

Technical Note

LP SDRAM Unterminated Point-to-Point System Design: Layout and Routing Tips

Introduction

Low-power (LP) SDRAM, including both low-power double data rate (LPDDR) and low-power single data rate (LPDDR), devices require a well-designed environment, package, and PCB, to maximize their performance for today's high-speed/low-power applications.

Designers can benefit from a set of proven layout and routing techniques for Mobile designs using unterminated point-to-point interfaces. Derived from electronics theory and Micron design experience, the guidelines in this technical note can enhance signal integrity (SI) optimization in unterminated point-to-point systems.

This technical note provides guidance for the development of multilayer board designs. It contains established system design recommendations to improve SI and reduce noise for LPDDR and LPDDR devices in unterminated point-to-point and point-to-multipoint systems.

The guidelines and examples in this technical note represent one of several acceptable methods and may not be applicable for all unterminated point-to-point designs.

Background

Previous Micron technical notes have focused on design, layout, and simulation techniques focused on standard SDRAM designs. Refer to TN-46-11, "Hardware Tips for Point-to-Point System Design: Termination, Layout, and Routing" and TN-46-14, "DDR SDRAM Point-to-Point Simulation Process."

The need for low power in LP SDRAM designs has influenced the adoption of programmable output drive strength and LVCMOS outputs. Though essential for LP designs, providing a programmable drive to replace external termination presents some unique design challenges. LVCMOS outputs eliminate the need for a V_{ref} supply but require near full-rail output transitions that can induce noise onto the DQ lines. The LP SDRAM interface is either x16 or x32, effectively doubling the number of DQ that transition during each cycle. The design of the memory signal environment can compensate for these differences to produce a more robust design.

Definitions

In this technical note:

- Vss refers to digital ground.
- Vssq refers to DQ and signal ground; the two are equivalent unless otherwise noted.
- Vdd is digital power for the device core.
- Vddq is DQ and I/O signal power; DQ and I/O power are also equivalent unless otherwise noted.
- “Unterminated” indicates that external series or parallel termination resistors are not utilized for signal termination. Instead, programmable drive strength is used to change signal impedance.
- Simultaneous switching outputs are referred to as SSO.
- Power delivery refers to power and ground layout, and decoupling techniques used to improve signal integrity.
- Memory signal environment refers to the layout-dependent characteristics that can affect signal integrity.

Programmable Drive Strength

Because LPDRAM devices are designed for mobile point-to-point applications, an option is provided to match drive strength to the impedance of the memory bus, eliminating the need for a termination voltage (V_{term}) and a series-termination resistor.

The programmable drive strength option supports four drive strengths: full (25Ω), three-quarter (37Ω), one-half (55Ω), and one-quarter (80Ω).

In low-power and low-cost applications, it is desirable to avoid using termination resistors, thus saving power and reducing costs. In point-to-point systems, termination resistors can be eliminated if the signal environment and driver are carefully selected.

Figure 1 on page 3 shows an example where the controller and LPDRAM device are positioned about 3 inches apart. The interconnect between the controller and the memory is designed so that its characteristic impedance is 50Ω .

For the first simulation, controller driver impedance is set to about 50Ω while the memory is configured for full drive strength (about 25Ω). Figure 2 on page 3 shows that the WRITE operation has good signal integrity with very little or no overshoot or ringback. Figure 3 on page 4 shows that during a READ operation, the strong driver and the unterminated line combine to produce large overshoot voltages and deep ringback. While the system may operate with such signaling, excessive noise reduces system margins and consumes more power than necessary during signal transitions.

This scenario can be improved by configuring the LPDRAM memory to use one-half drive strength. In such a case, the driver impedance closely matches the trace impedance and the signal is shown to be much more efficient. Figure 4 on page 4 shows that the reduced drive strength for the READ operation has eliminated the overshoot and ringback.

Figure 1: Data Point-to-Point




Figure 2: Memory WRITE with Full Drive Strength (25Ω)

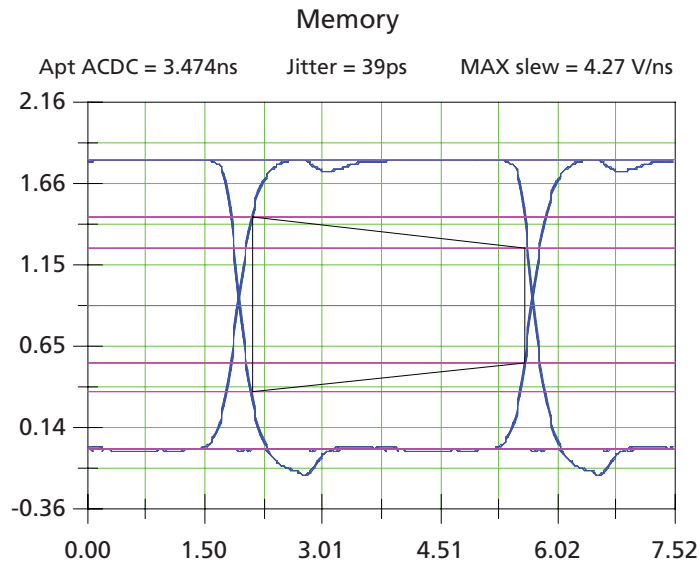


Figure 3: Controller READ with Full Drive Strength (25Ω)

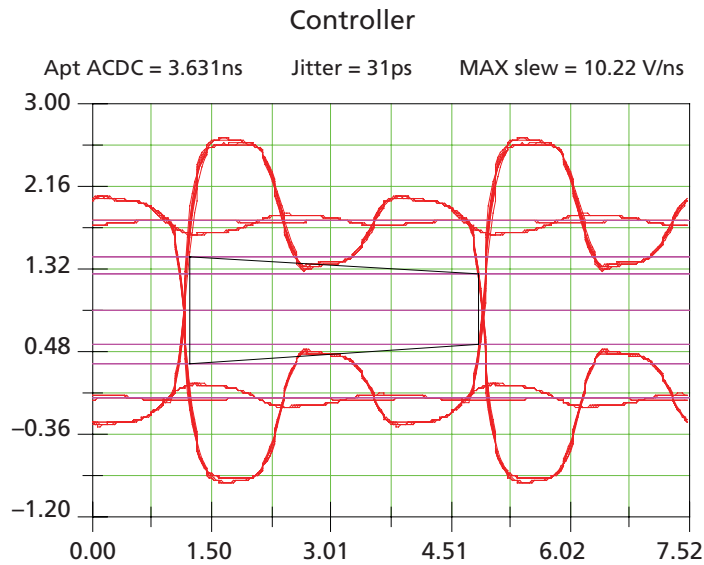
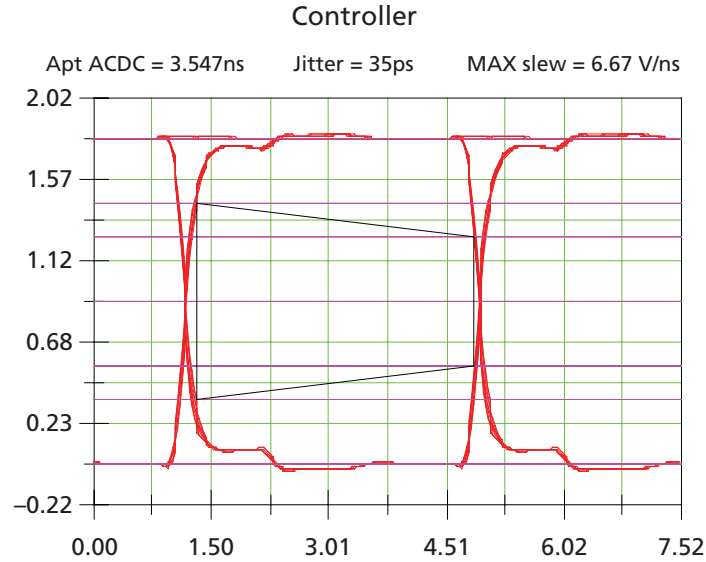


Figure 4: Controller READ with One-Half Drive Strength (55Ω)



Power Delivery

The LPDDR4 interface typically consists of either a x16 or a x32 data bus. Traditional SDRAM devices are usually x4 or x8.

LPDDR4 devices use LVCMOS output signals that switch to near full-rail levels for each transition. Traditional SDRAM uses SSTL-2 drivers with a reduced output-switching voltage range.

Doubling the number of I/Os and increasing the output-switching voltage range for LPDDR4 both contribute to SSO noise, particularly when full drive strength is selected for the interface. Increased switching noise can also require additional decoupling to filter the noise from the system.

The following examples illustrate how optimized power delivery can minimize SSO noise.

The waveforms in figures 5 and 6 are the result of a simulation done for a x32 LPDDR4 device on an unterminated mobile system.

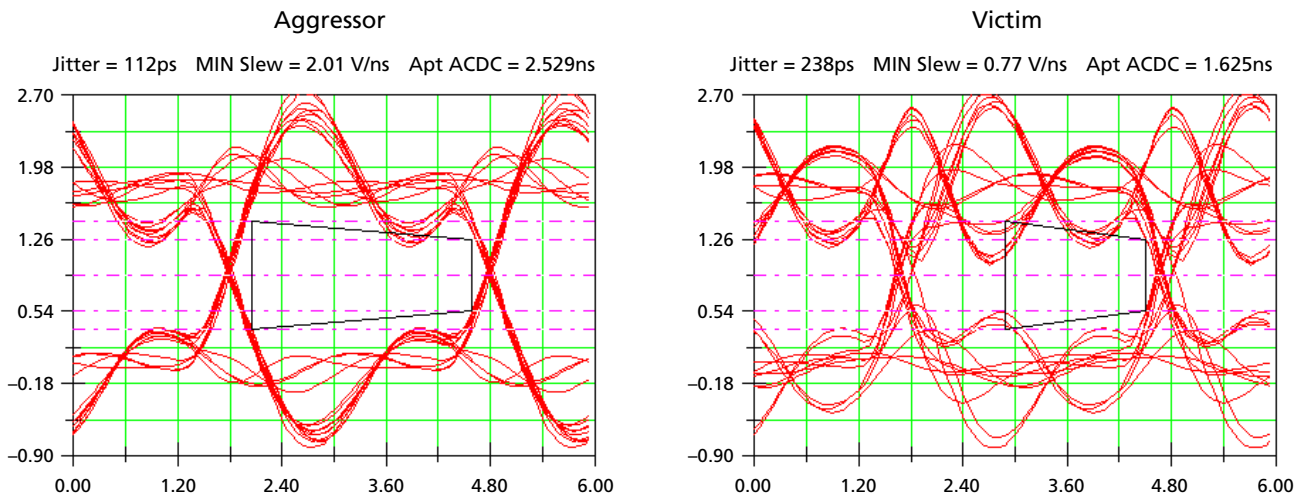
- The DQ bus is a 2-inch transmission line ($Z_0 = 50\Omega$) with a 5pF load.
- 31 DQ act as aggressors, all running the same PRBS pattern.
- The victim DQ runs a different PRB pattern.
- The four DQS pads are sent a recurring 0101... pattern.
- The data rate is 333 Mb/s (166 MHz).
- The driver is set to full drive strength.

The impact of poor power delivery (running at 333 MHz) on the DQ eye diagrams is shown in Figures 5.

Causes of poor power delivery include:

- Long, narrow metal traces for supply voltages
- Insufficient trace-routing spacing induces crosstalk
- Poorly placed or insufficient decoupling capacitors

Figure 5: DQ Eye with Full Drive Strength and Bad Power Delivery




The impact of good power delivery on the DQ eye when running at 333 MHz is demonstrated in Figure 6.

Techniques for good power delivery include:

- Minimizing inductance for supply voltages
- Spacing traces correctly to minimize crosstalk
- Providing dedicated supply voltage planes
- Using sufficient, well-placed decoupling capacitors

Achieving the best signal integrity in a system requires good designs for both trace routing and power delivery. A well-designed interconnect can be marginalized by a poorly designed power delivery mechanism. Maintaining low impedance for the power delivery provides the best signal integrity.

Figure 6: DQ Eye with Full Drive Strength and Good Power Delivery



Decoupling

Adequate power decoupling on the PCB is necessary to prevent excessive V_{dd} noise and the resulting memory errors in applications where power-supply draw can change by magnitudes in a single clock cycle.

Optimal capacitor placement and values can be defined through simulation or creation of a prototype. From this data, designers can determine the correct number of decoupling capacitors to use and their most advantageous placement. These factors—the number and location of capacitors—are of greater importance than the exact value of each capacitor.

As a general guideline, select the lowest equivalent series inductance (ESL) capacitor package suitable for your application. Use the highest capacitance available in that package and place as many capacitors as possible in the design.

Place capacitors as close as possible to each corner of each LPDDR or LPDDR device. Adjacent devices can share capacitors at the corners that fall between them. Backside components must have their own capacitors. This distributed decoupling approach minimizes capacitor ESL effects and localizes transient currents and returns.

See Micron technical notes TN-46-02, “Decoupling Capacitor Calculation for a DDR Memory Channel” and TN-00-06, “Bypass Capacitor Selection for High Speed Designs” for detailed decoupling discussions. Both are available on Micron’s Web site: www.micron.com.

Six-Layer PCB Stackup Examples

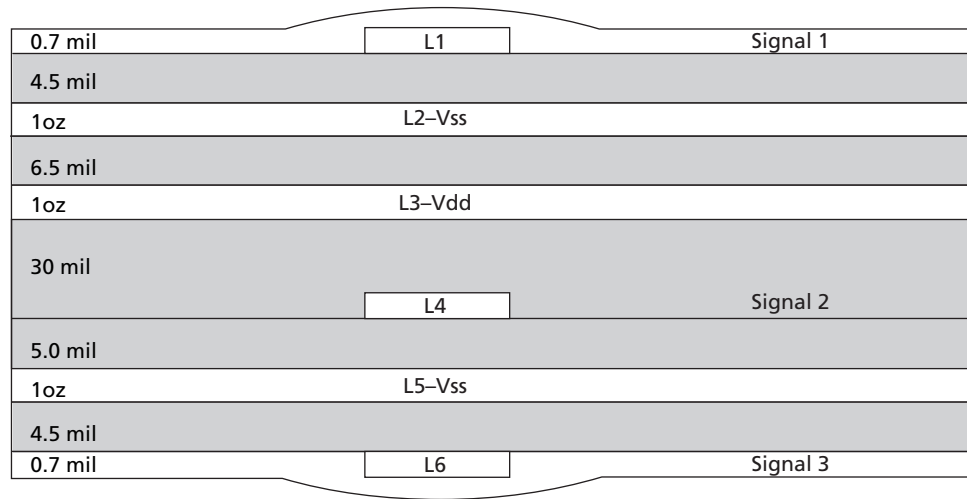
A well-designed PCB stackup is critical in eliminating digital switching noise. The ground plane must provide a low-impedance (Low-Z) return path for digital circuits.

Micron has experienced good results using a PCB design with a minimum of six layers: layers 1 (top) and 6 (bottom) for signals; layers 2, 3, and 5 for ground/power; and layer 4 for ground/power or for signals.

Poor SI and other problems render three- or four-layer PCBs unusable except in very limited cases.


A three-signal-layer stackup is shown in Figure 7, and a four-signal-layer stackup is shown in Figure 8. The required number of signal layers is determined by the number of signal groups to be routed and the required isolation between them. The number of devices to be routed and the size of the PCB dictate whether three or four signal layers are required. Simulation can provide feedback on signal integrity for a given application.

Figure 7: Example 1: Three-Signal-Layer PCB Stackup



- Notes:
1. Layer L1 has controlled impedance to L2 (Vss).
 2. Layer L4 has controlled impedance to L5 (Vss).
 3. Layer L6 has controlled impedance to L5 (Vss).
 4. Impedance on L1, L4, and L6 is 50–60Ω for a trace 5 mil wide.
 5. To balance copper weight in signal areas and reduce PCB warping, copper-flood the L4 layer where there are no signals.
 6. Solder mask thickness varies, but averages 0.7 mil.

Figure 8: Example 2: Four-Signal-Layer PCB Stackup



Packages

Typical routing for LPDDR4 components requires two internal signal layers, two surface signal layers, and two other layers (Vdd and Vss) as solid reference planes. Memory devices have Vdd and Vddq pins, which are both normally tied to the PCB Vdd plane. Likewise, component Vss and Vssq pins are tied to the PCB Vss plane. Each plane provides a low-impedance path to the memory devices to deliver Vssq. Sharing a single plane for both power and ground does not provide strong signal referencing.

Most LPDDR4 designs reference DQ, strobe, and clock signals to Vss; address, command, and control signals are referenced to Vdd. DQ signals are generally routed on the outer layers of the PCB.

Standard characteristic impedance (Z_0) of 50–60 Ω is recommended for all traces. The 60 Ω level also provides a good match to the output impedance of the controller/FPGA driver. Designers are advised to specify Z_0 , enabling board manufacturers to adjust dielectric thickness and line width to achieve the specification.

Minimizing the length of the signal return path (loop area) reduces transient current noise and electromagnetic interference (EMI). Micron recommends placing Vss and Vdd layers adjacent to each other on PCB layers 2 and 3 of a six-layer stackup and allowing the return signal on the plane immediately adjacent to the trace (such as plane 2 and plane 5) for the smallest loop area.

PCB Dielectric Material

The dielectric constant of PCB materials for most memory applications is 3.6 to 4.5, varying slightly with frequency, temperature, material, and the resin-to-glass ratio. FR-4, a commonly used dielectric material, averages 4.2 with signaling at 100 MHz. FR-4 is a copper-clad laminate that offers exceptional dimensional stability, dielectric thickness control, and high-quality manufacturing repeatability.

Micron has seen good results using FR-4 for PCBs, and FR-4 offers low cost, low moisture absorption, and low electrical conductivity in module fabrication.

Layout: Trace Widths, Intragroup Spacing, and Intergroup Spacing

There are two types of trace spacings that play a role in system SI: intragroup spacing and intergroup spacing.

Intragroup spacing (S1) is the distance between two adjacent traces within a related set of signals having similar or equivalent functionality. The control signals group, clocks, address bus, data bus, and data/strobes are all signal sets. The data bus is sometimes broken down into data bytes (sets of eight signals) plus the associated strobe and mask signal.

Intergroup spacing (S2) is the distance between the two outermost signals of different signal sets. For example, if the control signal set is routed together and adjacent to the address signal set, intragroup spacing is the distance between the two individual signals from the control and address sets that are closest together.

The difference between S1, S2, and trace width (S3) using the control and address groups is shown in Figure 9.

In general, Micron recommends an average S2 of 12 mil, with an 8 mil minimum. Closer spacing than the recommended minimum for S1 or S2 can increase costs (specific guidelines are shown in Table 1 on page 11). If all signals are routed at exactly 8-mil spacing for their full length, crosstalk is likely to disrupt SI, but if spacing limits are not met for short segments, SI is not likely to suffer.

Crosstalk is a function of trace spacing and slew rate; for systems with slew rates < 1 V/ns, trace spacing can be closer. Slower systems generally have more timing budget, which accommodates more crosstalk without affecting SI.


Trace Width (S3) Design Guidelines

Recommended S3 for functional signal sets:

- DQ lines = 4 mil minimum, 6 mil nominal
- DQS lines = 4 mil minimum, 6 mil nominal
- Address lines = 4 mil minimum, 6 mil nominal
- Command/control lines = 4 mil minimum, 6 mil nominal
- Clock lines = 4 mil minimum, 6–10 mil nominal

Supply voltages Vdd, Vddq, Vss, and Vssq must be composed of planes only, not traces. Short connections (≈ 8 mil) are commonly used to attach vias to planes in Micron designs. Any connections required from supply voltages to vias for device pins or decoupling capacitors should be as short and as wide as possible to minimize trace impedance. Micron recommends a 20 mil trace width.

Figure 9: S1, S2, and S3 Spacing



- Notes:
1. S3 is the trace width.
 2. Only RAS#, CAS#, and WE# control traces are shown. Other control signals are CS# and CKE.

Table 1: Intragroup and Intergroup Spacing Design Guidelines

Signal Set	Signals	Spacing Type	Min	Nom	Max	Unit	Notes
Data/Data strobe	DQ to DQ	S1	8	12	-	mil	
	DQ to DQS	S2	8	12	-	mil	
	DQS in a byte lane to DQS in a different byte lane	S1	-	-	-	mil	1
	DQ and DM	S2	8	12	-	mil	
Address	Adjacent address lines	S1	6	12	-	mil	
	Address lines	S2	6	12	-	mil	
Command/Control	CAS#, RAS#, WE#, CS#, CKE	S1	6	15	-	mil	
Clock	CK#-to-CK	S1	-	-	-	mil	2
	CK# (or CK in group of two) to DQS line	S2	-	-	-	mil	3
	Differential pair (CK, CK#) to any other signal	S2	8	12	-	mil	

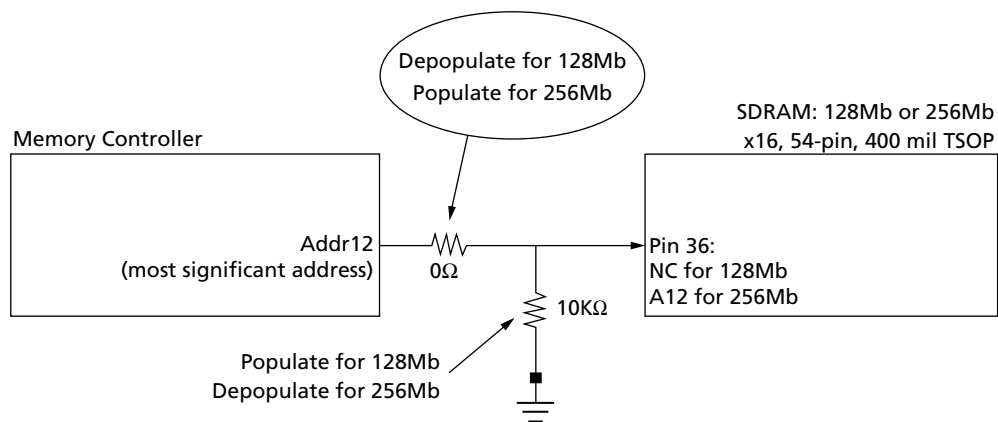
- Notes:
1. DQS signals are generally routed in the midst of related nibbles or bytes, so DQS-to-DQS spacing is not relevant.
 2. All CK and CK# signals lines should have differential characteristic impedance (Z_{diff}) of 100–120Ω.
 3. Generally not an issue as the CK# and DQS lines are not adjacent.

Expandability

Refer to Micron component data sheets for NC, DNU, or RFU pins in lower-density components that are used as upper-address pins in higher-density components. For example, ball H3 is NC in the 128Mb (8 Meg x 16) LPDDR4 component; it is A12 in the 256Mb (16 Meg x 16) LPDDR4 component. Designers can increase design expandability by placing a 0Ω series resistor between the pin and the controller such that it can be depopulated when switching to the lower-density component (shown in Figure 10 on page 11). This approach makes it possible to reuse the PCB without significant redesign.

Figure 10: Single PCB Layout to Accommodate Different Densities of Pin-Compatible SDRAM

Accommodating a 128Mb or 256Mb SDRAM component
(pinout difference is A12 only)



Some upper-column address lines change from “Don’t Care” for lower-density components to the most-significant column address inputs for higher-density components. If the controller does not account for these extra address pins, the system will experience addressing issues.

Routing

Though there are many signals on LPDDR or LPDDR components, most of them have similar functionality and work together. Groups of I/O signals have one of four purposes: to carry a binary address, to transmit or receive data, to relay a command to the device, or to latch in address/data or a command.

The address group consists of row/column address and bank address pins. The command group includes the row address strobe (RAS#), the column address strobe (CAS#), and write enable (WE#). The control group includes chip select (CS#) and clock enable (CKE). Each data group/lane contains 10 signals: the eight DQ (DQ[7:0]), the strobe (DQS), and the data mask (DM). Devices with x8 bus widths have only one data group, while x16 and x32 bus-width devices have two and four lanes, respectively.

To facilitate data lane routing, Micron recommends placing alternate, adjacent 10-line data lanes on different critical PCB stackup layers. To decrease crosstalk, when data and address/control tracks coexist on the same layer, they must be isolated from each other by at least 20 mil. If possible, place data tracks on different layers from address and control lanes.

Related functionality makes minimizing skew critical. This requires the signals of each group to be routed to similar electrical lengths. Routing address lines together on the same layer and isolating data lanes from the address, command, and control groups will also help minimize skew.

Match trace lengths for the data group within ± 50 mil of each other to diminish skew; serpentine traces (back and forth traces in an “S” pattern to increase trace length) can be used to match lengths. In addition, some controllers require byte lanes to have matched trace lengths. If the controller does not specify this, route byte lanes so that ± 500 mil is the largest trace-length difference relative to the clock group trace length.

Serpentine trace patterns contribute the desired delay, but be aware that there is some self coupling; that can change the propagation delay for a signal. Use simulations to validate timing.

Point-to-Point Clock Signals

LPDDR Devices

All LPDDR differential clock pairs (CK and CK#) must be routed on the same layer. Placing the clock signals on an internal layer minimizes EMI noise.

For LPDDR systems, match CK trace length to CK# trace length ± 20 mil, and CK/CK# trace lengths to DQS trace length ± 500 mil. If multiple clock pairs are transmitted from the controller to components, all clock-pair traces should be equivalent within ± 20 mil.

LPSDR Devices

LPDDR systems have only a single-ended clock (CLK), so the important trace-matching relationship is to the other groups rather than to a second differential clock trace. Match clock traces to data group traces within ± 500 mil. If multiple clocks are transmitted from the controller to components, all clock-pair traces should be equivalent to within ± 20 mil. Matching trace lengths to this level of accuracy helps minimize skew.

For both LPDDR and LPDDR, match clock traces to each signal trace in the address and command groups to within ± 400 mil. If clock traces cannot be matched to the trace lengths of these groups within 400 mil, then all clock trace lengths must be increased as a group. The longest-to-shortest trace-length difference must be ≤ 800 mil, so both longest and shortest traces determine how much length must be added to all clock lines.

Point-to-Multipoint Systems

Clock Routing in Multiple LPDDR Devices

For LPDDR systems, match CK trace length to CK# trace length ± 20 mil and CK/CK# trace lengths to DQS trace length ± 500 mil. If multiple clock pairs are transmitted from the controller to components, all clock-pair traces should be equivalent within ± 20 mil.

Recommended LPDDR routing topology for clock pairs is shown in Figures 11 and 12.

- If the trace lengths from split point to LPDDR components are less than ≈ 1 in (2.5cm), use a single 100–120 Ω resistor (R_t) at the split point (see Figure 11).
- If the lengths from the split point to the DRAM devices are greater than ≈ 1 in (2.5cm), use two resistors located near the respective LPDDR components (see Figure 12). These resistors are in parallel, so each R_t should be 200–240 Ω to keep the effective resistance at 100–120 Ω .

Figure 11: Single CK–CK# Differential Resistor Placement at Split Point

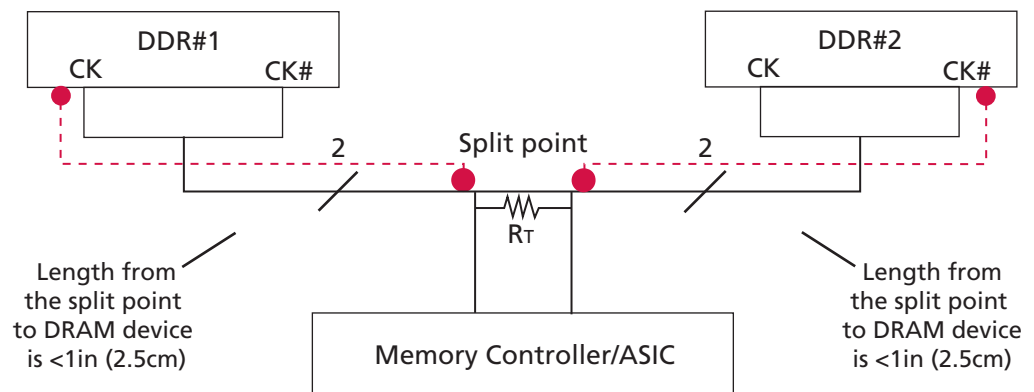



Figure 12: Dual CK-CK# Differential Resistor Placement at Component



Notes: 1. All CK and CK# signal lines have characteristic impedance (Z_0) of 100–120 Ω .

Match clock-pair traces to each signal trace in the address and command groups to within ± 400 mil. If clock traces cannot be matched to the trace lengths of these groups within 400 mil, then all clock trace lengths must be increased as a group. The longest-to-shortest trace length difference must be ≤ 800 mil, so both longest and shortest traces determine how much length must be added to all clock lines.

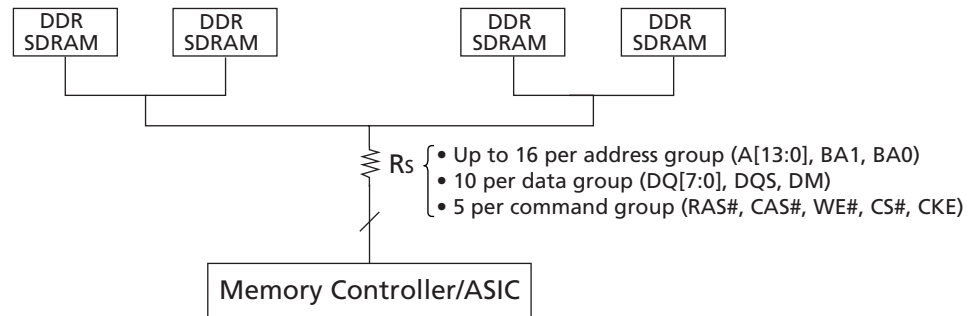
Control, Address, and Data Routing to Multiple LPSDRAM Components

If an application requires additional memory, multiple memory devices can be used. The additional load and routing will affect signal integrity. With standard LPDDR applications, the termination scheme would likely change to end termination. With LPSDRAM, other techniques may be feasible without using end termination.

When routing the interconnect between the controller and the memory devices, the use of a balanced-T topology is recommended (see Figure 13 on page 15). This will maintain an equal flight time for the signals going to each of the memory devices. The drive strength selection should still be chosen to match the trace impedance; however, signal integrity and timing margins could improve if the drive strength is set to a lower impedance.

It is important to simulate any multipoint design. When defining the source impedance, the selections available may not be ideal for your application. In such cases it may be necessary to add a series termination resistor to achieve the exact source impedance required. For unidirectional signals such as address and control lines, the series resistor should be placed as close to the driver as possible. With bidirectional signals, the series resistor can be placed at the midpoint of the trace, providing some benefit for all drivers on the network.

In a multipoint system design, address and control lines can take advantage of the fact that the bus is unidirectional. The equivalence simplification of circuits applies to transmission lines, so at the point of the T where the trace splits, the equivalent impedance of the parallel combined trace is one half that of the original impedance. This impedance discontinuity will cause some disturbance on the signals. To reduce this effect, the impedance of the trace from the split to the memory can be increased so that the parallel combination of the traces is closer to the original impedance of the trace. This reduces noise on the signal and improves signal integrity.

Figure 13: PCB Layout (Balanced-T) for Multiple Components


Miscellaneous Routing Recommendations

A 400 mil difference in address-, command-, or signal-group trace lengths equates to $0.4\text{in} \times (1,000\text{ps of propagation delay per } 6\text{in of trace})$, or a skew of $\approx 67\text{ps}$. If the timing budget can absorb this minor amount of lane-to-lane skew and other routing delays, the system will perform normally. Total routing-based delays must meet ${}^t\text{DQSCK}$, controller DQS recovery limits, and other data sheet AC timing parameters.

Regardless of bus type, all signal groups must be properly referenced to a solid Vssq or Vddq plane. For both READs and WRITEs, the key relationship is between CK/CK#, DQ, DM, and DQS signals (the LPDDR data group), which operates at twice the speed of other signal groups and makes SI more critical. DQ, DQS, and clock lines are best referenced to Vssq to minimize noise. If a Vssq layer is not easily accessible, address and command lines can reference a Vddq layer.

Keep traces as short as possible. If trace length (from controller pad to LPSDRAM pad) is $<2\text{in}$ (5cm) for both LPDDR and LPSDR point-to-point applications, routing is simpler and signal quality usually increases in proportion. In most cases, trace lengths $>2\text{in}$ (5cm) lead to more signal undershoot, overshoot, and ringing—all of which are detrimental to SI.

Additional Trace-Length Design Guidelines

- Match different DQ byte lanes to within 1in (2.5cm) of each other. A 1in trace-length difference equates to approximately 167ps of propagation delay. Thus, the timing budget must be able to absorb 167ps for a 1in difference in byte-lane matching.
 - Within a byte lane, match all DQ and DQS traces to within ± 50 mil.
 - Route data groups next to a Vss plane to minimize the return path/loop length.
- Maintain a solid ground reference (no split planes, etc.) for each group to provide a low impedance return path; high-speed signals must not cross a plane split.

Simulation

During the layout phase for a new or revised design, Micron strongly recommends simulating I/O performance at regular intervals. Optimizing an interface through simulation can help decrease noise and increase timing margins before building prototypes. Issues are often resolved more easily when found in simulation, as opposed to those found later that require expensive and time-consuming board redesigns or factory recalls.

Micron has created many types of simulation models to match the different tools in use. Component simulation models presently posted on Micron's Web site include IBIS, Verilog, VHDL, Hspice, Denali, and Synopsys.

Verifying all simulated conditions is impractical, but there are a few key areas to focus on: DC levels, signal slew rates, undershoot, overshoot, ringing, and waveform shape. Also, verifying that the design has sufficient signal-eye openings to meet both timing and AC input voltage levels is extremely important.

Summary

Signal integrity, power delivery, and decoupling are all major concerns when designing LPSDRAM applications.

Mobile LPSDRAM designs (LPDDR and LPSDR) provide an attractive alternative to traditional DRAM designs when used for mobile applications. The option to control the drive strength to match the impedance of the memory bus enables removal of the termination voltage (V_{term}) and series termination resistors. Mobile LPSDRAM designs can be used to reduce memory cost and power consumption in mobile applications.

Mobile applications, when properly designed and validated through simulations, can realize superior functionality and stability.

8000 S. Federal Way, P.O. Box 6, Boise, ID 83707-0006, Tel: 208-368-3900
www.micron.com/productsupport Customer Comment Line: 800-932-4992
Micron and the Micron logo are trademarks of Micron Technology, Inc.
All other trademarks are the property of their respective owners.



Revision History

Rev A.....	11/08
• Initial release.	