DDR3/4 Fly-by Topology Termination and Routing

Beyond Design

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

DDR3/4 fly-by topology is similar to daisy chain or multi-drop topology, but it includes very short stubs to each memory device in the chain to reduce the reflections. The advantage of fly-by topology is that it supports higher-frequency operation and improves signal integrity and timing on heavily loaded signals. If you are employing high-frequency DDR4, then the bandwidth of the channel needs to be as high as possible. However, with today's extremely fast edge rates, the sequencing of the stubs and the end termination, and the associate load, can make a measurable difference in signal quality. In this month's column I will look at how best to route DDR3/4 fly-by topology.

Reflections occur whenever the impedance of the transmission line changes along its length. This can be caused by unmatched drivers/loads, layer transitions, different dielectric materials, stubs, vias, connectors, terminations and IC packages. By understanding the causes of these reflections and eliminating the source of the mismatch, a design can be engineered with reliable performance. For perfect transfer of energy and to eliminate reflections, the impedance of the source must equal the impedance of the trace, as well as the impedance of the load.

As signal rise times increase, consideration should be given to the propagation time and reflections of a routed trace. If the propagation time and reflection from source to load are longer than the edge transition time, an electrically long trace will exist. If the transmission line is short, reflections still occur but will be overwhelmed by the rising or falling edge and may not pose a problem. But even if the trace is short, termination may still be required if the load is capacitive or highly inductive to prevent ringing.

Series termination is an excellent strategy for point-to-point routes, one load per net.

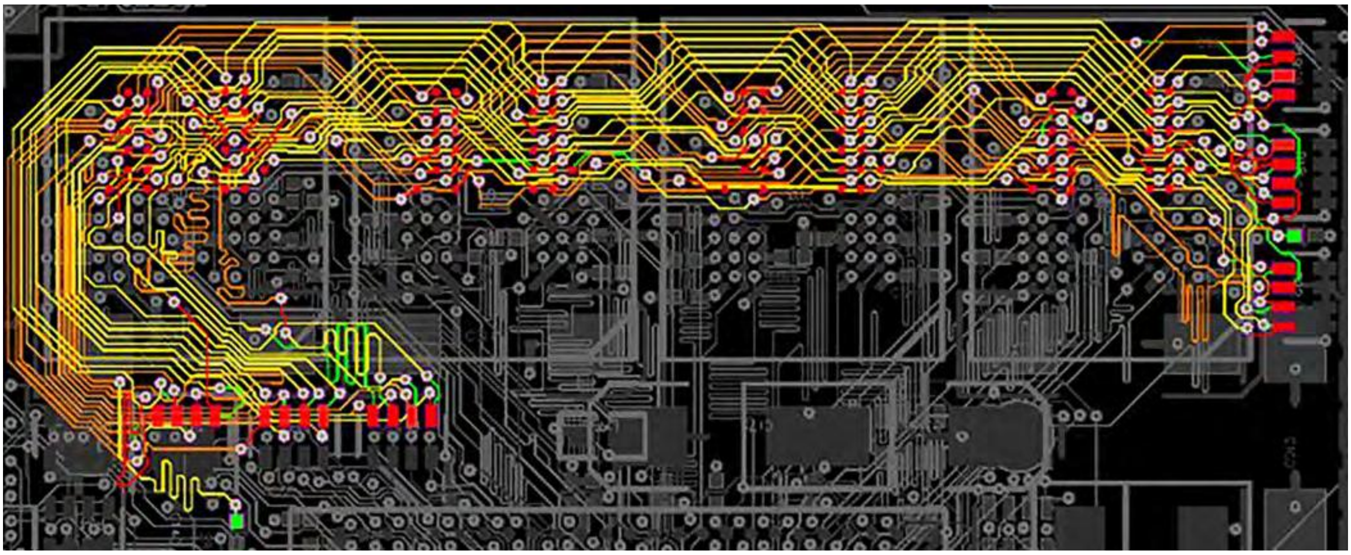

Figure 1: Fly-by topology for clock, address and command routing.

Whereas parallel (end termination), is preferred for busses with a number of loads in a multi-drop topology. For DDR3/4 layouts, a series termination is generally not required for on-board memory devices. However, if your design has plug-in memory then the data and data mask signal length may be excessive and require a series termination.

With the fly-by address, control and command (ACC) signals, the traces should be routed as practicable as possible to the memory device pins and the parallel termination placed at the end of the line as in Figure 1. Short stubs can be used to connect the passing signal to each memory device in sequence but the longer the stubs the higher the capacitance. This stub capacitance, along with the parasitic input capacitance of the receiver pin, creates an imperfection in the termination network.

When implementing parallel terminations, it is not always possible to place the termination after the final load in a multi-drop topology due to real estate restrictions. Figure 2 shows the schematic of this configuration. If the stub is very short and the signal edge rate is not excessively fast, then this may be acceptable. However, as the edge rate increases, the extra capacitance of the stub is significant and creates reflections.

When the first incident wave arrives from the driver, part of that wave, which is a small pulse, bounces off the imperfection and returns to the driver. This pulse bounces again off the low-output impedance driver and returns one round trip later to the receiver. What we observe at the receiver is the initial rising edge, followed one round trip later by a secondary pulse. If the initial reflected pulse is sufficiently

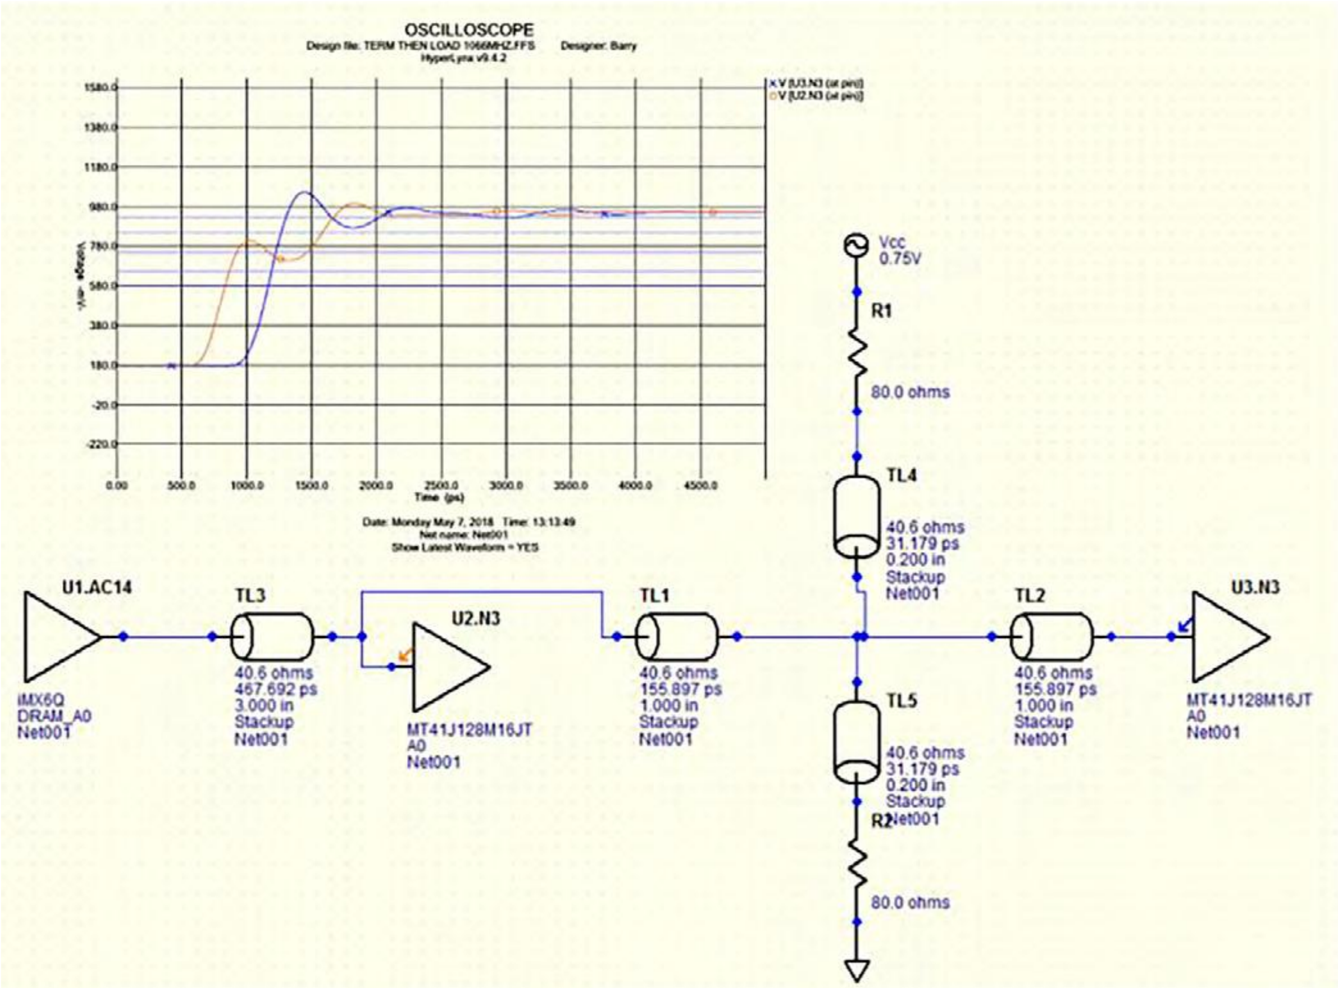

Figure 2: DDR3 address signal with termination before final load (all simulations performed in HyperLynx).

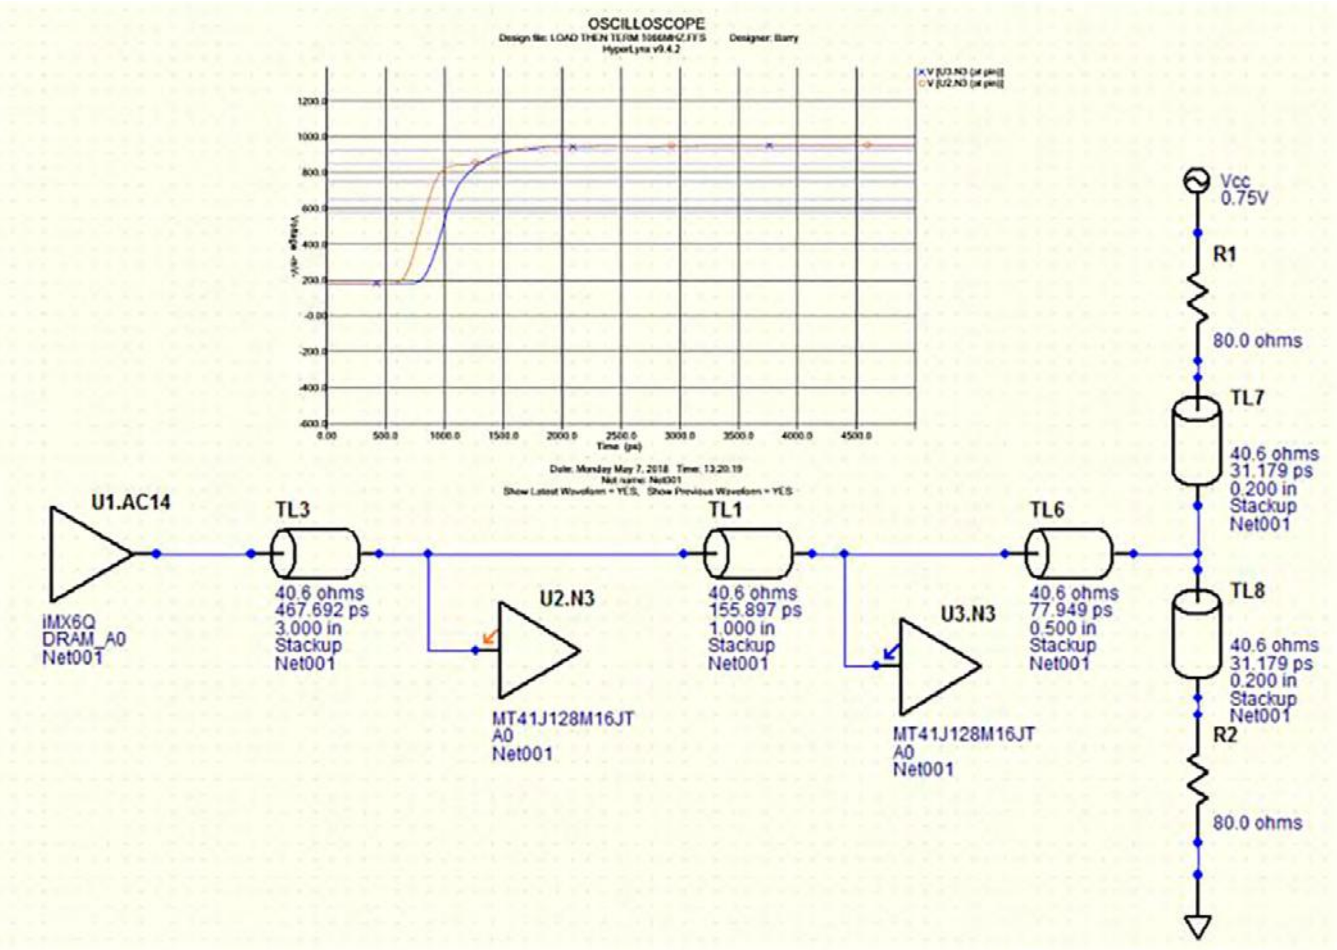

Figure 3: DDR3 fly-by topology with termination after final load.

small, all subsequent reflections will be negligible. However, in this case, the reflection causes a non-monotonic edge at the first receiver (orange waveform), which could lead to false triggering of the receiver. If for example, there are four memory chips on the address line, then this would occur on all receivers except the last to a similar extent.

Figure 3 illustrates a typical DDR3 fly-by topology with the termination at the very end of the final load. Also, the passing address signal trace goes directly to the receiver pins with no stub. This is the ideal scenario. In this case, there are no reflections, from the termination, which can be seen from the waveforms.

Now let's add a half inch stub from the passing address signal trace going to each receiver input pin, as I have seen many designers do. Figure 4 shows the schematic and the resultant waveforms. The reflections created by the additional capacitance, of the stubs, is not quite as bad as that of Figure 2, which had the termination before the final load. However, this demonstrates the impact of stubs on the ACC signals in the fly-by topology. Figure 5 shows the jitter and noise created by the half inch stubs on the signal.

After running a few more, quick simulations in HyperLynx LineSim with varying stub lengths, I found that the stub length can be no more than 200 mil, to alleviate the impact of the reflections, in this (typical) case. This results in the reasonably clear eye diagram of Figure 6.

In conclusion, DDR3/4 fly-by parallel terminations should be placed after the last receiver at the end of the transmission line and routed directly to the last receiver. Since there is no reflection from this topology, this implies that the length, from last receiver to the termina-

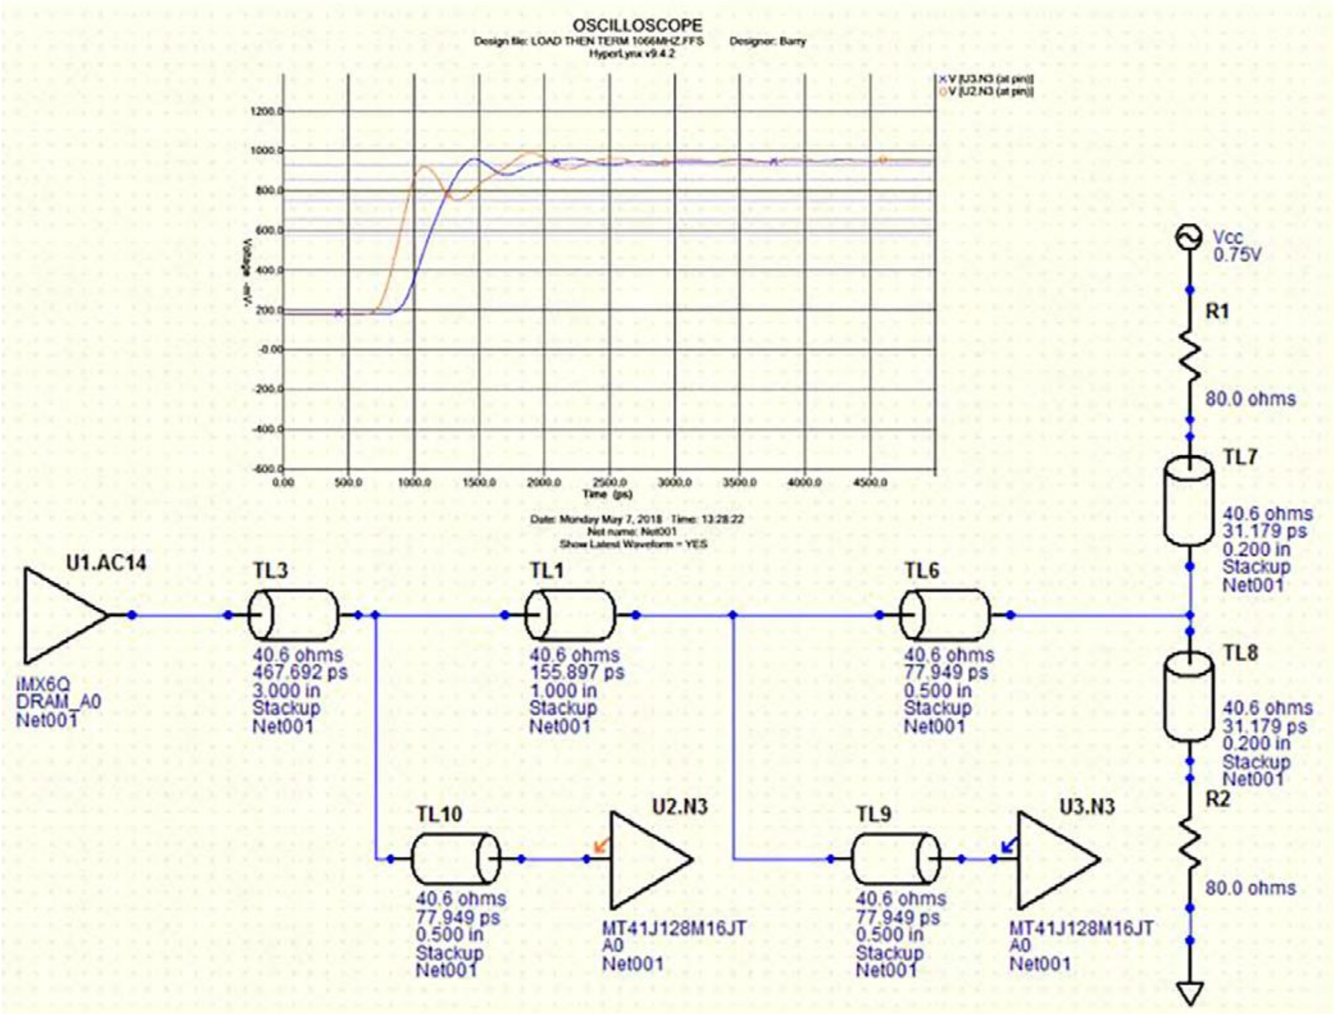

Figure 4: DDR3 fly-by topology with end termination and 500 mil stub.

tions, does not need to be matched to the delay/length of the other signals in the ACC group. Also, further simulations confirmed that the length can be rather long (up to 3 inches) without any noticeable degradation of signal quality. Stubs to each receiver input, from the passing signal, should be kept below 200mil in length to maintain a clear eye.

Key Points:

- Fly-by topology supports higher-frequency operation and improves signal integrity and timing on heavily loaded signals.
- The sequencing of the stubs and the end termination and the associate load can make a measurable difference in signal quality.
- Reflections occur whenever the impedance of the transmission line changes along its length.
- If the propagation time and reflection from source to load are longer than the edge transition time, an electrically long trace will exist.
- If the trace is short, termination may still be required if the load is capacitive to prevent ringing.
- Series termination is an excellent strategy for point to point routes, one load per net.
- Parallel (end termination) is preferred for busses with a number of loads in a multi-drop topology.
- Address, control and command (ACC) signals, traces should be routed directly to

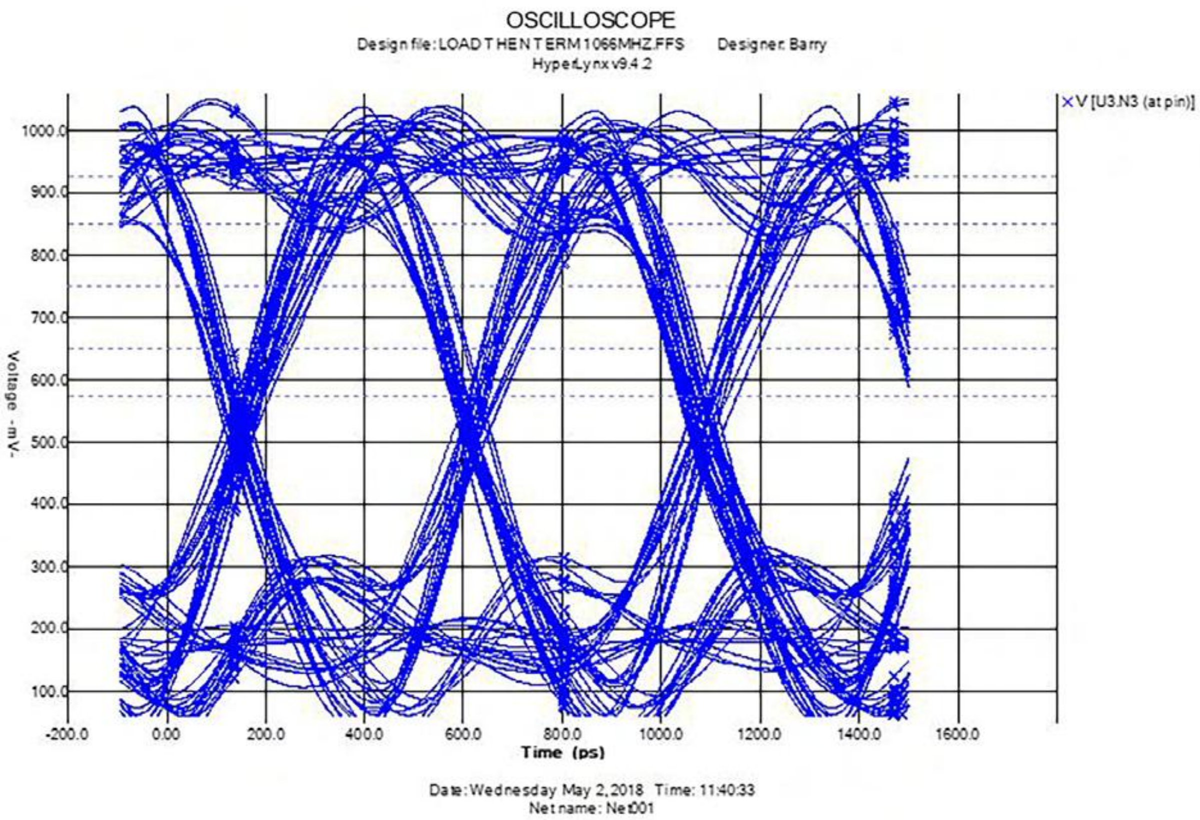

Figure 5: DDR3 fly-by topology with 500 mil stubs to receiver inputs.

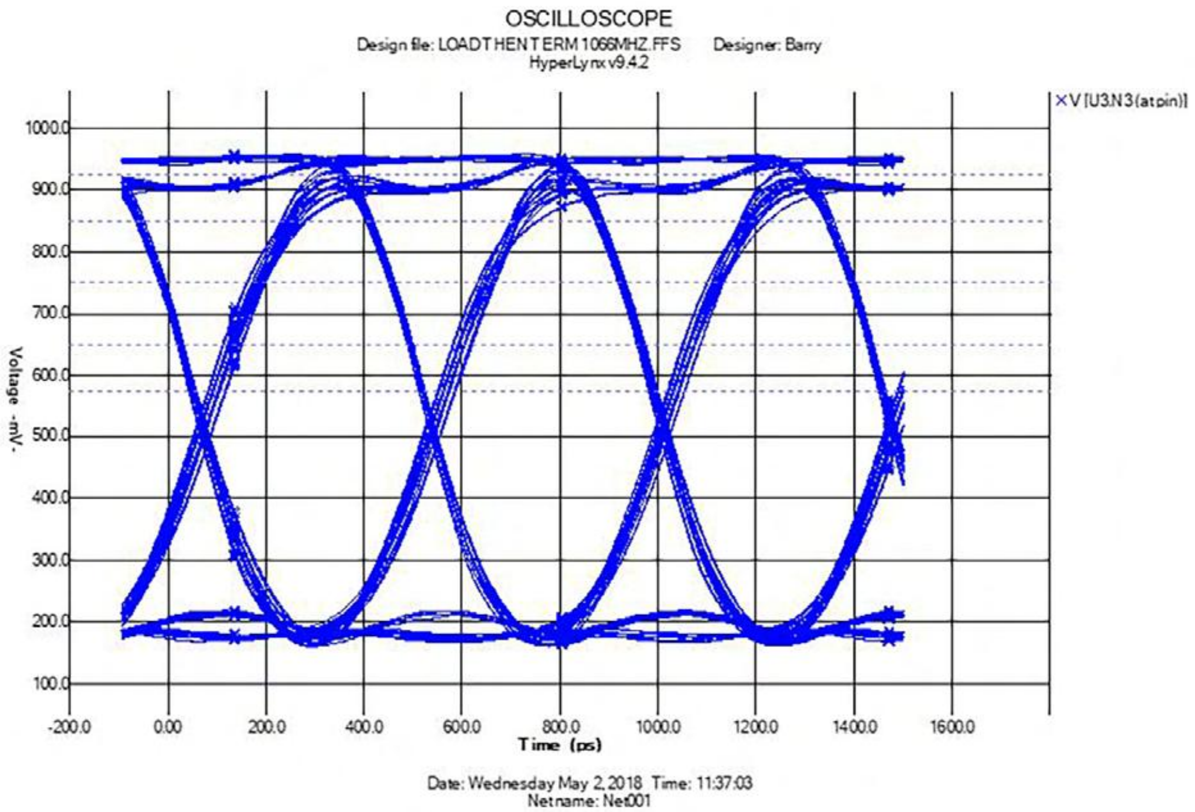

Figure 6: DDR3 fly-by topology with 200 mil stubs to receiver inputs.

the memory device pins and the parallel termination placed at the end of the transmission line.
- Stub capacitance, along with the parasitic input capacitance of the receiver pin, creates an imperfection in the termination network.
- There are no reflections when the passing address signal trace goes directly to the receiver pins with no stub and the termination is at the very end of the line.
- Stubs create reflections and have a detrimental impact on the address signal in the fly-by topology.
- Simulations suggest that the stub must be a maximum of 200mil to alleviate the impact of the reflections.

References:

1. Barry Olney's Beyond Design: DDR3/4 Fly-by vs T-topology Routing, Impedance Matching: Terminations.
2. High-Speed Signal Propagation, by Howard Johnson.

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 from www.icd.com.au. To contact Olney, or read past columns, click here.

Silicon Provides Means to Control Quantum Bits for Faster Algorithms

Research groups from Purdue University, the Technological University of Delft, Netherlands and the University of Wisconsin-Madison have discovered that silicon has unique spin-orbit interactions that can enable the manipulation of qubits using electric fields, without the need for any artificial agents.

"Qubits encoded in the spins of electrons are especially long-lived in silicon, but they are difficult to control by electric fields. Spin-orbit interaction is an important knob for the design of qubits that was thought to be small in this material, traditionally," said Rajib Rahman, research assistant professor in Purdue's School of Electrical and Computer Engineering.

The strength of spin-orbit interaction, which is the interaction of an electron's spin with its motion, is an important factor for the quality of a qubit. The researchers found more prominent spin-orbit interaction than usual at the surface of silicon where qubits are located in the form of so-called quantum dots–electrons confined in three dimensions. Rahman's lab identified that this spin-orbit interaction is anisotropic in nature, meaning that it is dependent on the angle of an external magnetic field, and strongly affected by atomic details of the surface.

The researchers published their findings on June 5 in Nature Partner Journals - Quantum Information. The Wisconsin-Madison team fabricated the silicon device, the Delft team performed the experiments and the Purdue team led the theoretical investigation of the experimental observations.

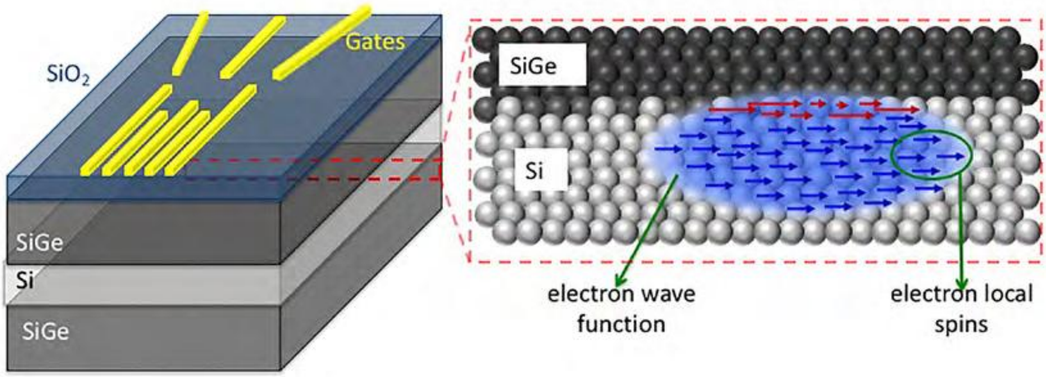

Upcoming work in Rahman's lab will focus on taking advantage of the anisotropic nature of spin-orbit interactions to further enhance the coherence and control of qubits, and, therefore, the scaling up of quantum computer chips.