



WP484 (v1.0) September 27, 2016

DDR2/DDR3 Low-Cost PCB Design Guidelines for Artix-7 and Spartan-7 FPGAs

By: Ravindra Gali

The Artix®-7 and Spartan®-7 families offer a low cost, small footprint array of highly efficient FPGAs, purpose-designed to address the special needs of the low-end market.

ABSTRACT

In an ongoing endeavor to increase throughput, designers have increasingly been pairing low-power, low-cost FPGAs like the Xilinx® Artix-7 and Spartan-7 class of devices with high-performance DDR2/DDR3 memories. In today's cost-sensitive systems, for example, one might find a low-cost FPGA moving data to and from a DDR3 memory at up to 1066Mb/s.

Given the cost-sensitive nature of these applications, system designers are often challenged to come up with a PCB design with the lowest bill of materials (BOM) cost. This white paper provides PCB designers with a set of pragmatic layout guidelines to tackle high-performance DDR2/DDR3 designs based on low-cost FPGAs. Also addressed are the cost trade-offs for designers opting for advanced PCB fabrication technologies to reduce the PCB layer count.

Introduction

Artix-7 and Spartan-7 devices come in a wide variety of packages that are designed for maximum performance and maximum flexibility. The Spartan-7 FPGA packages are available in small package footprint with package sizes ranging from 8mm to 27mm, while the Artix-7 FPGA packages vary from 10mm to 35mm. The packages are available in a 1.0mm, 0.8mm, and 0.5mm package pitches, respectively. Package pitch is defined as the distance between consecutive balls on a BGA package, measured from center to center, as shown in [Figure 1](#).

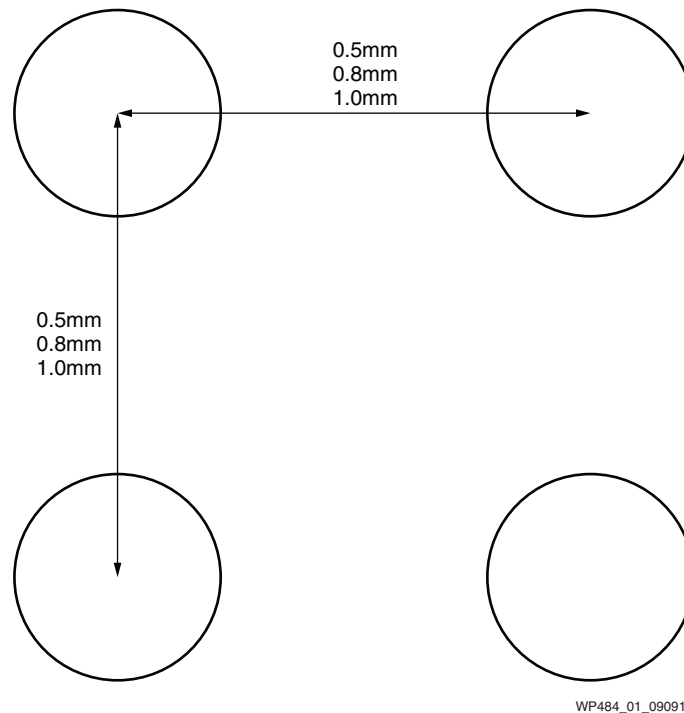


Figure 1: Package Pitch

In general, as the pitch size decreases, the challenges for PCB routing increase because there is less room to route traces and vias between the package balls.

Layer Count Estimation and Cost Trade-Offs

A quick way to estimate the number of routing layers required to fully break out signal pins from the FPGA is to use [Equation 1](#):

$$\text{Layers} = \frac{\text{Signal Pins (I/Os, MGTs)}}{\text{Routing Channels} \bullet \text{Routes per Channel}} \quad \text{Equation 1}$$

For Xilinx® cost-optimized FPGAs, the number of signal pins is approximately 60% of the total number of BGA balls. The other 40% includes power and ground signals that are most often routed directly down to the planes by vias. This is assuming full I/O utilization. If fewer I/Os are used, then the number of signals to route decreases accordingly.

Routing Channels are the total number of available routing paths out of the BGA—e.g., (Number of BGA balls on one side – 1) × four sides.

Figure 2 shows a sample 5 × 5 BGA ball out, resulting in a total of sixteen total routing channels.

Number of BGA balls on one side = 5
 Routing Channels = (5 – 1) × 4

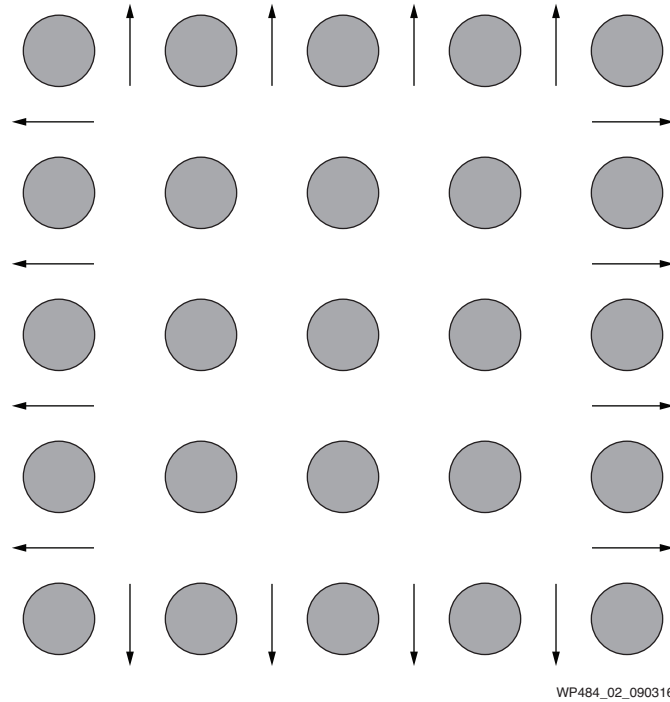
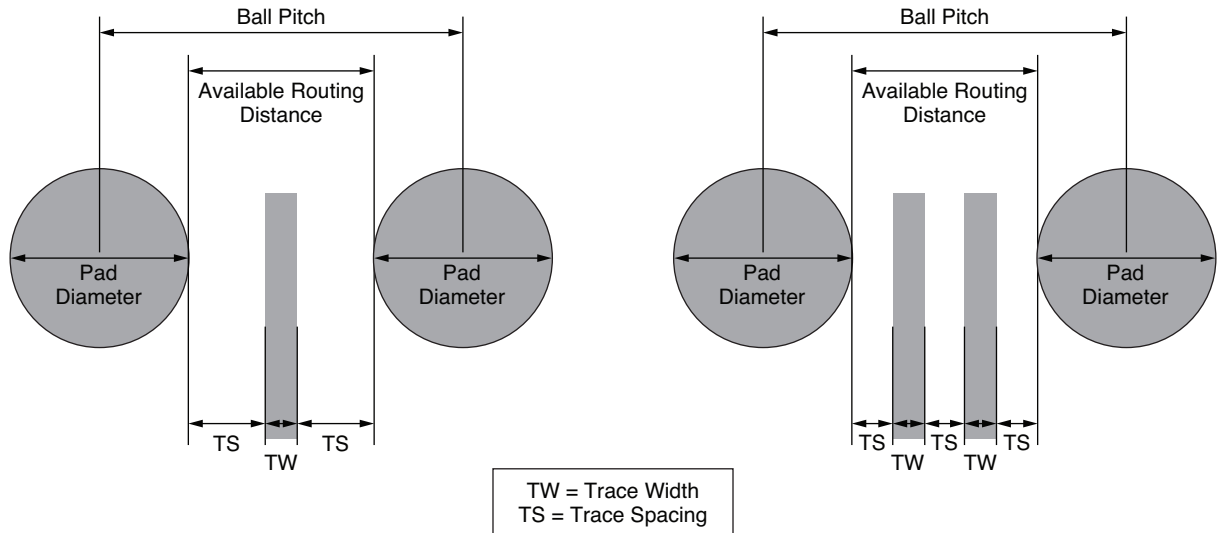


Figure 2: Definition of Routing Channel

Routes per Channel are typically one or two, depending on the number of traces that can be routed between the BGA pads on the top/bottom layer. From a signal integrity standpoint, adhering to design for manufacturing (DFM) guidelines is critical to assure meeting the nominal trace impedance requirements. Figure 3 shows a schematic representation of routes per channel. The routes per channel on the inner layers depend on the spacing between the vias, taking into account the drill-to-copper specifications.



WP484_03_090316

Figure 3: Definition of Routes per Channel on Top/Bottom Layers

For a 0.5mm pitch package, the pad size and the package pitch dimensions limit the PCB designer to a single trace between the BGA pads. However, PCB designers have the flexibility to go either for a single-trace or a dual-trace breakout when opting for a 0.8mm or a 1mm pitch package.

The approximate number of layers required to route Artix-7 and Spartan-7 FPGAs are shown in Table 1 and Table 2, respectively.

Table 1: Artix-7 FPGAs: Approximate Signal Layers per # of BGA Balls

BGA Balls	Ball Pitch (mm)	Signal Layer Counts (All Available I/Os Routed)	
		Routes per Channel:	
		Two Traces	One Trace
236	0.5	NA	3
256	1.0	2	3
324	0.8	2	3
325	0.8	2	3
484	0.8	3	4
484	1.0	2	4
676	1.0	3	5
1156	1.0	3	6

Table 2: Spartan-7: Approximate Signal Layers per # of BGA Balls

BGA Balls	Ball Pitch (mm)	Signal Layer Counts (All Available I/Os Routed)	
		Routes per Channel:	
		Two Traces	One Trace
144	0.5	N/A	2
196	0.5	N/A	2
225	0.8	2	3
324	0.8	2	3
484	1.0	2	4
676	1.0	3	5

The critical factors to consider when breaking out signals underneath a high-density BGA include:

- Dimensions of surface land pads
- PTH size and the corresponding pad/anti-pad dimensions
- Trace width and spacing requirements
- Number of signal layers available

The number of options available for a PCB designer during layout is primarily driven by the package pitch. PCB designers with a goal to minimize the PCB layer count at the expense of cost can use advanced fabrication techniques like micro vias, blind vias, and buried vias in addition to using thinner trace widths. However, these advanced fabrication techniques are not mandatory to ensure the success of a DDR3 design. Following [Figure 4](#) is a brief description of the various industry terms, along with an approximation of cost adders applied to the standard PCB fabrication cost. [Figure 4](#) shows the various via types.

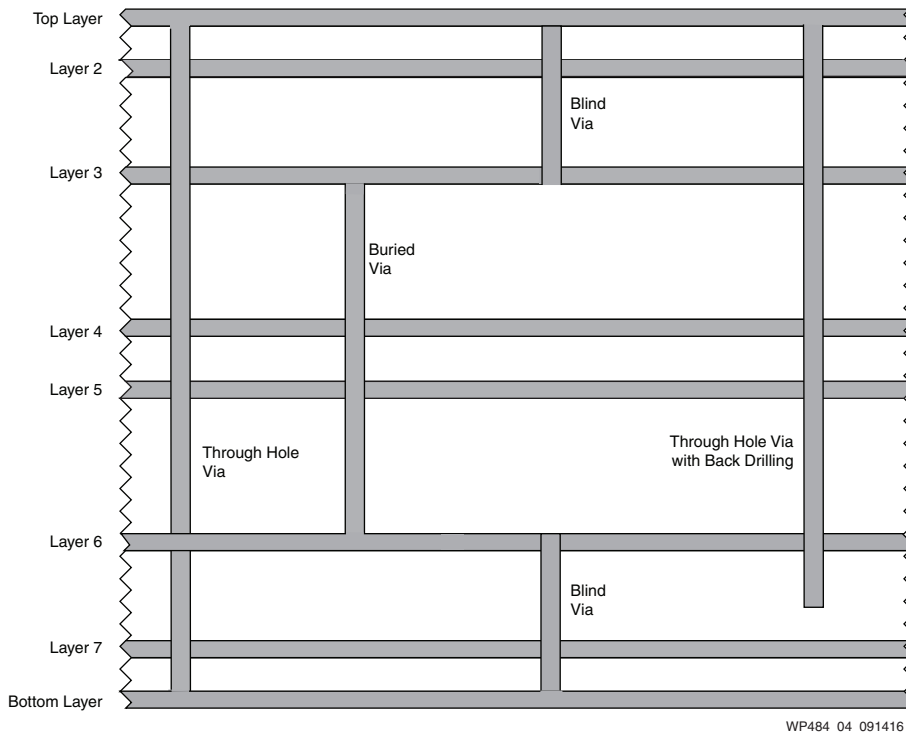


Figure 4: Different Via Types

Via Aspect Ratio – The ratio of PCB thickness to the smallest unplated via drill hole diameter. This is used as a guide to ensure that the PCB fabricator does not exceed the mechanical capabilities of the drilling equipment. A via aspect ratio of 10:1 is fairly common with the standard PCB fabrication. The via aspect ratio can be increased to 20:1 using advanced PCB fabrication while maintaining the design for manufacturing (DFM) rules.

Back-drilled Vias – A back-drilled via is a through-hole via that has a portion of its length “drilled out” such that it is no longer conductive. This improves signal integrity because it removes an unneeded stub from the route. The typical cost adder for back-drilling vias ranges from 5–10% of the total PCB fabrication cost.

Via-in-Pad – A Via-in-Pad is a via drilled directly beneath a pad. This removes the need for a separate metal trace (stringer) to drop down a via. This can help with breakout routing and improved signal integrity at the expense of higher board fabrication cost. The cost adder varies from +10–15% of the PCB fabrication cost and is dependent on the via aspect ratio.

Buried and Blind Vias – A buried via is located entirely inside the printed circuit board and does not touch the top or bottom layers; a blind via travels from either the top or bottom layer to an inner signal layer. Both types of vias free up room above or below for other routing. This differs from a through-hole via, which travels all the way from the top to the bottom layer. The cost adder for a buried or blind via depends on the number of different types of buried or blind vias that exist on the PCB. Each type of buried / blind via requires a separate lamination cycle, resulting in extra cost. For example, a PCB with three different types of buried / blind vias (L1 – L4, L16 – L12, L4 – L8) on a 16 layer PCB results in a 30% cost adder per each type of buried / blind via.

Micro Vias – A micro via is a form of blind via. The dimensions of a micro via are very small. They are formed using lasers and typically cannot penetrate more than one or two layers at a time. The cost adder is approximately 15% for each type of via.

Extra Layers – The cost of adding extra signal layers might be lower than the cost for some of the advanced via technologies described above. Hence, adding layers should not always be considered a negative alternative. The cost adder for two additional layers is typically 15–20% of the PCB fabrication cost.

With an advanced fabrication process, the PCB designer can spec traces as narrow as 2.5mils with 2.5mils spacing to achieve the target impedance specifications for optimum SI performance.

Xilinx has a detailed data sheet titled *Recommended Design Rules and Strategies for BGA Devices* [Ref 1] that provides specific guidelines on the dimensions of the surface pads, PTH sizes, trace width, and spacing recommendations for PCB breakout routing for various pin pitches. The user guide also includes snapshots of actual PCB layouts for a cost-optimized design, along with a design using an advanced PCB fabrication process for different pin pitches.

Low-Cost DDR3 Guidelines

Based on system requirements, DDR2/3 memories are connected to the Artix-7 and Spartan-7 FPGAs as either a set of discrete SDRAMs or as a DIMM module. Not all devices in these product families support all possible memory configurations. The exact memory configuration supported is dependent on the specific die / package combination.

Regardless of the topology, successful operation of the DDR2/3 interface at the highest possible data rate depends on its own microsystem of components and other factors. These factors include driver and receiver buffers, terminations, interconnect impedances, delay matching, crosstalk, and power integrity. A general comparison of the two memory types is shown in Table 3, while the signals common to both DDR2 and DDR3 are shown in Figure 5.

Table 3: Comparative Requirements of DDR2 and DDR3 Memory

Technology	DDR2	DDR3
Maximum supported clock frequency (MHz) / Data rate (Mb/s) on Artix-7 and Spartan-7	400/800	533 / 1066 on Artix-7 400/800 on Spartan-7
Power Requirements (Volts)		
V _{VDDQ}	1.80	1.50
V _{TT}	0.9	0.75
V _{REF}	0.9	0.75
Delay Matching Requirements		
Match ADDR/CMD/CTRL to Clock	Yes	Yes
Match Data group (DQ), DM to corresponding strobe pair (DQS)	Yes	Yes
Match DQS to Clock Loosely	Yes	Not required

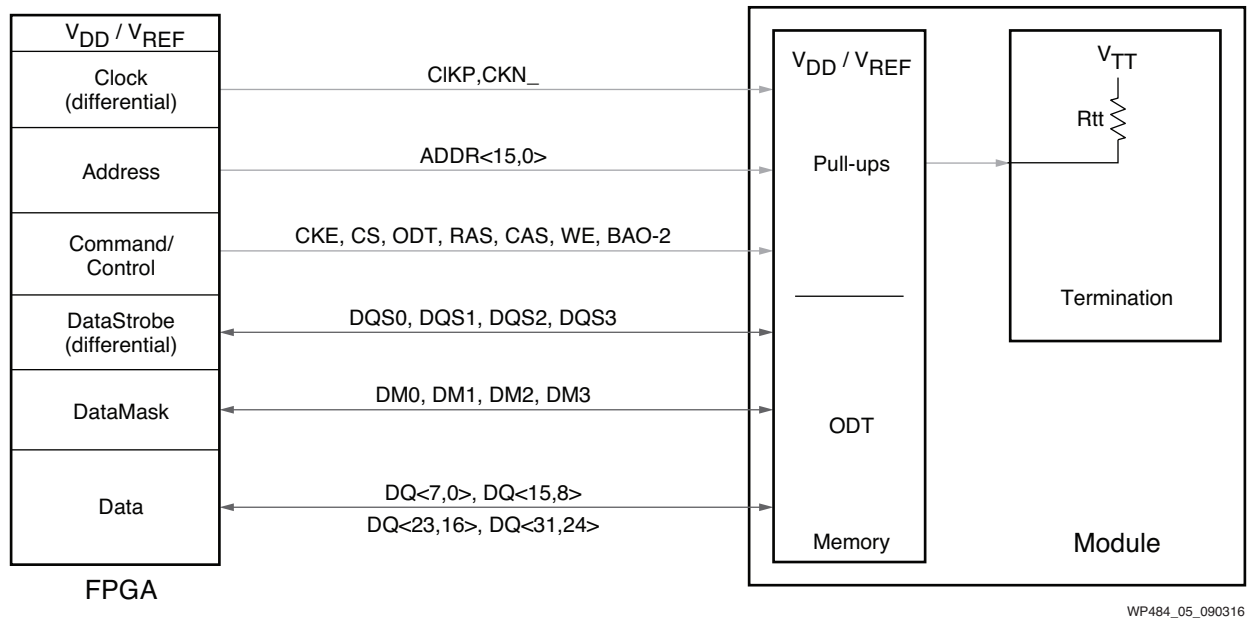


Figure 5: Architecture and Interface Technology Common to DDR2 and DDR3 Memory

This section provides high-level layout guidelines for enabling a low cost PCB design. The key challenge to a successful memory layout is:

- Breaking out all the data and address signals on minimum number of routing layers
- Ensuring a robust signal integrity by minimizing crosstalk, signal reflections due to impedance discontinuities, etc.

Waveform Integrity

DQ, DM, DQS nets are typically point-to-point connections. These nets are bidirectional, with data being latched on both the rising and falling edges of their associated data strobe signals. Xilinx recommends the following:

- Choose a FPGA driver setting with an output impedance closest to the transmission line impedance.
- Route these signals with a 50Ω characteristic impedance on the PCB all the way from the FPGA to the memory device.
- Enable the on-die termination (ODT) setting that is closest to 50Ω on the DRAM to minimize reflections during WRITE operation.
- Enable termination on FPGA during READ operation to ensure a matched termination for bidirectional high data rate operation

To minimize crosstalk, it is always advisable to space the signals far apart and minimize the via length during layer transitions. However, the area under the FPGA and DRAM device are space-constrained, making it difficult to space the signals far apart. To ease PCB layout, Xilinx allows a minimum spacing—i.e., 1X spacing—in the breakout region. 1X spacing refers to the air gap between the traces equal to the trace width. This spacing can be maintained provided the trace

length is less than 1in when breaking out of the FPGA/DRAM device. To further ensure reliable signal integrity, the following guidelines must also be followed:

- 2X or greater spacing after the breakout region.
- Total PCB interconnect length of 4in from the FPGA to the DRAM using an FR4-type substrate.

CLOCK, ADDR, CMD, and CONTROL Signals

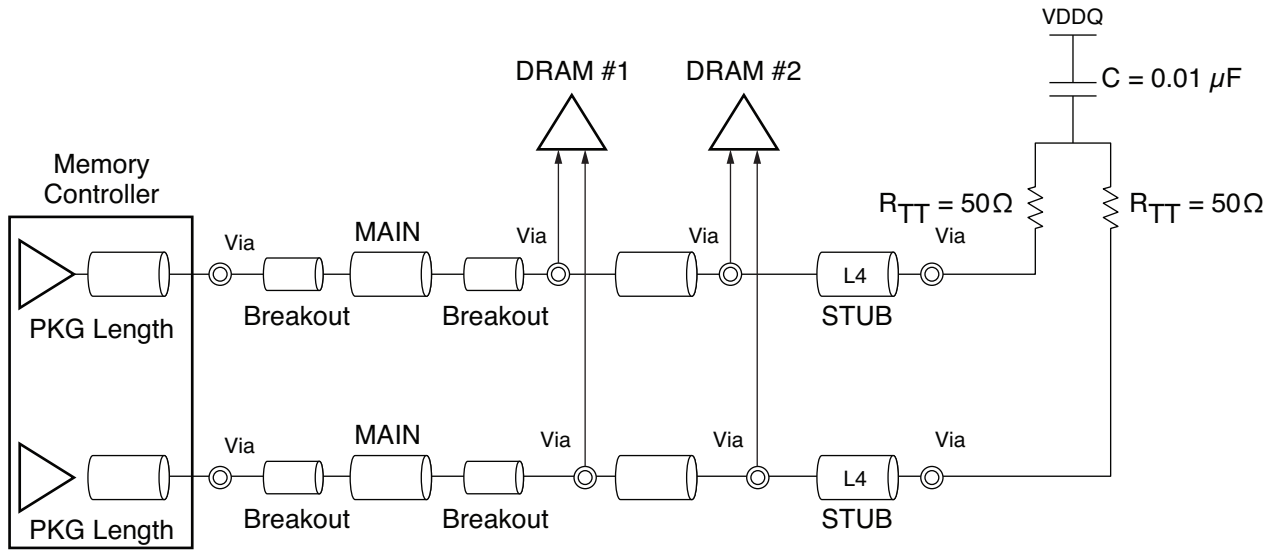
The CLOCK, ADDR (address), CMD (command), and CONTROL signals are typically point to multi-point connections and require a unique topology termed Fly-by. Fly-by can be envisioned as a daisy-chain connection without stubs. These signals are unidirectional and are driven from the FPGA to the memory device. The differential clock net is used as the reference signal for timing analysis.

External Termination

ODT is not available for these nets, and an external discrete termination is required. The recommended form typically consists of a resistor placed at the far end, past the last memory device, and pulled up to V_{TT} ($V_{VDDQ}/2$). The value of the pull-up resistor and the impedance of the interconnecting traces depends on the number of devices on the net. These values are usually optimized through simulation. Xilinx provides the necessary FPGA IBIS models for customers who want to perform their own simulation analysis to further optimize these values.

Xilinx recommends the following guidelines for low cost design.

- Xilinx recommends routing the ADDR/CMD/CTRL signals with a 50Ω characteristic impedance trace using a 50Ω fly-by termination resistor to V_{TT} ($V_{VDDQ}/2$) at the far end. This should be applicable for most cases.
- For the CLK differential pair, it is recommended to implement a differential trace impedance of 100Ω using two separate 50Ω pull-up resistors to V_{VDDQ} and a DC blocking capacitor. [Figure 6](#) shows a schematic representation for clock fly-by termination.



WP484_06_090316

Figure 6: Fly-By Termination

Figure 7 shows V_{TT} capacitor placement to ensure reliable power integrity. Xilinx recommends placing at least one $0.1\mu\text{F}$ capacitor tied to V_{TT} for every four termination resistors.



WP484_07_090316

Figure 7: V_{TT} Capacitor Placement

In addition, Xilinx recommends keeping the interconnect length between the DRAMs and the stub length from the last DRAM to fly-by termination resistor within 0.75in.

Delay Matching of Signal Nets

While trace length, impedance, and terminations can be designed for optimal waveform integrity, it is also important to ensure that the delay between the synchronous nets be matched very closely. All DQ and DM nets in a byte lane must be matched to their associated DQS nets, taking into account the package flight time differences.

Xilinx recommends the following:

- All DQ/DM nets should be matched to their associated DQS nets to within ± 15 ps for DDR2/DDR3 interfaces at 800Mb/s
- All DQ/DM nets should be matched to their associated DQS nets to within ± 10 ps for DDR2/DDR3 interfaces at 1,066Mb/s
- For unidirectional signals, all ADDR/CMD/CTRL signals must be matched to the CLK signal. It is a good design practice to match each transmission line segment (FPGA to DRAM1, FPGA to DRAM2, FPGA to fly-by termination resistor, etc.) to within a reasonable tolerance of ± 25 ps.

Conclusion

Xilinx Artix-7 and Spartan-7 devices are proven to interoperate with DDR2/3 speeds at up to 1,066Mb/s and 800Mb/s, respectively. The purpose of this white paper is to provide high-level guidance on layer count estimation and the cost implications of using advanced fabrication technologies. In addition, the paper presents high-level layout guidelines for low-cost PCB design, which help optimize I/O performance and reduce the risk of performance issues. For complete details, refer to the device data sheets.[\[Ref 2\]](#)[\[Ref 3\]](#)

References

1. Xilinx User Guide [UG1099](#), *Recommended Design Rules and Strategies for BGA Devices*
2. Xilinx Data Sheet [DS181](#), *Artix-7 FPGA Data Sheet: DC and AC Switching Characteristics*
3. Xilinx Data Sheet [DS189](#), *Spartan-7 FPGA Data Sheet: DC and AC Switching Characteristics*

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
09/27/2016	1.0	Initial Xilinx release.

Disclaimer

The information disclosed to you hereunder (the "Materials") is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available "AS IS" and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

Automotive Applications Disclaimer

AUTOMOTIVE PRODUCTS (IDENTIFIED AS "XA" IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE ("SAFETY APPLICATION") UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD ("SAFETY DESIGN"). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.