DDR2 Advantages for Dual Processor Servers

White Paper

August 2004

Overview

Technology innovation across platform subsystems is necessary to keep computing platforms in balance. With processors continuing to achieve performance advantages at the rate of Moore's Law, memory and I/O subsystem improvements must keep pace. 2004 features a major step forward for the memory subsystem with the transition from DDR to DDR2 system memory.

As part of a broad technology launch across Intel® architecture platforms, 400MHz DDR2 is being intro duced in dual processor servers and workstations this summer with the Intel® E7525 chipset (previously code named Tumwater), Intel® E7520 chipset (previously code named Lindenhurst), and the Intel® E7320 chipset (previously code named Lindenhurst-VS). This initial launch in enterprise platforms will be followed up with a DDR2 introduction in quad processor server platforms with the introduction of the Twin Castle chipset. DDR2 offers greater overall platform solutions for IT managers concerned about performance, power, and platform flexibility/cost. Furthermore, DDR2 will be the primary DRAM technology for multiple generations ensuring longer-term supply and price efficiencies. This white paper aims to provide a background on why the industry is making a transition to a new memory technology, expectations for the industry's readiness for DDR2, and expectations for platform advancements with the deployment of DDR2-enabled systems.

Background

Dual processor servers have traditionally featured x4 or x8 DRAMs. The distinction between x4 or x8 is determined by different number of data outputs per DRAM; with the total amount of memory available per DIMM being the same (i.e., a x4 DRAM has twice the addresses as a x8, but each is half as big). ECC DIMMs feature 72 data bits (64 data bits plus eight ECC bits). Therefore, a DIMM with x4 devices requires 72/4 or 18 total DRAM chips. DIMMs featuring x8 devices require 72/8 or nine total chips.

This 72-bit unit of devices (18 or nine) is called a rank. All the devices on a given rank are tied to a single, unique chip select. It is possible for DIMMs to comprise two ranks, called "dual rank DIMMs." DIMMs that feature 36 DRAMs (two ranks of x4 components) are called "High Density" or "stacked" DIMMs (This does not necessarily mean that the DRAMs are physically stacked but can be all on one planar).

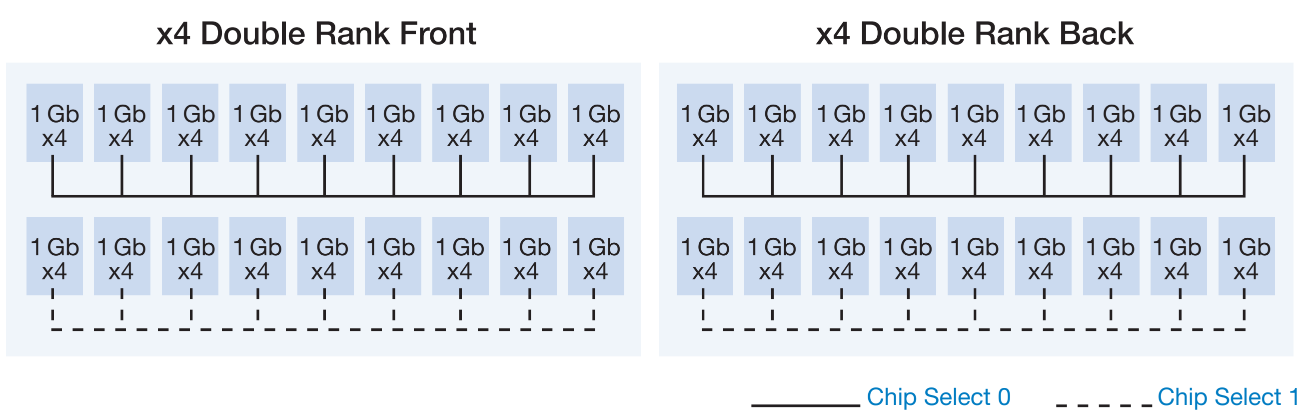

Figure 1: Conceptual Example of Chip Select on a Dual Rank x4 1 Gb (Illustrative purposes only)

DDR2 DRAM memory cell arrays are separated into equal parts called banks. DDR2 256 Mb and 512 Mb DRAM devices have four banks and DDR2 1 Gb and future 2 Gb DRAM devices have eight banks. The time to read the data, activate the bank, and pre-charge the bank are each determined by the latencies and specified in terms of clock cycles. With a DDR2 400 3-3-3 DRAM, it takes three clock cycles to start reading the data, another three cycles to activate the bank, and another three cycles to pre-charge the bank. Intel has validated DDR2 400 DRAMs 3-3-3 and 4-4-4; however, the validated DDR2 400 Reg-DIMMs have been all based on 3-3-3 DRAMs. (developer.intel.com/technology/memory/)

Industry Readiness

DDR2 transition may be the “fastest memory ramp ever” (as described by Samsung). Unlike the DDR transition, DDR2 is first available to the industry on Intel® platforms to include server, workstation, desktop and mobile concurrently. Intel has worked with the memory industry in the laborious process of validation of DDR2 DRAM components prior to launch of DDR2-enabled systems and has validated DRAMs (256 Mb, 512 Mb, and 1 Gb) from all major DRAM suppliers. Intel has also validated DDR2 Registered DIMMs from at least 15 DIMM suppliers including 256 MB, 512 MB, 1 GB, and 2 GB densities. Further validation efforts will continue as the memory industry introduces new speeds and densities of DRAM and DIMM components. The current listing of vendors’ various DDR2 DRAMs and DIMMs that have been validated by Intel is publicly available at developer.intel.com/technology/memory/.

Comparisons of DDR2 400 vs. DDR333 in Enterprise Platforms

Performance

DDR2 400 offers additional performance advantages over DDR333. This is evident in the increase of sustainable memory bandwidth in DDR2 400 platforms as measured as a gain of up to 10 percent over DDR-enabled systems as measured by STREAM TRIAD .

Of course, overall system performance is ultimately determined by multiple factors (memory type, application workload, …). Below are some rules-of-thumb suggestions specifically for increased performance related to memory.

1. Utilization of all four ranks per channel is always preferred for optimal performance; performance should not differ based on whether four ranks are spread over four DIMMs or four ranks are consolidated onto two DIMMs.
2. Identical DIMMs—Dual Rank (DR) or Single Rank (SR)—in one system are preferred because the system’s chipset can equally distribute memory addresses. Configuration with DIMMs of different ranks (i.e., mix of SR and DR DIMMs per channel) will offer sub-optimal performance versus identical DIMM configurations. If a memory upgrade is required, then the upgrade of all DIMMs is preferred over partial upgrade.
3. Dual channel implementations are always preferred over single channel implementations for performance advantages.

Lower Power

DDR2 offers significant advantages in power consumption over DDR. DDR2 has lower I/O and core voltage of 1.8V (versus the 2.5V core voltage with DDR), 4-bit prefetching to reduce core cycle time (versus 2-bit prefetching with DDR), and reduced activate and powerdown power. This combines to provide a DDR2 advantage of ~30 percent lower power than DDR266 and ~40 percent lower power than DDR333 (based on analysis from current and throughput). This decreased power consumption of several watts per DIMM (conceivable total savings of 20-30W per system) is especially important in dense rack (2U/1U) servers, and blade configurations.

Platform Flexibility

DDR2 400 supports four ranks per channel that can be spread over four DIMMs whereas DDR333 which supports four ranks per channel but is limited to only three DIMMs. This additional DIMM capacity enables more flexibility in server design.

DDR2 400: Four DIMMs/channel

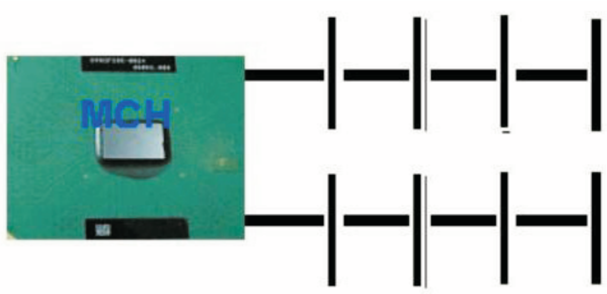

DDR333: Three DIMMs/channel

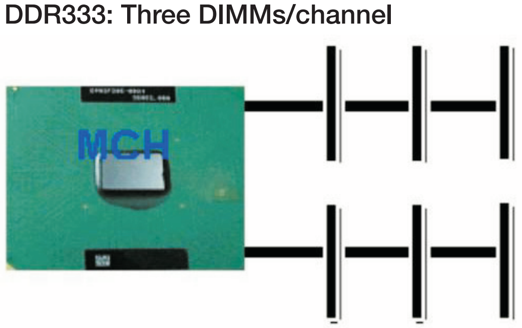

Lowering Memory Configuration Costs

In keeping with historical memory transitions, the industry expects DDR2 to be priced at a premium initially. However, given an expected broad industry adoption and a robust supplier base, DDR2 premiums are expected to decline precipitously. Most vendors predict price parity between DDR2 and DDR to occur by the middle of 2005.

Memory costs are primarily driven by DIMM capacity and by DRAM densities (1 Gb, 512 Mb, 256 Mb) rather than the nature of DIMM (SR or DR). For example, at this time a 2-GB DIMM based on SR x4 1 Gb (uses 18 DRAMs) will be more expensive than a 2-GB DIMM based on DR x4 512 Mb (uses 36 DRAMs) because 1-Gb devices are more than twice as expensive as 512-Mb devices (per bit price parity between 1Gb and 512 Mb expected sometime after 2005). Thus, implementations of >4 GB per channel require much more expensive DIMMs based on 1-Gb DRAMs rather than the lower cost 512-Mb or 256-Mb DRAMs.

Capacity per Channel	DIMMs per Channel	DIMM Capacity	DRAM Density
8 GB	4	2 GB	1 Gb
8 GB	3	Mix 4 GB & 2 GB	1 Gb
4 GB	4	2 GB	512 Mb
4 GB	3	Mix 2 GB & 1 GB	512 Mb
2 GB	4	512 MB	256 Mb

Table 1: Example of Relationship between DDR2 400 Channel Capacities and DRAM Densities

However, for implementations requiring high *per channel* memory capacities DDR2 may be a less expensive alternative than DDR at the point of DDR2 introduction. As mentioned, the increased flexibility (four DIMMs/channel) for DDR2 400 configurations may enable DDR2 400 to meet a given memory requirement in a more cost-effective manner than DDR333.

For example, for a system demanding 4 GB per channel each of the DDR2 400 DIMMs can be single rank 512 Mb based whereas the DDR333 configuration requires one dual rank 512-Mb-based DIMM. Memory pricing of a dual rank 512-Mb-based DIMMs should be at least twice the cost of a single rank 512-Mb-based DIMM.

4 GB per Ch	DIMM 1	DIMM 2	DIMM 3	DIMM 4
DDR2 400	1 GB	1 GB	1 GB	1 GB
DDR333	2 GB	1 GB	1 GB	Not Supported

Similarly, each of the DDR2 400 DIMMs can be single rank 1 Gb based whereas the DDR333 configuration requires a high density dual rank 1-Gb-based DIMMs.

8GB per Ch	DIMM 1	DIMM 2	DIMM 3	DIMM 4
DDR2 400	2 GB	2 GB	2 GB	2 GB
DDR333	4 GB	2 GB	2 GB	Not Supported

Note: Although the picture above is of discrete memory controller implementations, the *per channel* benefits of DDR2 400 over DDR333 still apply regardless if the DDR333 configuration is on an Integrated Memory Controller (IMC).

Other Advantages of DDR2

In addition to the advantages expressed above, Intel's initial DDR2 enterprise platforms now offer increased reliability with ECC memory, x4 Single Device Data Correction (SDDC), DIMM sparing, and DIMM scrubbing.

One example of DDR2-400 flexibility is on die termination (ODT) rather than on the baseboard as with DDR. ODT saves board real estate that may allow for more DIMMs per system for some implementations. For example, for the standard SSI form-factor ODT allows eight DIMMs configurations (versus six DIMMs without ODT).

Summary of DDR2 Advantages

DDR2 400 offers a complete set of advantages for 2P Servers over the older DDR technology. The table summarizes these advantages (vs. DDR333):

DDR2 Features	DDR2 Benefits
DRAM Technology for Multiple Generations	Increased Platform Longevity
Higher Bandwidth	Higher System Performance
Lower Power	Increased Server Density
On Die Termination	Increased Memory Density
Four DIMMs per Channel	Lower cost memory configurations

Comparisons of DDR2 400 vs. DDR400 in Enterprise Platforms

Lower Power

DDR2 400 offers significant advantages in power consumption over DDR400. DDR2 400 SDRAM offers 50 percent less power consumption than DDR400 SDRAM (based on analysis of power and throughput.) This can offer savings of several watts per DIMM lowering the overall memory power budget especially important in dense rack (2U/1U) servers, and blade configurations.

Lower Cost Memory Configurations

DDR2 400 supports four ranks per channel that can be spread over four DIMMs whereas any DDR400 DIMM implementation would likely be limited to no more than two DIMMs.

DDR2 400: Four DIMMs/channel

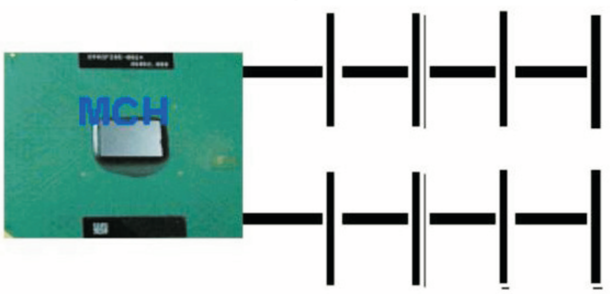

DDR400: Two DIMMs/channel

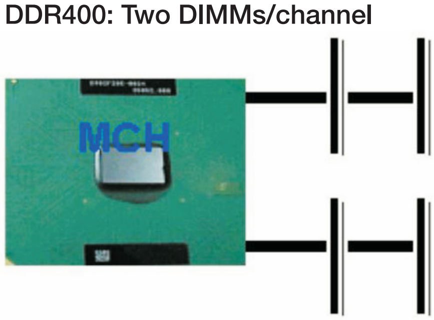

The DDR2 400 increased flexibility may enable DDR2 400 to meet a given memory requirement in a more cost-effective manner than DDR400. For example, for a system demanding 4 GB per channel, each of the DDR2 400 DIMMs can be single rank 512 Mb based whereas a DDR400 configuration requires two dual rank 512-Mb DIMMs. Memory pricing indicators highlight that dual rank 512-Mb-based DIMMs are expected to be more than twice the cost of a single rank 512-Mb-based DIMM.

4 GB per Ch	DIMM 1	DIMM 2	DIMM 3	DIMM 4
DDR2 400	1 GB	1 GB	1 GB	1 GB
DDR400	2 GB	2 GB	Not Supported	Not Supported

Similarly, each of the DDR2 400 DIMMs can be single rank 1 Gb based whereas the DDR400 configuration requires two high-density dual rank 1-Gb-based DIMMs. High-density DIMMs have historically cost more than twice DDR2 DIMMs.

8 GB per Ch	DIMM 1	DIMM 2	DIMM 3	DIMM 4
DDR2 400	2 GB	2 GB	2 GB	2 GB
DDR400	4 GB	4 GB	Not Supported	Not Supported

Note: Although the picture above is of discrete memory controller implementations, the per channel benefits of DDR2 400 over DDR400 still apply regardless if the DDR400 configuration is on an Integrated Memory Controller (IMC). The IMC implementation may compensate by offering more memory channels provided that the second CPU is populated.

Summary of DDR2 Advantages

DDR2 400 offers a complete set of advantages for 2P Servers over the older DDR technology. The table summarizes these advantages (vs. DDR400):

Features	Benefits
DRAM Technology for Multiple Generations	Increased Platform Longevity
Significantly Lower Power	Increased Server Density
On Die Termination	Increased Memory Density
Four DIMMs per Channel	Lower cost memory configurations

For More Information

For more memory information such as webcasts, animations, presentations see the Memory Implementers Forum Web site (www.memforum.org). Membership signup is free and Web based; no commitments to products or marketing required.(>140 companies signed up as of July 28, 2004)

IV—Appendix:

DDR2 400 DIMM configurations (one channel)

DIMM Configurations	DIMM 1	DIMM 2	DIMM 3	DIMM 4
1 Single Rank	Empty	Empty	Empty	Single Rank
1 Dual Rank	Empty	Empty	Empty	Dual Rank
2 Single Rank	Empty	Empty	Single Rank	Single Rank
1 Dual Rank, 1 Single Rank	Empty	Empty	Single Rank	Dual Rank
2 Dual Rank	Empty	Empty	Dual Rank	Dual Rank
3 Single Rank	Empty	Single Rank	Single Rank	Single Rank
1 Dual Rank, 2 Single Rank	Empty	Single Rank	Single Rank	Dual Rank
4 Single Rank	Single Rank	Single Rank	Single Rank	Single Rank

DDR333 configurations (one channel)

DIMM Configurations	DIMM 1	DIMM 2	DIMM 3
1 Single Rank	Empty	Empty	Single Rank
1 Dual Rank	Empty	Empty	Dual Rank
2 Single Rank	Empty	Single Rank	Single Rank
1 Dual Rank, 1 Single Rank	Empty	Single Rank	Dual Rank
2 Dual Rank	Empty	Dual Rank	Dual Rank
3 Single Rank	Single Rank	Single Rank	Single Rank
1 Dual Rank, 2 Single Rank	Single Rank	Single Rank	Dual Rank

Note: DIMM1 is closest to MCH; DIMM4 is farthest from MCH. “Empty” implies that a connector is on the board put not populated.