

Technical Note

DDR2 (Point-to-Point) Package Sizes and Layout Basics

Introduction

Point-to-point designers face many challenges when laying out a new printed circuit board (PCB). The designer may need to arrange groups of devices within a certain area of the PCB, to place components away from critical keep-out zones, include additional solder pads for future component upgrades, or even ensure the design is robust enough to support a potentially faster memory speed. Additional features—like test points, sockets, and alignment holes—may also be needed. Combining these requirements with an already complex PCB floor plan can make a point-to-point layout extremely difficult to complete.


DDR2 SDRAM memory can make floor planning easier with the use of JEDEC-standard FBGA packages. The predefined ball out with a simple addressing scheme that supports all densities and configurations helps make it easy to design for future memory upgrades. Another layout advantage of DDR2 memory is support for on-die termination (ODT). ODT improves signal quality while eliminating most of the external parallel termination resistors. By eliminating termination resistors for memory I/O, available board space is increased and the number of through-hole vias is reduced.

This document does not go into detail on PCB stack-ups, types of traces (stripline vs. microstrip), or topology characterization for the various nets. This technical note does provide general guidelines for developing the PCB floor plan, points out some of the key features of DDR2 technology, and identifies what to consider when starting a new point-to-point design that uses DDR2 SDRAM devices. A Micron DDR2 design guide containing detailed routing information will soon be available.

Getting Started—Understand the Packages

One of the first steps to a good design is understanding the package variations that may exist between DDR2 devices. Although all standard packages fall within the JEDEC-defined limits, various package sizes and ball arrays are available. For example, x16 components use either an 84-ball array or a 92-ball array, depending on the overall package size. Likewise, the x4 or x8 configurations use either a 60-ball array or a 68-ball array. Both the 92-ball and 68-ball packages include outrigger balls in each of the four corners. These additional outrigger balls provide mechanical support for the larger package sizes. The center section, or electrical ball array, is identical for the package with outriggers and the package without outriggers. The outrigger balls are true no connects and don't have any electrical connections to the substrate.


Figure 1: 60-Ball and 68-Ball Package Compatibility – used only for x4/x8 devices



PCBs using DDR2 x4 or x8 devices should be designed to accept both 60-ball and 68-ball packages

Actual package size can vary by density, vendor, or configuration. For example, the Micron 512Mb, 8-bit-wide DDR2 device is 10mm (wide) x 10mm (tall) where another vendor may produce their 512Mb, 8-bit-wide DDR2 device in an 9mm x 11.5mm package. The two packages have approximately the same area, but one is wider and shorter, and the other is narrow and taller. Even though both packages have the identical 60-ball electrical array, the finished board layout may not be able to accommodate both vendor packages if the layout designer did not compensate for the slight variation between these package sizes. For a look at other package-related design considerations, see Micron technical note TN-47-08, “DDR2 Package Sizes and Layout Requirement.” The subjects covered include the maximum, JEDEC-allowable package sizes; a common landing pattern, which was developed to accept all possible package combinations; and a functional cross-reference table for all signals/balls.

Figure 2: 84-Ball and 92-Ball Package Compatibility – used only for x16 devices



PCBs using DDR2 x16 devices should be designed to accept both 84-ball and 92-ball packages

Special Signals and Designing for Flexibility

DDR2 SDRAM has many more features than previous DRAM memory, and to accommodate these features, several balls are multifunctional. Layout is easier if these unique multifunctional balls are identified before starting the design (see Table 1 on page 4).

The complement strobe (DQS#) is an optional signal which is only active when the DDR2 SDRAM is enabled for differential strobes. If differential strobes are not enabled, this ball becomes a not-used (NU) connection. A NU ball is defined as having an internal connection to the die, but does not have any function.

DDR2 (x8) memory supports a redundant set of strobes to allow the x8 configuration to replicate two x4 DDR2 devices. These two signals – redundant strobe (RDQS) and complement redundant strobe (RDQS#) – are optional and are only active if enabled. RDQS is enabled in the extended mode register (EMR), with bit 11 and RDQS# is enabled by EMR bit 10. If they are not enabled, these two balls become NU signals.

Unless the design is for the maximum-density DDR2 part, it is essential to incorporate a few additional signals to guarantee future flexibility for higher density devices. The addressing scheme is straightforward, and support for higher density parts is easy if the additional addresses are available at the package. These key addresses include:

- Bank address 2 (BA2) is needed for all 8-bank devices, including the 1Gb- and 2Gb-density parts
- Row/Column address 13 (A13) is required on x4/x8 devices starting with the 512Mb density and on the 2Gb x16 device
- Row/Column address 14 (A14) is required on the 2Gb x4/x8 devices

Each of these address pins is considered a no connect (NC) on the devices that do not require them. As an NC pin, there is no internal connection to the die. Micron suggests routing these pins to the memory controller so higher density parts can be supported. If the memory controller does not support these higher order addresses, they could be tied to ground through a resistor. This will allow the higher density parts to function when soldered onto the PCB.

Note that if these pins are tied to ground the higher density parts will function, but the density is limited to the number of address pins actually driven by the controller. Also note the higher density parts may have slightly different refresh timing requirements.

Table 1: Key Signal Functionality Table

Signal Name	Pin Nomenclature	Signal Type	Function
A13	NC	No connection to internal die. Okay to run trace to PCB pad.	Used on x4/x8 512Mb and 1Gb devices and all configurations of the 2Gb or 4Gb.
A14	NC	No Connection to internal die. Okay to run trace to PCB pad.	Used on x4/x8 2Gb devices and all 4Gb configurations.
A15	NC	No Connection to internal die. Okay to run trace to PCB pad.	Used on 4Gb (x4/x8) devices.
BA2	NC	No Connection to internal die. Okay to run trace to PCB pad.	Used on all configurations of the 1Gb, 2Gb, and 4Gb.
DQS#	NU	May have internal connection to die. If not enabled DO NOT USE.	Optional for use on the x4 and x8 configurations. Only use when the differential strobe function is enabled.
UDQS#, LDQS#	NU	May have internal connection to die. If not enabled DO NOT USE.	Optional for use on the x16 configurations. Only use when the differential strobe function is enabled.
RDQS	NU	May have internal connection to die. If not enabled DO NOT USE.	Optional for use on the x8. Only use when the redundant strobe function is enabled.
RDQS#	NU	May have internal connection to die If not enabled DO NOT USE.	Optional for use on the x8 configuration. Only use when the redundant strobe and differential strobe functions are enabled.


Design for Debug and/or Test

Unlike previous DRAM technologies where most devices were offered in a TSOP or SOJ package and all pins were exposed for external test points, DDR2 devices are only available in FBGA packages. The FBGA package uses very small solder balls on the bottom of the device to make both an electrical and physical connection to the PCB. Once the FBGA package is soldered to the PCB, the interconnections are inaccessible. Due to these hidden nodes on DDR2 devices, it is virtually impossible to access the DRAM signals unless the board is well designed. In addition, because of DDR2's high speeds, including trace stubs and large test points could affect the quality of individual signals.

Initial debug and board prototyping needs are different from long term manufacturing requirements. For the initial board debug and for verifying signal quality, a lab technician can scrape the protective coating off of a trace or solder a thin wire into an exposed via to make a reliable scope/probe connection. But, in a production environment, most test engineers are going to want to verify the integrity of the solder joints by using a "bed-of-nails" type of tester. All of these requirements can be achieved with careful planning.

For example, make all nodes accessible, at least for the key signals (clocks, strobes, commands, and control), and do not use blind or buried vias. When possible, keep signals on an outer layer and place small test points in-line rather than adding stubs to traces (see Figure 3). Where applicable, use termination resistors as test points.

Figure 3: Typical Routing of Test Points for DDR2 Devices



If inner layers are required for routing, try to include the test points at one of the existing vias. Staggering nodes between the top- and bottom-sides of the PCB will help make room for test points. For example, place the test point for address (A2) near one device and place the test point for address (A0) near a different device.

Stub Series Terminated Logic (SSTL₁₈)

DDR2 uses SSTL₁₈ logic for all inputs. A simple way to describe SSTL is that each input is a differential receiver with one leg tied to a reference voltage (V_{REF}). When the voltage offset between V_{REF} and the input voltage exceeds the threshold of the input receiver, the differential receiver will drive to either a HIGH or LOW logic level.

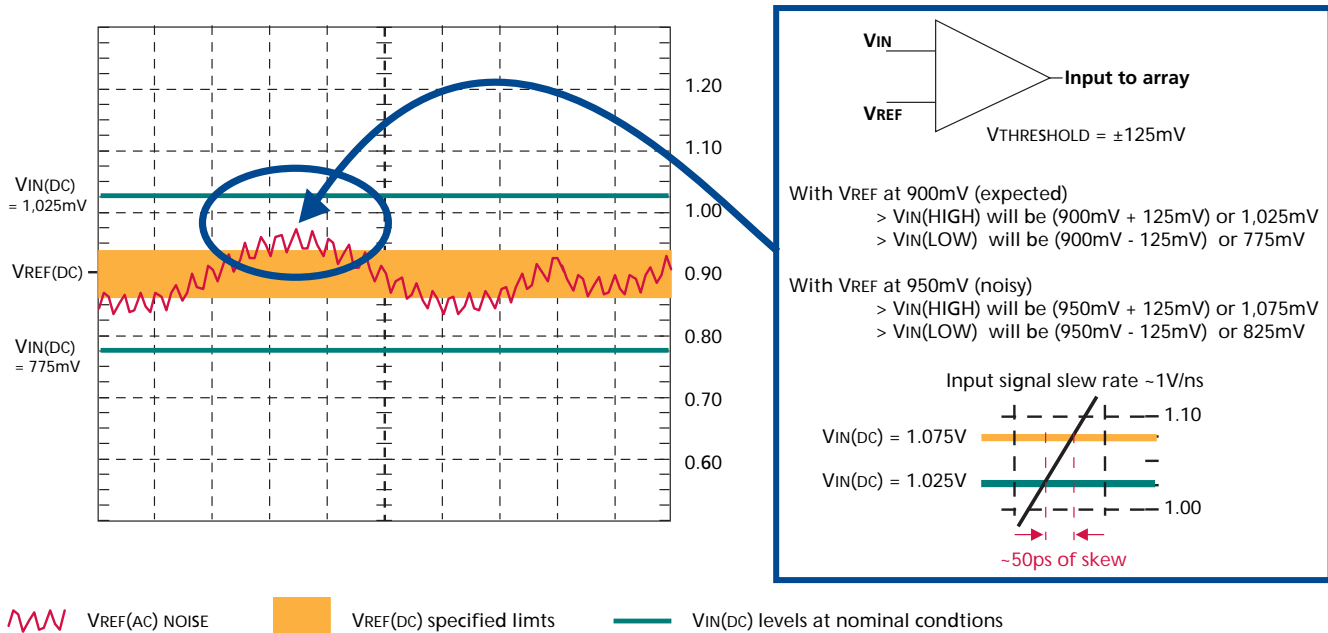
Due to the fast switching requirements of DDR2, the DC threshold is 125mV. This means that when the input to the receiver approaches $V_{REF} + 125mV$ or $V_{REF} - 125mV$, the device will start to detect a logic level. By specification, V_{REF} equals 50 percent of V_{DDQ} and it must track the slight variations in V_{DDQ} to within 1 percent of the nominal value. The AC portion of V_{REF} is limited to 2 percent of the DC nominal value.

As can be seen in Figure 4 on page 6, it is imperative to keep V_{REF} free of noise and within the specified values. If V_{REF} is noisy or is not centered at $V_{DDQ}/2$, the switching points of the receiver may change and may drastically affect the input receiver's setup and hold timing in addition to other DRAM timing.

Figure 4: Setup and Hold Margins can Significantly be Reduced When VREF is Noisy

With the ideal conditions, the receiver will detect a high logic level at 1,025mV and a low logic level at 775mV.

If VREF does not track VDDQ or if VREF is noisy, the receiver may detect logic levels at different points. There could also be timing skew.



General DDR2 Layout Guidelines

A successful layout requires the use of good design rules, an understanding of critical device parameters, and the verification of signal integrity and timing through simulation.

The following topics should familiarize the layout engineer with the primary factors to consider when starting a new DDR2 project.

Match the I/O Drive

DDR2 SDRAM has two output drive levels: full drive strength and reduced drive strength. The full drive has a target output impedance of about 18Ω and the reduced drive has a target impedance of approximately 40Ω. This feature is available on all DDR2 densities and configurations (x4, x8, and x16). Most point-to-point layouts will use the 40Ω value since it best matches the target trace impedance of approximately 40Ω to 55Ω.

Keep the traces short

When developing the PCB floor plan, the proximity of the DDR2 device to the memory controller is an important factor. If the memory is close to the controller, the layout is usually easier. For example, with short traces the address, control, and command signals may not require both parallel (RTT) and series (RS) termination, or at worst case, only require a small series resistor (RS) of about 10Ω or less. This RS is not for impedance matching, but is used to dampen the signals. If the RTT resistors are not required, this frees up more board space for signal routing and eliminates the need for a VTT power

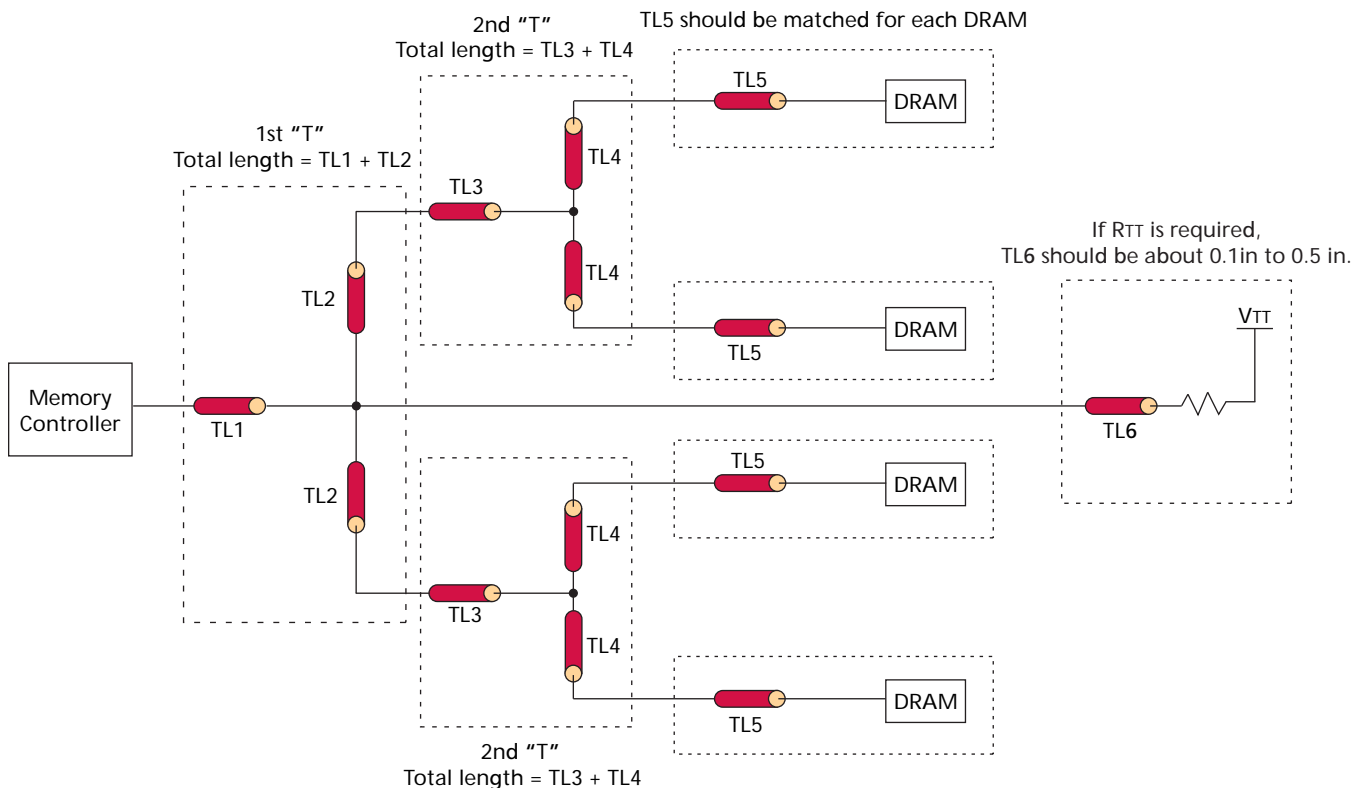
source and dedicated voltage plane. With longer traces it is likely that using a pull-up resistor, R_{TT} to V_{TT} ($V_{DDQ}/2$), will be required. To avoid the use of R_{TT} termination on high-speed DDR2 the target address trace length should be 2.5in (63.5mm) or less.

Use the Correct Topology

Not all point-to-point designs use only one memory device for each controller; many layouts will use two or four memory devices. If four x16 devices are used in parallel, this equates to a 64-bit data bus. In this scenario, the data and strobe signals are point-to-single-point, but the address, control, and command signals are point-to-four-point. When the controller is driving more than one device, a matched “TREE” type routing pattern should be used. See Figure 5 for an example of “TREE” routing.

The advantage of a “TREE” topology is that all trace segments are balanced for each path. The total trace length to the first DRAM is identical to that of the last DRAM, and flight times for each DRAM are also the same. Not only does this help with matching timing, it also helps with transmission line effects by controlling reflections. In “TREE” topologies, the first segment is typically the longest, and as the segments get closer to the DRAM, they get shorter.

Figure 5: Matched “TREE” Routing Pattern



Watch Out for Crosstalk

All DDR2 signaling is SSTL_18, which means the receivers are optimized to switch when they detect a very small input variation from the reference voltage (V_{REF}). The $V_{IN(DC)}$ levels are specified as $V_{REF} \pm 125\text{mV}$. If a fast switching data line is coupled with a nearby address or control line, the result could be system failure due to crosstalk. Crosstalk can occur even with short parallel traces.

DDR2 supports reduced drive I/O, which will have slightly slower edge rates. Slower edge rates will help to control crosstalk. It is also important to keep adequate spacing between signals that run parallel, and to provide a good solid reference plane so there is a current return path.

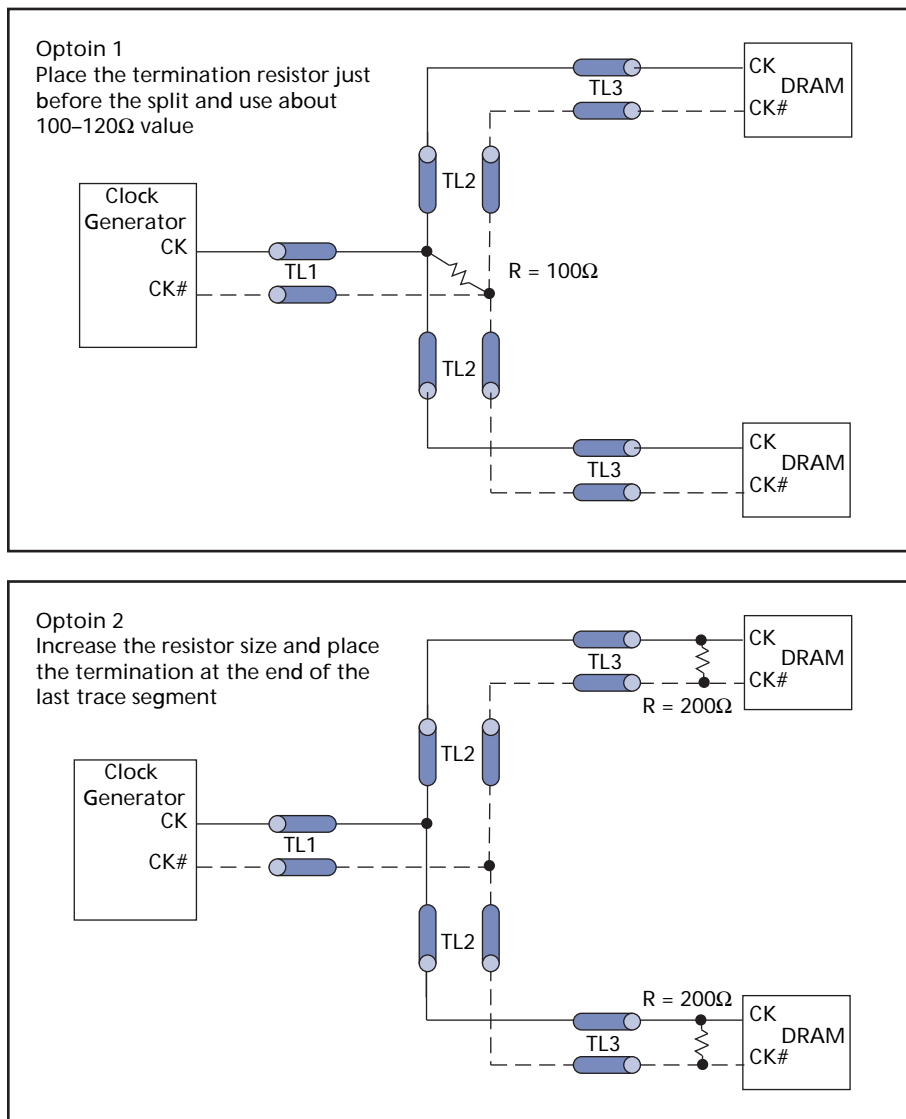
Methods of Termination

Three signal types are used on DDR2 memory: true differential signals (clocks), bidirectional signals (data), and single-ended signals (address/command/control). Each type of signal has unique termination requirements.

True Differential Signals (Clocks)

The clocks (CK and CK#) are a differential signal and should be terminated at the end of the line between CK and CK# with a 100–120 Ω resistor. If more than one DRAM is placed, then the clock pair may be split into several segments. In the case of multiple segments, the termination resistor should be placed at the first split, or the resistor value should be increased and a resistor placed at the end of each segment.

Figure 6: DRAM Clock Routing and Placement of the Termination Resistor(s)



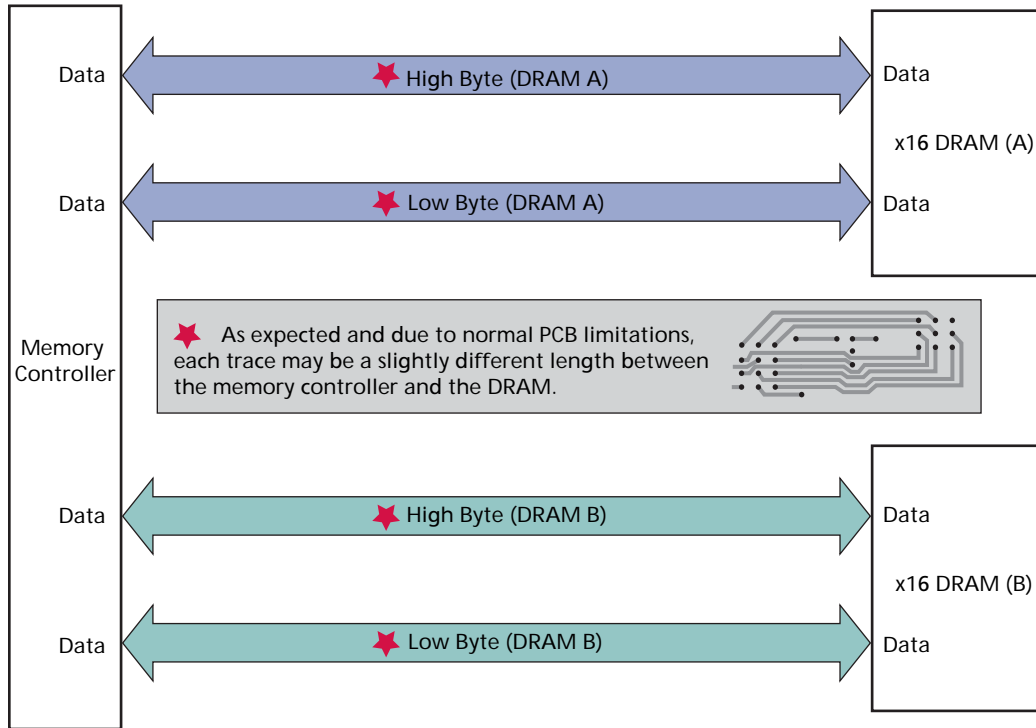
Bidirectional Signals (Data)

The I/O type signals include data (DQ0–DQ15), strobes (DQS, RDQS, LDQS, UDQS, DQS#, RDQS#, LDQS#, UDQS#), and data mask (DM, LDM, UDM). These signals all support ODT equivalent values of 50Ω, 75Ω, or 150Ω. With the use of ODT and a reasonable trace length, additional termination may not be needed. For these signals the critical factor is the timing within and between each byte lane. Depending, of course, on the timing budget for each design. A good starting point is keeping all traces within a byte lane between 15–20ps. This includes all data lines and the associated strobe. Timing between different byte lanes can usually be a bit more relaxed and are typically within 60–70ps.

Note: Values can vary between designs and will be dependant on actual controller requirements and timing budget.

Figure 7: Length Matching (Timing) Between Individual Bits and Byte Lanes – 32-Bit Bus Example

General guidelines (starting point):
Traces within each byte lane should not have a delta between each other of more than 15–20ps.
From byte lane to byte lane there may be up to 60–70ps of difference.



Notes: 1. The values for each design will vary due to controller requirements, device speed grades, and actual timing budget.

Single-Ended Signals (Address/Command)

Address, command, and control signals are single-ended and are driving only from the controller. Command (RAS#, CAS#, WE#) and control (CS#, ODT, CKE) signals are latched into the DRAM on each clock edge. Some controllers will run in 2T mode. In 2T mode, the address signals are driven for two clock cycles but only latched on the second rising clock edge. This allows the signals to settle out before they are latched into the DRAM. 2T addressing can also help to reduce the need for termination since the switching times are cut in half. Note that control signals should always run in 1T mode. As described earlier, the address, command, and control signals may not require any termination if the topology is optimized and the trace lengths are kept around 2.5in (63.5mm) or less. It is best to simulate all nets to determine the best method of termination and the exact termination values.

Conclusion

Using DDR2 SDRAM devices can help alleviate the complexity of the PCB floor plan with its ultra-small package size, common ballouts, and the need for fewer termination resistors. DDR2 technology supports two primary electrical ball arrays – one for x4/x8 and one for x16. The electrical ball array for the x4/x8 configurations is a standard 60-ball array, or 60-ball plus outriggers (68 balls total) for the longer package sizes. The x16 configuration uses the standard 84-ball array, or 84-ball array plus outriggers (92 balls

total) for the extended package sizes. Also, with the dynamic support of ODT on all I/O lines and reduced output drive levels, such PCB designs may not require additional termination to ensure good signal quality. The combination of compact ball arrays, miniature package sizes, and less termination placements all free up board space for routing traces.

Most PCB layouts using DDR2 SDRAM devices should have approximately 40–55 Ω trace impedances. Where signals are routed to more than one DRAM, balanced “TREE” trace segments should be used. If a parallel termination resistor is needed to VTT, it is best to place it at the end of the line or on a short segment of its own.

If the designer understands the basic principles of DDR2 technology, follows the recommendations of this technical note, and completes a full simulation, the resulting design should be a success.

For additional DDR2 information, and for the latest data sheets, refer to Micron’s Web site at www.micron.com



8000 S. Federal Way, P.O. Box 6, Boise, ID 83707-0006, Tel: 208-368-3900

prodmktg@micron.com www.micron.com Customer Comment Line: 800-932-4992

Micron, the M logo, and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners.



Revision History

Rev A.....	5/06
• Initial release	