

Getting Zen with Soldered-down Memory Layout for AMD Based Embedded Systems

禪 Zen—*Enlightenment through Exhaustive Study*

DDR3 memory layout is tough enough for standard memory modules because of the myriad guidelines which include trace length matching, spacing, and impedance requirements, but soldered down memory adds a whole new dimension to the problem for embedded systems. Layout for soldered down is often exponentially tougher because many desired placement topologies are not documented by AMD's Motherboard Design Guide (MBDG) and therefore not supported in the verifications tools, specifically AMD's Net Tool. The memory recommendations in AMD Motherboard Design Guides originated for PC applications which always implemented memory modules (uDIMM, SODIMM). They are conservative, due to the need to assure, to the highest degree possible, successful operation across millions of systems, with a plethora of memory module suppliers, with module PCBs from hundreds of manufactures, with memory devices from several memory manufactures; a significant challenge to say the least. Soldered down memory needs primarily come from the embedded space. AMD Motherboard Design Guides have recently been updated and now document recommendations for a limited number of component placement topologies. But as it turns out, many embedded designers have requirements for topologies that do not match the documented topologies.

Because guidelines do not exist for many soldered down topologies, the embedded design engineer and layout designer are mostly on their own when they need to stray from the guidelines. This is not to say that it can't be successfully done, far from it. This paper can help the embedded design engineer and layout designer become Zen with the guidelines. Utilizing the Zen insights, they can create quality memory down layouts that are not documented in AMD MBDG. This paper will equip the embedded designer with expanded guidelines and provide justification. Once armed with this expanded knowledge, the embedded designer will be able to properly evaluate soldered down placement topologies and be confident that the selected topology will produce a quality layout.

AMD Memory Layout Guidelines Basics

AMD's memory layout guidelines are based on memory module topologies. Both DIMM and SO-DIMM modules fundamentally utilize the same topology. For modules, Address/Command/Control/Clock (ACC/CLK) routing is based on what is called "flyby". See Figure 1. This means these signals are connected to the memory device on the far edge of the board first and then connected in succession to each device and terminated at the end. On the modules, the trace length between devices is tightly controlled to create a uniform timing delay device to device. The termination means that these signals are of the incident wave transmission type. From a propagation standpoint each successive device sees the transition on the signals later and later. This aspect of the "flyby" must be accounted for in the DRAM controller using interface training. The JEDEC standard for memory devices specifies one training procedure and how this training is supported in the DDR3 devices. AMD DRAM controllers support this training and two additional training procedures to establish the relative timing differences for each device. This training will be covered in more detail later.

Data signals are connected on the memory module with very tightly controlled lengths as well, keeping the delay from connector edge to device nearly identical. The DDR3 devices have on-die termination

(ODT) to support incident wave transmission. AMD DRAM controllers have training capabilities allowing a high degree of byte-lane-to-byte-lane length difference while requiring tight length matching within a byte lane. This will also be covered in more detail later.

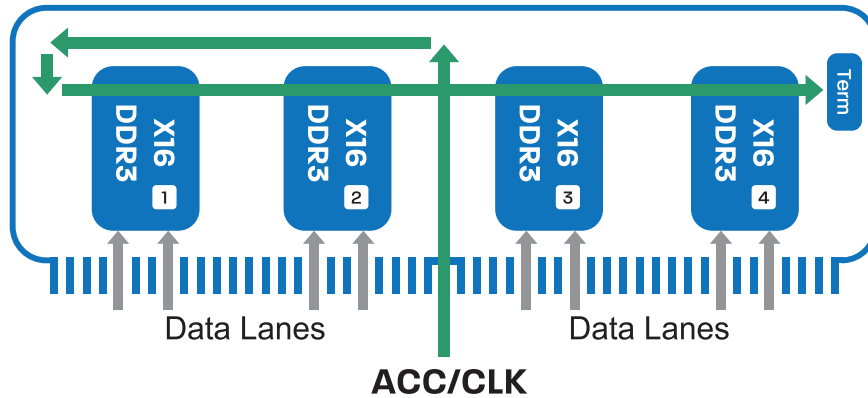


Figure 1 – Basic Routing Paths of SO-DIMM and DIMM Memory Modules

Key Zen Insights Regarding DDR3 Memory Modules

- 禪 The ACC/CLK signals are upwards of 220mm in length on a DIMM and upwards of 190mm in length on a SO-DIMM.
- 禪 The ACC/CLK signals are much longer than the Data signals.
- 禪 From Figure 1 the DDR3 device 1 has the shortest read/write time and DDR3 device 4 has the longest read/write time.
- 禪 Interface training compensates for these timing mismatches.

Memory Interface Training

AMD DRAM controllers perform three training steps to compensate for the timing differences to each DDR3 memory device:

- ▶ Write Leveling
- ▶ DQS Receiver
- ▶ DQS Position

Write Leveling

The Write Leveling procedure allows the DRAM controller to adjust for the timing skew (effective trace length difference) between DQS and CLK at the DRAM. This procedure takes advantage of a feature already in the DDR3 memory devices. The AMD algorithm uses a seed value per byte lane as a beginning setting in the DRAM controller. Seeds have been pre-determined for UDIMM and SO-DIMM

by characterization of AMD reference designs and are documented in the BIOS Kernel Developers Guide (BKDG) for each AMD processor. These will work for the majority of DIMM and SO-DIMM designs; however, for soldered-down memory reliable initial seed values are not known. Typically an iterative process is used to determine the best seed value for a given design.

The AMD DRAM controller issues a command to the DDR3 memory devices to enter write leveling mode. In this mode the DRAM controller drives the clock and all DQS signals with an initial delay on the DQS set by the seed value. With DIMM and SO-DIMM modules the clock trace is always quite a bit longer than the DQS due to its routing on the module itself (Figure 1). Without any programmed delay, the DQS will arrive at the DDR3 memory device before the clock. Starting with the seed value delay for DQS, the DRAM controller will output a rising clock followed by a rising DQS. The DQS signal is used by the DDR3 memory device to sample the clock signal. Memory data (DQ) is used to output the sampled value. If the clock signal sampled low, that means that the DQS arrived at the DDR3 memory before the rising clock edge. If sampled high, DQS arrived after the rising clock edge. Since the DQS signal is shorter and hopefully the seed programmed skew is not too much, the DRAM controller will see a low on the DQ. It will then add more delay and try again. The DRAM controller will iterate until it sees a high on the DQ, and when that happens the DRAM controller now has established the timing skew between DQS and clock. Most DDR3 memory devices only use the least significant bit of each byte lane to output the sample but the AMD DRAM controller looks for a transition on any bit of a byte lane. This allows the embedded designer to route any bit of a byte lane from the AMD processor to any bit on a byte lane on the DDR3 memory device, or to put it another way, to swizzle the byte lane connections.

DQS Receiver Training

DQS Receiver Training determines the propagation delay of DQS from the DRAM back to the controller for read timing optimization. This depends on the DQS signal lengths, loading, CLK signal length to each DDR3 device, and CLK to DQS delay within each DDR3 device. Each byte lane has a seed and, just as in write leveling, the recommended values for DIMM and SO-DIMM have been determined by characterization. For soldered-down design typically the DQS signal length is shorter than when modules are used, so the seed values are typically smaller.

DQS Position

DQS Position coordinates the timing of the data (DQ) writes and reads with the strobes (DQS). For reads the DDR3 device outputs data and strobe coincident. The AMD DRAM controller internally delays the strobe to latch the data midway between data transitions. This training determines the amount of delay. For writes, the AMD DRAM controller adjusts the data relative to the strobe so that the strobe arrives midway between data transitions. This is a brute force training process and does not use seed values. Failures here can often be traced back to other problems such as poor length matching somewhere, excessive crosstalk, or poor signal impedance control. The test positions the DQS strobe between the DQ (Data) signal transitions. This training process is the only one that fully exercises the entire interface.

Address/Command/Control/Clock (ACC/CLK)

ACC/CLK Length Matching Basics

Figure 2 shows the ACC/CLK connections to two ranks of x8 DDR3 devices. This figure is not meant to imply any component placement but does show the basis of the AMD ACC/CLK routing guideline for this memory organization. Depending on the number of devices in an actual design, devices may be placed only on top or on top and bottom. For top and bottom placement, devices are usually mirrored for better routing control. If no bottom DDR3 device placement, ACC/CLK can be routed completely microstrip except for L4.

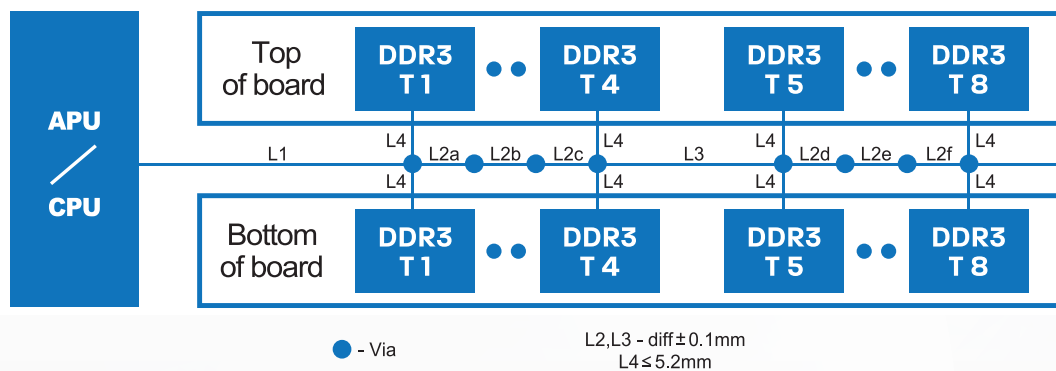


Figure 2 – ACC/CLK Length Matching

AMD has length matching recommendations for a specific soldered-down topology in the MBDGs that allow for a length differential min to max of 6.35mm (250mil) to the first connected DDR3 device. The length matching recommendations via to via, L2x and L3, is 0.1mm (~4mil). According to the MBDGs, L4 must be less than or equal to 5.2mm. By the letter of the guideline:

- ▶ The longest to shortest ACC, to the first two DDR3 devices, L1+L4, can be up to 11.55mm (~450mils).
 - a. The longest L1 to the shortest L1 just meets the recommendations of 6.35mm
 - b. The shortest L1 signal is connected directly to T1 or B1 at 0mm
 - c. The longest L1 signal is connected to T1 or B1 at 5.2mm.
- ▶ While this via to ball connection length may be impractical and probably not the guideline intent, a Zen insight would suggest some length matching latitude.

Also consider that the MBDGs express L2x and L3 in absolute lengths dramatically reducing component placement latitude, thus with the advantage of Zen insight becomes overly conservative.

- ▶ Nominally L1 can be as short as 25.4mm and as long as 127mm.
 - a. Length matching of the ACC signals to the first via, L1, is specified to be 6.35mm
 - b. The supportable absolute length can vary by ~100mm, L1 max minus L1 min.
- 禪 *Zen insight:* The specified absolute lengths for soldered-down memory between the DDR3 devices are actually arbitrary, simply revealing their connection to memory module layouts where the absolute lengths are vital for interoperability.

- 禪 *Zen insight:* Length matching from APU ball to DDR3 device ball calculated for each DDR3 device actually supplies the important guideline making the absolute length to each device not important except for total length.

ACC/CLK Length Matching with Zen Insights

Now that we have some basic Zen insights laid down, we can expand the application to ACC/CLK length matching.

- ▶ ACC/CLK length matched APU/CPU ball to DDR3 ball: $\pm 6.35\text{mm}$ (250mil) calculated for each DDR3 device.
 - a. If ACC/CLK is routed microstrip or stripline for the entire path, apply rule normally.
 - b. *Zen insight:* If ACC/CLK is routed microstrip only in the bus channel, also apply rule to via where ACC becomes stripline which should be near the first connected DDR3 device.
 - c. *Zen insight:* If the suggestions cannot be met, then the differing flight times of microstrip and stripline must be taken into account. Rule of thumb is 5.91ps/mm for microstrip, 7.09ps/mm for stripline.
- 禪 *Zen insight:* Minimum ACC/CLK length (APU/CPU ball to any DDR3 ball) is 38.1mm (1500mil). This allows the DQ/DQS/DM signals minimum length requirements to be met. See DDD length matching with Zen insights (See final point under “DDD Length Matching With Zen Insights”)
- 禪 *Zen insight:* Maximum ACC/CLK length (APU/CPU ball to last DDR3 ball): max 254mm (10000mil).
 - a. Based the routing to UDIMM.
 - b. Being conservative, the maximum length on the motherboard is specified to be 101.6mm, plus the additional length of the UDIMM routing. See Figure 1.
- ▶ CLK differential pair average length APU/CPU ball to DDR3 ball is longer than the shortest ACC trace by no more than 3.25mm calculated for each device.
 - a. If CLK is routed microstrip or stripline for the entire path, apply rule normally.
 - b. If CLK is routed microstrip only in the bus channel, also apply rule to via where ACC/CLK becomes stripline near the first DDR3 device.
- ▶ CLK differential pair average length APU/CPU ball to DDR3 ball is shorter than the longest ACC trace by no more than 3.25mm calculated for each DDR3 device.
 - a. If CLK is routed microstrip or stripline for the entire path, apply rule normally.
 - b. If CLK is routed microstrip only in the bus channel, also apply rule to via where ACC/CLK becomes stripline near the first DDR3 device.

ACC/CLK Trace Spacing

Spacing guidelines established in the MBDGs should be met to keep crosstalk levels at an acceptable level.

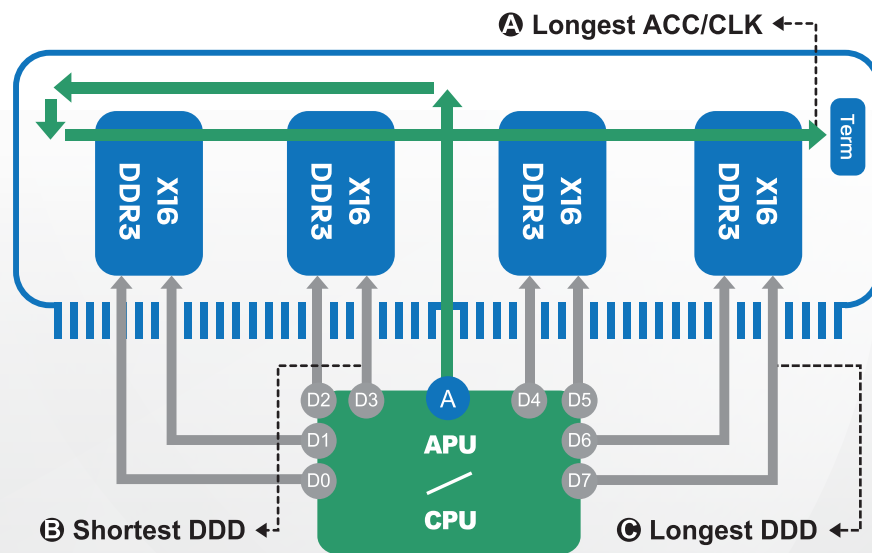
ACC/CLK Trace Impedance

Impedance guidelines established in the MBDGs should be met to achieve acceptable signal quality.

Data

DQ/DQS/DM (DDD) Length Matching Basics

On the memory modules Data/Data Strobe/Data Mask (DQ/DQS/DM or DDD) signals are much shorter than ACC/CLK signals. See Figure 3. Within a byte lane the trace matching requirements are very tight, $\pm 3\text{mm}$, but byte lane to clock, which indirectly includes byte lane to byte lane, is significantly looser. When combining the board requirements with the reality of the module trace lengths, the ACC/CLK signals are typically much longer than the DDD signals. To the last DDR3 device on the module the difference can easily exceed 200mm (~7500mil).



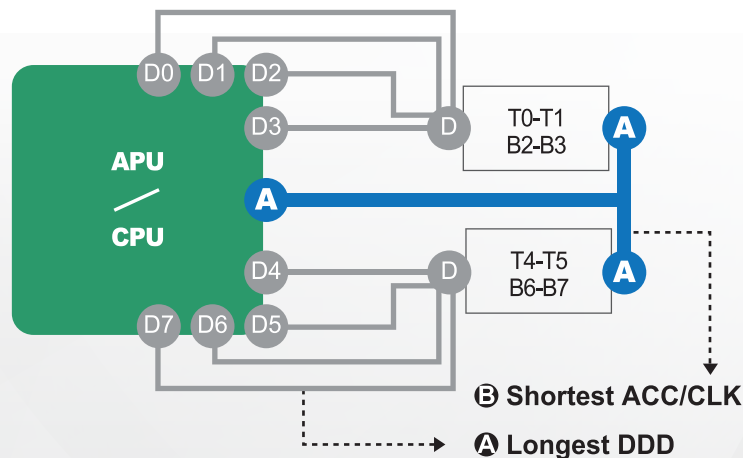
Zen recommendation 3 : $A - B < 200\text{mm}$
 Zen recommendation 4 : $A + C < 381\text{mm}$

Figure 3 – Zen Recommendations for Delay Loop for ACC/CLK and DDD

DDD Length Matching with Zen Insights

Continuing the application of Zen insights, we can move on to DDD length matching.

- ▶ DDD length matching within a byte lane must meet the MBDG guideline of $\pm 2.54\text{mm}$.
- 禱 *Zen insight:* For each DDR3 device, each DDD byte lane averaged length should be shorter than the corresponding CLK connection by at least 12.7mm (500mil). This is based on the combination of the board length matching requirements and the realities of the DDR3 memory modules.
- 禱 *Zen insight:* For each DDR3 device, each DDD byte lane averaged length should be no shorter than 200mm of the corresponding CLK connection. This is based on the combination of the board length matching requirements and the realities of the DDR3 memory modules.
- 禱 *Zen insight:* The sum of each DDD byte lane plus the corresponding CLK signal should be shorter than 381mm (15000mils). Need to limit overall loop delay.
- 禱 *Zen insight:* The above insights create an overall bounded loop delay recommendation. These three recommendations combined create a safe operational envelope for the write leveling training. The loop delay recommendation for the trace lengths of each DDD byte lane and ACC/CLK is bounded by the loop delay of the worst case timing loop of the motherboard lengths plus the DDR3 DIMM module and the best case timing loop delay defined in the MBDG for soldered-down memory, star routing of ACC/CLK signals. See Figure 3 and Figure 4.
- ▶ Minimum DDD length (APU/CPU ball to DDR3 ball) is 25.4mm (1000mil). (Referenced by Zen insight under “ACC/CLK Length Matching with Zen Insights”.)



Zen recommendation 2 : $B > A + 12.7\text{mm}$

Figure 4 - Zen Recommendations for Minimum Trace Lengths

DDD Trace Spacing

Spacing guidelines established in the MBDGs should be met to keep crosstalk levels at an acceptable level.

DDD Trace Impedance

Impedance guidelines established in the MBDGs should be met to achieve acceptable signal quality.

Example Placement Topologies

The following example placement topologies shown in Figure 5 and Figure 6 are not supported by the MBDG but will meet the Zen adjusted recommendations. Both of the examples violate the MBDG guidelines regarding the DDD trace lengths relative to CLK. The soldered-down MBDG recommendations indicate that each DQS must be no shorter than to the first device CLK is connected to by 25.4mm or no longer than 50.8mm. DQS to the furthest DDR3 device could easily be 50.8mm longer.

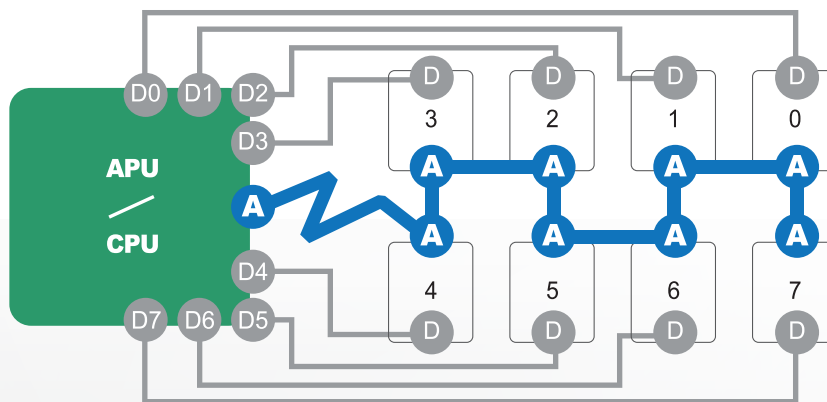


Figure 5 - Eight x8 DDR3 Device Placement Topology

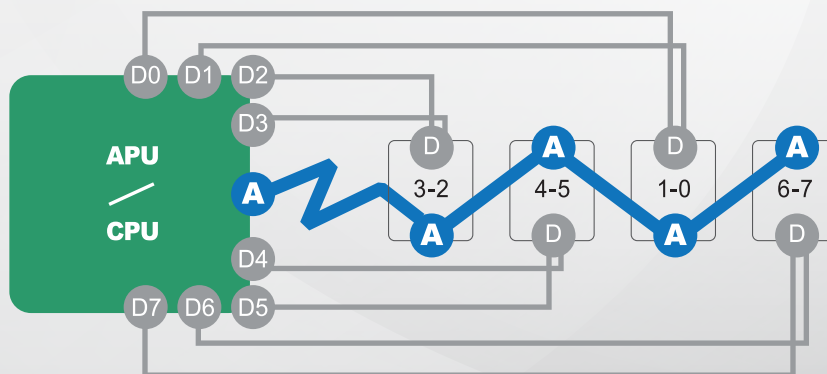


Figure 6 - Four x16 DDR3 Device Placement Topology



DQS receiver training seed values will likely be impacted by these two component topologies. Each byte lane has a seed value, and in most cases the same value can be used for each byte lane. However, the examples may need different seed values due to ever-expanding loop from byte to byte. Both the ACC/CCLK signals and the DDD signals are getting longer creating a large loop difference min to max. In the supported topologies only the ADD/CLK signals are getting longer byte to byte.

Both examples only show devices on the top side of the board; however, additional devices could be placed on the bottom side for a larger memory array.

Conclusion

AMD's MBDGs establish layout guidelines for DIMM, SO-DIMM, and very specific soldered-down designs specifically targeted at traditional PC's. If followed precisely, with no warning from the layout evaluation tool, Net Tool, high quality layouts with significant operational margin can be expected. However, many embedded systems require soldered-down topologies that are not documented in the MBDGs. Using the Zen-inspired enhanced recommendations provides the embedded designer the path to enlightenment, enabling non documented soldered-down topologies with the expectation that the necessary operational margin can be achieved.

DISCLAIMERS

The contents of this document are provided in connection with Advance Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sales, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's products could create a situation where personal injury, death or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

© 2013 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD arrow logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc. All other names used in this publication are for informational purposes only and may be trademarks of their respective owners. PID: 52881-A