variance in dielectric constant of each layer which changes the velocity of the signals propagation. The difference is graphically displayed in the ICD Stackup Planner's new Relative Signal Propagation dialog. Even the 11ps stripline variance is more than enough to offset the timing, particularly using high-speed DDR3 and DDR4 devices, regardless of the routing topology.

In conclusion, fly-by topology supports higher frequency operation, reduces simultaneous switching noise, reduces the quantity and length of stubs and consequently improves signal integrity and timing. And, most importantly, from a PCB designer's point of view, it eases routing of memory devices dramatically. However, no matter what topology is implemented, one should pay strict attention to the signal propagation, on each layer, ensuring the total flight time of the critical signals match, regardless of length.

Points to Remember

- Fly-by topology supports higher frequency operation, reduces simultaneous switching noise, reduces the quantity and length of stubs and consequently improves signal integrity and timing.
- The controller must add the write leveling delays to each byte lane to maintain the strobe to clock requirement at the SDRAMs.
- T-topology can be challenging to route but it can be advantageous when using multi-die packages with high capacitance loads. Whereas, fly-by topology eases routing of DDR3/4 devices.
- The double T-topology was used for DDR2 and had a downside in that the impedance discontinuities, due to branching along the traces, caused obvious margin losses.
- With conventional T-topology, the trace stub is lengthened with an increase in number of memory device loads.
- The clock traces should be routed to a longer delay than the strobe traces per byte lane.

- Designing a memory interface is all about timing closure.
- The portion of the timing budget, consumed by the controller IC and SDRAM devices are fixed and cannot be influenced by the PCB designer.
- If you work on 10ps, then you can be assured you are within the margin allowing for any derating.
- The margin limits can be increased, if the memory interface is not operating at the maximum frequency and/or if a fast memory device is used.
- Matched length does not equal matched delay. The most dramatic difference being that of microstrip (outer layers) to stripline (inner layers). **PCBDESIGN**

References

1. Barry Olney's Beyond Design columns: [PCB Design Techniques for DDR, DDR2 & DDR3 Parts 1 & 2](#); [Matched Length does not equal Matched Delay](#)
2. JEDEC Specifications JESD 79F, JESD79-2E, JESD79-3F & JESD79-4
3. [The Xilinx Zynq-7000 PCB Design Guide](#)
4. EDN article by Chang Fei Yee: [DDR4 memory interface; Solving PCB design challenges](#)
5. Micron's [TN-41-08: Design Guide for Two DDR3-1066 UDIMM Systems Introduction](#)
6. The SI List, available at [Freelists.org](#)
7. The ICD Stackup and PDN Planner software can be downloaded from [www.icd.com.au](#)



Barry Olney is managing director of In-Circuit Design Pty Ltd (ICD), Australia. This PCB design service bureau specializes in board-level simulation, and has developed the ICD Stackup Planner and ICD PDN Planner software. To read past columns, or to contact Olney, [click here](#).