TECHNICAL BULLETIN 00005

Subject: AMD Opteron Server Memory Configurations
Author: Sharon Yen
Location: Celestica - Salem NH
Date: 08-06-2004

The purpose of this Tech. Bulletin is to highlight the important memory configuration rules that should be observed when populating memory DIMMs in server systems based on AMD Opteron Processor architecture. The content of this technical bulletin is derived from Celestica A8440 AMD Opteron server product user's manual. The same rules apply to A2210 AMD dual Opteron processor server. Both A8440 and A2210 are designed with four DIMMs slots per processor.

Modifying the memory configuration on a server is fairly simple provided the user understands and follows a few basic rules. Some definitions are required:

Bank – one or more DIMMs, logically arranged to form a given memory bus access width (64 or 128 bit).

Node – a single processor, its memory controller, and all of its associated memory DIMMs and DIMM sockets. Therefore "Node" represents a given processor's memory array, as viewed from the HyperTransport™ link.

In general, always install DIMM slots in groups of two with identical DIMMs. Node 0 (CPU 0, the boot node) must be loaded with some amount of memory. In the case of single DIMM memory configuration in the server, the DIMM must be installed on the boot node, CPU 0. One cannot add memory DIMMs without the presence of the associated CPU Node.

Single DIMM per Node Configuration:

- If a given node is to receive a single DIMM, it should be stuffed in either slot DIMM0 or DIMM2. Note that this creates a 64-bit wide bank, which forces the memory controller to do two accesses to form a 128-bit element.

Two DIMMs per Node Configuration:

- If a given node is to receive two DIMMs, which are not identical, then they should be installed in **DIMM0** and **DIMM2**. This configuration creates two 64-bit wide banks, which cannot be formed into a single 128-bit bank.

- If two DIMMs are installed into adjacent slots, they must be identical. In other words, a DIMM placed in slot **DIMM1** must be identical to the DIMM placed in slot **DIMM0**. A DIMM placed in slot **DIMM3** must be identical to the DIMM placed in slot **DIMM2**. The memory controller will create 128-bit wide banks from these adjacent pairs.

Four DIMMs per Node Configuration:

- If a given node is to receive four DIMMs, which are two pairs of two DIMMs, the above rules must be observed.

Erroneous Configurations:

- It is not permitted to install only three DIMM slots on a given node. Memory controllers can be in either 128 or 64-bit mode, not both.

Additional General Information:

- Memory controllers adjust the memory clock frequency to match the slowest DIMM placed in slots DIMM0 and DIMM2. One slow DIMM will slow down the rest.

- It is recommended that non-identical DIMMs be placed on separate nodes. For example, if the user has two PC2100 DIMMs, and two PC2700 DIMMs, the user may wish to place the two PC2100 DIMMs on a less used node, and place the two PC2700 DIMMs on a more heavily used node. This allows the system to optimize the timing for each type of DIMM, rather than reduce the memory interface speed to the lowest common denominator, which would be the case if the DIMMs were installed on the same node.

- **Bank Interleave** can be manually enabled via BIOS setup and is only applied to nodes that contain two identical 128-bit banks (four identical DIMMs). There is no Bank Interleave support for 64-bit wide banks. **Bank Interleave** is defined as follows – memory controller can group two 128-bit DIMM banks into one large array. Every other 128-bit word is stored in a given DIMM pair such that even word addresses are stored in the bank composed of slots DIMM0, DIMM1 and odd word addresses are stored in the bank composed of slots DIMM2, DIMM3. If enabled in BIOS Setup, each node with four identical DIMMs is set up to use Bank Interleave.

- **Node interleave** can be manually enabled via BIOS and is only relevant to servers with even numbers of nodes (2 or 4) that have memory arrays of the same size. This implies the two following example: A four-processor server configuration with all four nodes installed with the same amount of memory, or a two-processor server configuration with both nodes installed with the same amount of memory. **Node Interleave** is defined as follows - Node based interleaving causes the system to group even numbers of nodes into one large array. In the case of two-way node interleaving (two processors present), every other 128-bit word is stored on a given node. Four-way node interleaving results in every fourth 128-bit word being stored on a given node. Node Interleave is not compatible with Microsoft's .NET SRAT Table definition. If enabled in BIOS Setup and if all loaded nodes have the same amount of memory, the system is set up to use node interleave.