



WP420 (v1.0) June 26, 2012

Xilinx Virtex-6/Spartan-6 FPGA DDR3 Signal Integrity Analysis and PCB Layout Guidelines

By: Syed Bokhari (Fidus Systems, Inc.) and Romi Mayder (Xilinx, Inc.)

In the ongoing endeavor to increase throughput, designers have been increasingly pairing Xilinx® Virtex®-6 and Spartan®-6 FPGAs with the latest high-performance DDR2 and DDR3 memories, constantly pushing the envelope of the speeds at which such devices can operate. In today's high-end systems, for example, one might find a Virtex-6 FPGA (-2 or -3 speed grade) moving data to and from a DDR3 memory at up to 1,066 Mb/s. [\[Ref 1\]](#)

As the designer considers turning these dreams of ultra-high throughput into reality, however, a nightmarish design/debug cycle might seem inevitable. This white paper strives to provide designers with a set of pragmatic tools with which to tackle a high-performance design based on Virtex-6/Spartan-6 FPGAs.

Introduction

Based on system requirements, memories are connected to FPGAs as either a set of discrete SDRAMs or as a single DIMM module, as shown in [Figure 1](#) and [Figure 2](#).




Figure 1: **Implementation of DDR3 Memory as a Set of Four SDRAMs**





Figure 2: **Implementation of DDR3 Memory as a SODIMM Module**

Regardless of the topology, successful operation of this interface at the highest possible data rate depends on its own microsystem of components and other factors. These ultimately determine whether the waveform integrity and delay allow the interface to operate as intended. Components and factors that characterize the operation of an FPGA-based DDR2/3 system include driver and receiver buffers, terminations, interconnect impedances, delay matching, crosstalk, and power integrity — this final factor being often overlooked by designers.

A general comparison of the two types is given in [Table 1](#), while the signals common to both DDR2 and DDR3 memories are shown in [Figure 3](#).

Table 1: Comparative Requirements of DDR2 and DDR3 Memory

Technology	DDR2	DDR3
Max Clock Frequency (MHz)/Data Rate (Mb/s)	533/1,066	933/1,866
Power Requirement:		
V_{DD} (Volts)	1.8 ± 0.1	1.5 ± 0.075
V_{TT} (Volts)	0.9 ± 0.04	0.75 ± 0.03
V_{REF} (Volts)	0.9 ± 0.018	0.75 ± 0.015
Delay Matching Requirement:		
Match ADDR/CMD/CNTRL to Clock Tightly	Yes	Yes
Match DQ<7,0>, DM0 to DQS0 Tightly	Yes	Yes
Match DQ<15,8>, DM1 to DQS1 Tightly	Yes	Yes
Match DQ<23,16>, DM2 to DQS2 Tightly	Yes	Yes
Match DQ<31,24>, DM3 to DQS3 Tightly	Yes	Yes
Match DQS0-3 to Clock Loosely	Yes	Not Required



WP420_03_042512

Figure 3: Architecture and Interface Technology Common to DDR2 and DDR3 Memory

This document provides guidelines that are applicable to a large majority of designs based on signal integrity (SI) simulations that use IBIS models for Virtex-6 and Spartan-6 devices. Links to documents containing additional details can be found in the [References](#) section.

Waveform Integrity

DQ, DM, and DQS

DQ, DM, and, DQS nets are typically point-to-point connections, except in situations that involve multi-rank configurations. In such cases, the memory devices can be in the form of a stacked die, or two memories can be placed back-to-back on a PCB. In effect, however, such configurations emulate a point to-point connection. These nets are bidirectional, with data being latched on both the rising and falling edges of their associated data strobe signal. Consequently, for a 533 MHz data strobe signal, the data rate is 1,066 Mb/s. On-die termination (ODT) is invariably used at the memory device on a WRITE operation, and Digitally Controlled Impedance (DCI) is activated within the Xilinx FPGA during a READ operation to ensure a matched termination for bidirectional high data rate operation.

Data WRITE in Single-Ended DQ and DM Nets

In the WRITE case, the driver is within the FPGA and the receiver is within the SDRAM. The Virtex-6 FPGA provides an SSTL 1.5V I/O standard (IBIS model name: *Virtex6_SSTL15_DCI_O*). The SDRAM buffer must provide ODT. Typically, ODT values are selectable between 40Ω and 60Ω, resulting in a common interconnect impedance also in the range of 40Ω–60Ω. (The case of 60Ω impedance with ODT 60 termination shows slightly better noise margin than the 40Ω case.)

In a system involving only a few discrete SDRAMs, the interconnect trace length is usually kept in the range of 500 mils–2,000 mils (1 mil = 0.001 inch). However, successful operation is possible at lengths up to 6,000 mils; trace lengths of this magnitude are often seen in applications using one or more DIMMs. Since the circuit is properly terminated, waveform integrity remains excellent over the typical range of trace impedances and lengths. Refer to [Figure 4](#) for typical simulation results for a 1,066 Mb/s random data stream across a trace length of 3,000 mils. [Figure 4](#) shows a wide-open data eye that meets all waveform integrity requirements of the DDR3 JEDEC standard [\[Ref 2\]](#) with very little pattern-dependent jitter. Simulation of both Fast and Slow driver cases results in similar waveforms.




Figure 4: Data WRITE Operation Recommendations

Data READ in Single-Ended DQ and DM Nets

In the READ case, the driver is within the SDRAM and the receiver within the FPGA. In this scenario, the Virtex-6 FPGA provides an SSTL I/O standard (IBIS model name: *Virtex6_SSTL15_DCI_I*) and includes a split Thevenin termination. The SDRAM I/O buffers usually provide two different drive strengths identified by their output impedance values, 34Ω and 40Ω . While both lead to nearly identical waveforms, the best noise margin is achieved using the 34Ω driver, which is recommended.

The recommended interconnect impedance and trace lengths are the same as those determined for the WRITE case. Refer to Figure 5 for typical simulation results for a 1,066 Mb/s random data stream across a trace length of 3,000 mils. Figure 5 shows a wide-open data eye that meets all waveform integrity requirements of the DDR3 JEDEC standard [Ref 2] with very little pattern-dependent jitter. Simulation of both Fast and Slow driver cases results in similar waveforms.




Figure 5: Data WRITE Operation Recommendations

Data WRITE and READ in Differential DQS Nets

To reduce delay uncertainty, the models for both the FPGA and SDRAM must be the same as those used for the DATA signals. The resulting simulations of the DQS nets are shown in Figure 6 and Figure 7 for the WRITE and READ cases, respectively. DQS uses differential technology and the excitation is periodic. As in the DQ and DM nets, trace lengths of 500 mils–6,000 mils and differential impedances of 80Ω–120Ω are acceptable. In practice within this range, 100Ω of differential impedance is easiest to achieve, and is therefore recommended. It can be seen in Figure 6 and Figure 7 that both rising and falling edges are monotonic and that there is adequate noise margin. Fast and slow driver case waveforms also show similar behavior, and over/undershoot limits are met in all cases.




Figure 6: DQS WRITE Operation Recommendations

In the case of a DIMM-based topology (shown in Figure 2), recommendations for the FPGA and memory I/O, trace impedances, and trace lengths are precisely the same as those described earlier in this section. The signals do need to propagate through the DIMM connector (the only difference when compared to SDRAM usage), but the connector presents only a minimal discontinuity at data rates of 1,066 Mb/s. Consequently, the DATA and DQS waveforms are largely similar to those shown in Figure 4 through Figure 7. For layout convenience, the designer can keep the single-ended trace impedance at 50Ω and the differential impedance at 100Ω.




Figure 7: DQS READ Operation Recommendations

CLOCK, ADDR, CMD, and CONTROL

Fly-By Topology and Maximum Data Rates

The CLOCK, ADDR (address), CMD (command), and CONTROL nets are point-to-multipoint 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 SDRAM. The differential clock net is used as the reference signal for timing analysis. The ADDR, CMD, and CONTROL signals are latched only at the rising edge of the positive clock signal; consequently, the maximum effective data rate is 533 Mb/s for a 533 MHz clock signal.

External Termination

ODT is not available for these nets, and an external discrete termination is required. The recommended form consists of a resistor placed at the far end, past the last memory device, and pulled up to $\frac{1}{2}V_{DD}$. The value of the pull-up resistor and the impedance of the interconnecting traces depend on the number of devices on the net. These values are optimized through simulation. A single-ended trace impedance of 50Ω and a pull-up resistor value of 50Ω is adequate for the ADDR, CMD, and CONTROL nets in most cases. The RESET and CKE signals are not terminated. [Ref 1]

For the CLOCK differential pair, a trace differential impedance of 100Ω and the use of two separate pull-up termination resistors of 50Ω is recommended.

For these unidirectional signals, Virtex-6 FPGAs provide an SSTL I/O standard (IBIS model name: *Virtex6_SSTL15*). The memory devices also have a unique input buffer at the receiver. The interconnecting trace can be broken down into three parts:

- TL1 between the FPGA and the first memory
- TL2 between each memory
- TL3 between the last memory and the termination

Typical timing relationships along the interconnecting trace are shown in Figure 8 and Figure 9.




Figure 8: **CLOCK WRITE Operation Recommendations**




Figure 9: **ADDR, CMD, and CNTRL WRITE Operation Recommendations**

A short stub is present at each memory device pin. Simulations show that optimal waveform integrity is achieved when the lengths of TL1, TL2, TL3, and the stubs are *kept as short as possible*. Typical practical values are 1,000 mils–3,000 mils for TL1, less than 1,000 mils for TL2 and TL3, and less than 100 mils for the stubs. Simulated results shown in Figure 8 and Figure 9 derive from values of:

- TL1 = 3,000 mils
- TL2 and TL3 = 800 mils
- Stub Length = 100 mils

Device models both with and without DCI produce largely identical waveforms; therefore, either model can be used. Figure 9 shows a wide-open data eye that meets all waveform integrity requirements of the DDR3 JEDEC standard [Ref 2] with very little pattern-dependent jitter. Simulation of both fast and slow driver cases results in similar waveforms.

In the case of a DIMM (shown in Figure 2), the implementation is much simpler, as the DIMM contains all the required termination circuitry. The only difference is that the signals need to propagate through the DIMM connector, which has little effect at clock (data) rates of 533 MHz (533 Mb/s).

For convenience, the designer can keep the single-ended trace impedance at 50Ω and the differential impedance at 100Ω.

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 synchronous nets be matched very closely. All DQ and DM nets in a byte lane must be matched to their associated DQS nets. The simulated results shown in Figure 4 and Figure 6 are superimposed in Figure 10. V_{REF} (0.75V) and 0V are used as the thresholds for DQ and DQS, respectively.




Figure 10: Delay Between DQS and DQ at Receiver for a Trace Length of 3,000 mils

It can be seen that the differential signal (offset by -2V for display clarity) has a slightly smaller delay compared to the single-ended signals. This is due to the reference levels

used. To maximize timing margins, it is recommended that the transmission line of a DQ/DM net be matched to its associated DQS net within ± 5 ps, which is easily achievable.

For unidirectional signals, all ADDR, CMD, and CONTROL signals must be matched to the CLOCK signal. The simulated results shown in Figure 8 and Figure 9 are superimposed in Figure 11. Even though trace lengths have been matched exactly between the two topologies, the relative delay is not consistent. This is due to the different loading conditions for the single-ended and differential cases. A maximum delay uncertainty of 28 ps exists between the first and last SDRAM.




Figure 11: Delay between CLOCK and ADDR at Each Receiver

Minimizing trace and stub length is critical in PCB layout. Any increase in these lengths above recommendations increases the delay uncertainty. Since the data rate of these nets is lower, a relative delay tolerance of ± 25 ps at each receiver is recommended.






Further, *piece-wise delay matching* is preferred: that is, each transmission line segment — for example, DTL1 on the clock — must be matched to the corresponding transmission line segment TL1 on the ADDR, CMD, and CONTROL nets. A summary of the delay matching requirement is given in Table 2 and Table 3.

Table 2: Delay Matching Requirement for Figure 1

Net/Group	Topology	Match	To	Tolerance (±ps)
Clock/ ADDR	Treat as Master WP420_Tab2_A_041812			
ADDR[0-15], CKE, CS, ODT, RAS, CAS, WE, BA[0-2]/ ADDR	Treat as Master WP420_Tab2_B_041812	TL1 TL2 TL3 STUBS	DTL1 DTL2 DTL3 DSUBS	5 5 5 2
DQS[0]/Byte_0	Treat as Master WP420_Tab2_C_041812			
DQ[0-7], DM[0]	WP420_Tab2_D_041812	TL	DTL	5
DQS[n]/Byte_n	Use same rule as for group Byte0. DQS[0-3] need not match to each other.			

In Table 3, DIMM CONN denotes the DIMM connector. The DIMM module already has tolerance associated with its onboard delay matching. Thus, when designing the main board, it is better to use a tighter tolerance (e.g., ±5 ps) for *all* nets. This tighter tolerance is still easily achievable.

Table 3: Delay Matching Requirement for Figure 2

Net/Group	Topology and Min/Max Constraints	Match	To	Tolerance (±ps)
Clock/ ADDR				
Additional Clocks/ ADDR		CD	AB	5
ADDR[0-15], CKE, CS, ODT, RAS, CAS, WE, BA[0-2]/ ADDR		EF	AB	5
DQS[0]/Byte_0				
DQ[0-7], DM[0]		IJ	GH	5
DQS[n]/Byte_n	Use same rule as for group Byte0. DQS[0-3] need not match to each other.			


To ensure a matched delay, trace length matching is carried out during PCB layout. It is important to pay attention to four important factors (Figure 10) that require compensation.

- Depending on the FPGA pin assignment, a substantial amount of skew within the package traces can exist. This should be compensated for on the PCB by appropriate lengthening or shortening of the associated PCB traces. Package trace lengths of each Virtex-6/Spartan-6 FPGA are available using the Xilinx PARTGen utility. The length in millimeters should be converted to a time delay value in picoseconds by multiplying it by 6.5. [Ref 1]

- The velocity of propagation of a microstrip trace cm is greater than that of a strip line trace cs . Consequently, velocity compensation is required when traces in a group are routed as both microstrip (exposed) traces of total length lm and stripline (buried) traces of total length ls . The total delay is $((lm/cm) + (ls/cs))$.

Advanced PCB layout software has features to compute the total delay based on the trace type, making this requirement easy to ensure. Otherwise, compensation must be implemented by computing the total delay and then elongating or shortening the microstrip or stripline parts appropriately.

- Bending traces in a trombone shape is a simple technique to increase trace length. However, the electrical delay of a trombone-shaped trace is smaller than that of a straight trace due to coupling between the parallel trace segments. This coupling can be reduced by ensuring that the spacing between parallel segments (L3 in Figure 12) is, in most cases, ≥ 25 mils.



WP420_12_041812

Figure 12: Factors Leading to Differences in Delay

- Layer-jumping using vias is typically unavoidable. The delay of a trace containing a via is greater than that of a straight trace by about 10 ps. This arises from the loading effect of a via; it depends upon the geometric parameters of the via, anti-pad, layer stackup, and the location of return vias.

It is recommended that GND planes be used as reference planes for all signals, in which case the return via is a GND via. Consequently, it is important to ensure that a signal via has at least one GND via in close proximity. The ideal situation is where each signal via is surrounded by four closely spaced GND vias leading to what is termed a *controlled impedance via*. However, this is often impractical for dense layouts. It is acceptable to place a few GND vias close to the region of signal vias. Further, all ADDR, CMD, and CONTROL nets can be referenced to a V_{DD} power plane if needed.

Mitigating Potential Crosstalk and Power Integrity Problems

Crosstalk adds to jitter. Jitter reduces the eye opening, which then leads to data dependency. This can become significant for single-ended data traces if they are routed as microstrip lines.

Trace Spacing

Crosstalk is reduced by increasing the spacing between adjacent traces in long parallel runs. This has the drawback of increasing the total trace length; therefore, a reasonable value must be chosen.

The distance between a trace and its nearest reference plane (dr) plays an important role in this determination. Typically, the edge-to-edge spacing between traces should be $> 2dr$ for stripline traces and $> 7dr$ for microstrip traces. Increasing the number of GND vias and maximizing the use of stripline routing is recommended to help keep crosstalk levels manageable. The same trace-spacing rule is also applicable to differential signals, such as the clock and DQS.

Maintaining Power Integrity

In the context of this paper, power integrity refers to meeting the circuit board's power supply tolerance requirements (see [Table 1](#)) under worst-case (maximum) switching conditions. Failure to address this requirement can lead to a number of problems, such as increased jitter and crosstalk, all of which eventually reduce timing margins. This, in turn, ultimately forces reduced data-rate operation.

Decoupling theory is very well understood, and usually starts with definition of a target impedance that must be met over a predetermined frequency range. [\[Ref 3\]\[Ref 4\]](#) The three power rails of concern are the V_{DD} , V_{TT} , and V_{REF} . The tolerance requirements on the V_{DD} and V_{TT} rails can be met in several different ways; using a shape on a plane layer with the required number of decoupling capacitors of three to five different values is recommended. It is important to design the capacitor pad mounting structure for minimal mounted inductance. [\[Ref 3\]\[Ref 4\]](#)

The V_{REF} rail has a tighter tolerance than V_{DD} and V_{TT} ; fortunately, it draws very little current. Its target impedance is easily met using narrow traces and one or two decoupling capacitors in the range of 0.01–0.1 μF . *It is important that these capacitors be placed very close to the device pins.*

Spartan-6 FPGAs

In the simulations previously described, Virtex-6 FPGA IBIS models have been utilized. Where the FPGA is a Spartan-6 device, available SSTL15 models provide different options.

For example, SSTL 1.5V output buffers are available with untuned output resistance termination values of [no termination], 25 Ω , 50 Ω , and 75 Ω , and V_{CCO} and V_{CCAUX} values of 2.5V and 3.3V.

SSTL 1.5V input buffers are available with untuned split termination values of [no termination], 25 Ω , 50 Ω , and 75 Ω , and V_{CCO} and V_{CCAUX} values of 2.5V and 3.3V.

For DATA WRITE operations at 800 Mb/s, the *Spartan6_SSTL15_OT25_LR_33* or *Spartan6_SSTL15_II_LR_33* model results in eye opening values that are close to those

obtained using Virtex-6 FPGAs. These are 1.5V SSTL output buffers with or without 25Ω of untuned termination and an auxiliary voltage of 3.3V.

Note: *LR* in the model name refers to the left/right banks; the same buffer is also available for the top/bottom banks.

For DATA READ operations, use of the *Spartan6_SSTL15_IN50_LR_33* model (1.5V SSTL input buffers with 50Ω of untuned split termination and an auxiliary voltage of 3.3V) results in eye opening values that are close to those obtained using Virtex-6 FPGAs.

For the CLOCK, ADDR, CMD, and CONTROL nets, use of the *Spartan6_SSTL15_II_LR_33* model (unterminated Class II SSTL buffers with an auxiliary voltage of 3.3V) results in eye opening values close to those obtained using Virtex-6 FPGAs.

All other memory settings and implementation guidelines remain the same as those for Virtex-6 devices.

Guidelines for Implementation

PCB Layer Stackup

Achieving a good balance between cost, performance, and complexity is the ultimate goal when defining any PCB layer stackup. This potentially iterative task is limited, however, because the layer count is almost always driven by the breakout requirements of the Virtex-6/Spartan-6 FPGA package and of other high-density component packages in the design. In [Figure 13](#) and [Figure 14](#), 12-layer and 14-layer examples are shown as an optimal starting stackup that can be conveniently expanded, if required, to accommodate more layers.

12-Layer Stackup

The 12-layer stackup model shown in [Figure 13](#) is constructed with low-cost, RoHS-suitable FR4 material with an assumed relative permittivity of 4.2. The computed trace widths and spacings shown provide an easily realizable single-ended impedance of 50Ω and a differential impedance of 100Ω. The stackup uses solid GND planes as references for all routing layers to ensure uniform trace impedance. Where “layer-jumping” is necessary, return path continuity is easily achieved by inserting GND vias close to the signal vias.

Power layers are located in the middle of the board, sandwiched by solid GND planes. This enables power-plane splitting without affecting signal routing. In addition, this topology allows convenient placing of decoupling capacitors on either the top or bottom layer of the board, as the effective length of their vias is nearly the same.

Layer #	Thickness (mils)	Drill	Cross Section Diagram	Layer Type	Layer Definition	Single-Ended Line		Impedance Ref. Layer	Edge-Coupled Diff	
						Width(mils)	Impedance		Width Space	Impedance
	0.5			mask						
	1.4			plating						
L01	0.600			foil	TOP	7.0	50.0 Ω	L02	5W 9 sp	100.0 Ω
L02	5.0			Prepreg						
	0.6				GND1					
L03	5.0			0.5/0.5 Core	SIG1 - X	5.5	50.0 Ω	L02,L05	5W 10sp	100.0 Ω
L04	4.0			Prepreg						
	0.6				SIG2 - Y	5.5	50.0 Ω	L02,L05	5W 10sp	100.0 Ω
L05	5.0			0.5/0.5Core						
	0.6				GND2					
L06	4.0			Prepreg						
	0.6				PWR1-SPLIT					
L07	5.0			0.5/0.5 Core						
	0.6				PWR2-SPLIT					
L08	4.0			Prepreg						
	0.6				GND3					
L09	5.0			0.5/0.5 Core	SIG3 - X	5.5	50.0 Ω	L08,L11	5W 10sp	100.0 Ω
L10	4.0			Prepreg						
	0.6				SIG4 - Y	5.5	50.0 Ω	L08,L11	5W 10sp	100.0 Ω
L11	5.0			0.5/0.5 Core						
	0.6				GND4					
L12	5.0			Prepreg						
	0.600			foil	BOTTOM	7.0	50.0 Ω	L11	5W 9 sp	100.0 Ω
	1.4			plating						
	0.5			mask						
Total:	62		Finish thickness $\pm 10\%$							

WP420_13_041812

Figure 13: 12-Layer PCB Stackup with FR4 Dielectric

Dual stripline construction is employed to reduce layer count. Routing in these layers should be perpendicular for crosstalk reduction. This layer stackup is a reasonably low-cost, easy-to-use, practical solution for many medium-complexity designs.

Nevertheless, the designer must always be aware of excessive trace parallelism — for example, when designing with a DIMM. In such a case, it is better to use *single stripline* construction, shown in the 14-layer stackup in Figure 14.

14-Layer Stackup

Layer #	Thickness (mils)	Drill	Cross Section Diagram	Layer Type	Layer Definition	Single-Ended Line		Impedance Ref. Layer	Edge-Coupled Diff	
						Width (mils)	Impedance		Width Space	Impedance
	0.5			mask						
	1.4			plating						
L01	0.600			foil	TOP	8.0	50.0 Ω	L02	6W 9sp	100.0 Ω
L02	5.0			Prepreg						
	0.6				GND1					
L03	4.0			0.5/0.5 Core	SIG1	4.0	50.0 Ω	L02,L04	3.5W 10sp	100.0 Ω
L04	4.0			Prepreg						
	0.6				GND2					
L05	4.0			0.5/0.5Core	SIG2	4.0	50.0 Ω	L04,L06	3.5W 10sp	100.0 Ω
L06	4.0			Prepreg						
	0.6				GND3					
L07	2.0			ZBC2000						
	0.6				PWR1-SPLIT					
L08	4.0			Prepreg						
	0.6				PWR2-SPLIT					
L09	2.0			ZBC2000						
	0.6				GND4					
L10	4.0			Prepreg						
	0.6				SIG3	4.0	50.0 Ω	L09,L11	3.5W 10sp	100.0 Ω
L11	4.0			0.5/0.5 Core						
	0.6				GND5					
L12	4.0			Prepreg						
	0.6				SIG4	4.0	50.0 Ω	L11,L13	3.5W 10sp	100.0 Ω
L13	4.0			0.5/0.5 Core						
	0.6				GND6					
L14	5.0			Prepreg						
	0.600			foil	BOTTOM	8.0	50.0 Ω	L13	6W 9sp	100.0 Ω
	1.4			plating						
	0.5			mask						
Total:	62		Finish thickness $\pm 10\%$							

WP420_14_041812

Figure 14: 14-Layer PCB Stackup with Nelco 4000-13EP or Isola FR408HR Dielectric

The 14-layer stackup model shown in [Figure 14](#) uses material with a lower dielectric constant, such as Nelco 4000-13EP or Isola FR408HR. The required trace impedances are conveniently realized using common trace widths. Improved power integrity can be expected with this model compared to the 12-layer stackup. This is attributable to the use of thinner dielectric layers between the power and ground planes, resulting in high-frequency decoupling benefits.

Due to individual process and material variation, the displayed stackups must be reviewed, verified, and potentially altered by the PCB shop prior to fabrication.

BGA Breakout Optimization

In high-density designs, following placement, land pattern breakout is often the most time-consuming part of the layout process. Breakout is defined as the planned escape from the package pin to outside of the package body (for example, the BGA pad array). In the context of DDR2 and DDR3 connectivity to Virtex-6/Spartan-6 FPGAs, this consists of associating the appropriate FPGA pin with the appropriate DDR2/DDR3 signal within the limitations imposed by the FPGA pin mapping. [\[Ref 1\]](#) With the flexibility of the Virtex-6/Spartan-6 FPGAs, this is most easily done by first placing components in the desired location on the board and then routing (optimally) backwards from the DIMM or discrete SDRAMs towards the FPGA, stopping the routes just short of the previous breakout demarcation points. The FPGA pins can then be carefully swapped to ensure that the route can be completed with little or no crossover.

Completing Placement and Routing

After an appropriate stackup and optimized BGA breakout is established, critical components are then placed.

V_{REF} decoupling capacitors should be placed close to the device pins. V_{TT} pull-up resistors and decoupling capacitors can be grouped near the last memory. Smaller value V_{DD} decoupling capacitors can be distributed near the pins of each device. *Pay careful attention to the placement of these components to avoid blocking routing channels.* Larger value and bulk decoupling capacitors can be placed at convenient locations away from most routing.

All decoupling capacitors should use a reduced-inductance footprint. [\[Ref 3\]](#)[\[Ref 4\]](#) A reduced-inductance footprint is typically realized using two short wide traces from each pad perpendicular to the capacitor length, pinned down with a low-inductance via. For identification and certainty of deployment, a reduced inductance footprint can be assigned within the schematic and driven into the layout through the netlist. The board file should be set up for trace impedance, min-max length, delay matching, and spacing constraints, as described in sections [Waveform Integrity](#) and [Delay Matching of Signal Nets](#). First, attempt coarse routing with reduced stub length using the minimum trace spacing required to meet crosstalk rules; then, implement fine length matching, ensuring that multiple ground vias are inserted where signal traces jump layers.

Verification Checklist

Table 4: Verification Checklist

	Task	Verified
1	All decoupling capacitors use a reduced-inductance footprint	✓
2	V _{REF} decoupling capacitors are placed close to V _{REF} pins	✓
3	Trace widths and spacings produce the correct impedance	✓
4	Correct P/N skew constraints exist on all differential pairs	✓
5	CLOCK ADDR delays match tolerance	✓
6	DQS/DATA and DM delays match tolerance for all byte lanes	✓
7	GND vias are present near signal vias	✓
8	Adequate spacing exists between parallel traces of a meander line	✓
9	Reference planes are solid GND and do not contain large slits or slots	✓

Conclusion

Xilinx Virtex-6 and Spartan-6 devices are proven to interoperate with DDR2/3 speeds at up to 1,066 Mb/s and 800 Mb/s, respectively. The purpose of this white paper has been to illustrate best practices and design guidelines to optimize I/O performance and reduce risk of performance issues in first-article prototypes. For additional information or design support, contact your Authorized Xilinx Distributor or [Fidus Systems, Inc.](#)

References

1. Xilinx: [UG406](#), *Virtex-6 FPGA Memory Interface Solutions User Guide*
2. JEDEC: [JESD79-3](#), *DDR3 SDRAM Standard*
3. Xilinx: [UG072](#), *Virtex-4 FPGA PCB Designer's Guide*
4. Xilinx: [WP411](#), *Simulating FPGA Power Integrity Using S-Parameter Models*

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
06/26/12	1.0	Initial Xilinx release.

About Fidus Systems

Fidus Systems, a Xilinx Premier Design Services Alliance Member, provides Electronic Product Development, Staff Augmentation, and Consulting Services to companies across a wide range of industries. Fidus' extensive experience includes System Level, FPGA/DSP, RF/Wireless, Signal/Power Integrity, Hardware, PCB Layout, Embedded Software, and Mechanical Design. As a design partner, Fidus offers companies greater flexibility and capability in their product development with access to the expertise, process, and tools to successfully move their products to market. Fidus has delivered on more than 800 products and projects for 215 customers across North America. For more information, please visit www.fidus.com.

Notice of 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 the Limited Warranties which can be viewed at <http://www.xilinx.com/warranty.htm>; 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 Critical Applications: <http://www.xilinx.com/warranty.htm#critapps>.