AMD Opteron™ 6200 Series Processor

What's new in the AMD Opteron™ 6200 Series Processor (Codenamed "Interlagos") and the new "Bulldozer" Microarchitecture

Platform	Processor	Socket	Chipset
Opteron 6000 *(codenamed "Maranello")*	Opteron 6200 *(codenamed "Interlagos")*	G34 *(with BIOS update)*	56x0 / 5100

In their first microarchitecture rebuild since Opteron, AMD introduces the Bulldozer module, implemented in both the Opteron 4200 and 6200 Series. The modules are engineered to emphasize throughput, core count, and parallel server workloads. A module consists of two tightly coupled cores, in which some core resources are shared. The effect is between a core with two threads and a dual-core processor, in which each core is fully independent.

Performance, scalability, and efficiency for today's applications.

Processor Features	Platform Features
▪ 32nm architecture ▪ Up to 16 Bulldozer cores ▪ 2 die per package ▪ Up to 32M combined L2 + L3 cache ▪ Compatible with existing G34 socket w/BIOS update	▪ 4 DDR3 memory channels, LRDIMM, RDIMM, UDIMM up to 1866 GT/s • Low-voltage DRAM 1.2xV to 1.5V ▪ 4 HyperTransportTM technology 3.0 links up to 6.4 GT/s ▪ Advanced Platform Management Link (APML) for system management ▪ AMD-V 2.0 technology

Product Specifications

Process	TDP[1] / ACP[2]	Model	Clock Frequency / Turbo Core Frequency (GHz)	Cores	L2 Cache	L3 Cache	Max Memory Speed	System Bus Speed
32nm	140W / 105W	6282SE	2.6 / 3.3	16	16 x 1MB	16MB	DDR3 1600 MHz	6.4 GT/s
	115W / 80W	6276	2.3 / 3.2	16	8 x 2MB	16MB	DDR3 1600 MHz	6.4 GT/s
		6274	2.2 / 3.1					
		6272	2.1 / 3.0					
		6238	2.6 / 3.2	12	6 x 2MB			
		6234	2.4 / 3.0					
		6220	3.0 / 3.6	8	4 x 2MB			
		6212	2.6 / 3.2					
		6204	3.3 / na	4	2 x 2MB			
	85W / 66W	6262 HE	1.6 / 2.9	16	8 x 2MB	16MB	DDR3 1600 MHz	6.4 GT/s

[1] *TDP stands for Thermal Design Power.*

[2] *ACP stands for Average CPU Power.*

Bulldozer Resource Sharing

The Bulldozer module has shared some components to reduce cost, power, and space (i.e., pack in more cores). The balance of the components remains dedicated to each core.

The OS sees the module as 2 cores.

Shared between the cores:

- Instruction Fetch
- Decode
- L1 instruction cache
- L2 cache
- Two 128-bit FMAC floating-point pipelines

Dedicated to each core:

- Integer Scheduler unit
- L1 data cache
- Load Store unit

Chip Level:

- L3 cache
- HyperTransport links

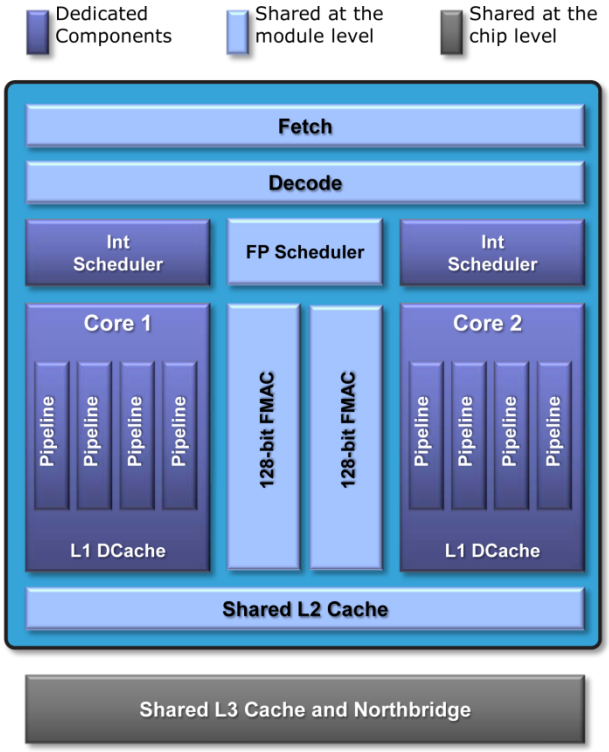

TDP Power Cap

AMD Power Cap Manager (Lisbon / Magny-Cours)

The AMD Power Cap Manager with the Lisbon and Magny-Cours processors cut down power to the processor by locking out the top P-state – also limiting CPU's ability to get to highest frequencies. High frequencies improved processing performance, so implementing Power Cap could force the processors to work longer to get the job done.

Power Capping Power Thresholds (Bulldozer)

The new Power Capping Power Thresholds with Bulldozer will allow the user to set a custom TDP via BIOS or APML. In normal circumstances (running at 40-70% load), the response times will be about the same as with no capping. If your workload does not exceed the new modulated power limit, you can still get top speed because you aren't locking out the top P-state just to reach a power level.

You benefit from:

- Making power consumption predictable
- Flexibility to set power limits without capping CPU frequencies

Bulldozer Flex FP: Flexible 256-bit FPU

The Flex FP unit is built on two 128-bit FMAC units, shared between two integer cores in a module. With each cycle, either core can operate on 256 bits of parallel data via two 128-bit instructions or one 256-bit instruction, OR each of the integer cores can execute 128-bit commands simultaneously. This operation can change with each processor cycle to meet the needs of the moment. In typical data center workloads, floating point operation utilization is typically much lower than integer operations, so if a core has a set of FP commands to be dispatched, there is a high probability that it will have all 256 bits to schedule.

By sharing one 256-bit floating point unit per every 2 cores, the power budget for the processor is held down, and more integer cores can be added within the power budget.

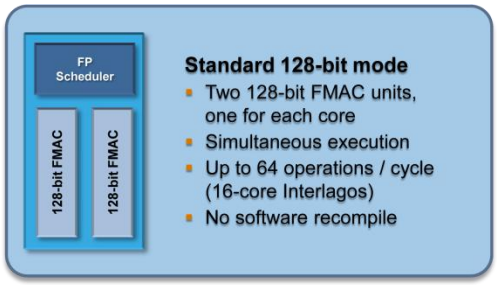

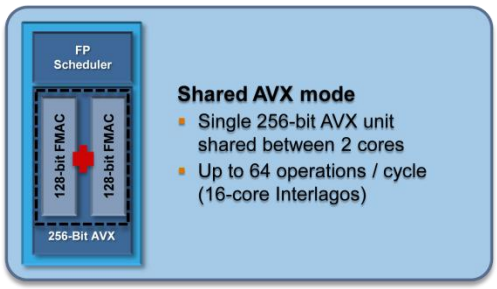

C6 Power State

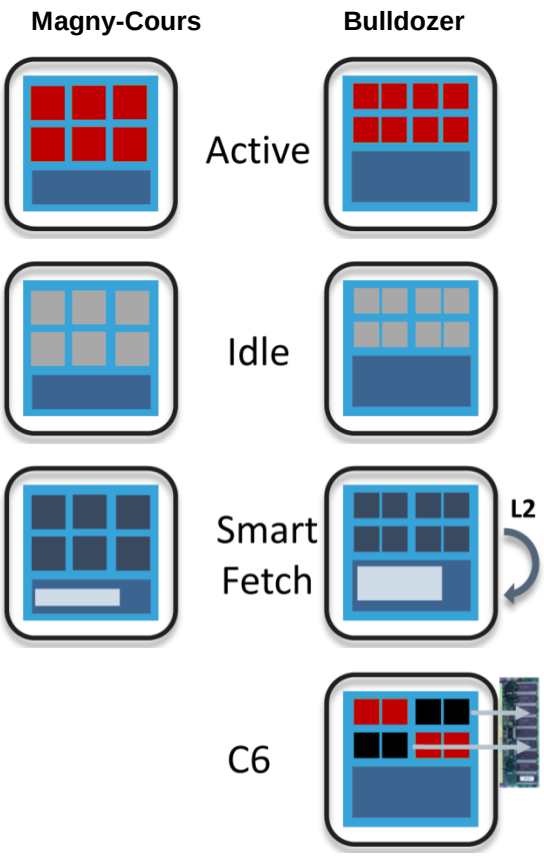

All cores running workloads; core/module frequency can run independently to save power

No cores running workloads; core/module frequency reduced to 800MHz to save more power

After a set idle time, L2 cache is flushed to L3, allowing cores to “sleep” to save power while maintaining MP coherency

In Bulldozer-based processing, any idle module can independently enter C6, completely gating power for considerable savings; module state is saved to DRAM

The new C6 power state reduces processor power consumption at idle by up to 46%.

AMD Turbo Core Technology

Base Frequency with
TDP Headroom

All Core Boost Activated

When there is TDP headroom in a given workload, AMD Turbo Core technology is automatically activated and can increase clock speeds by up to 500 MHz **across all cores.**

Maximum Turbo Activated

When a lightly threaded workload sends half the Bulldozer modules into C6 sleep state but also requests max performance, AMD Turbo Core technology can increase clock speeds by up to 1 GHz across **half the cores.**

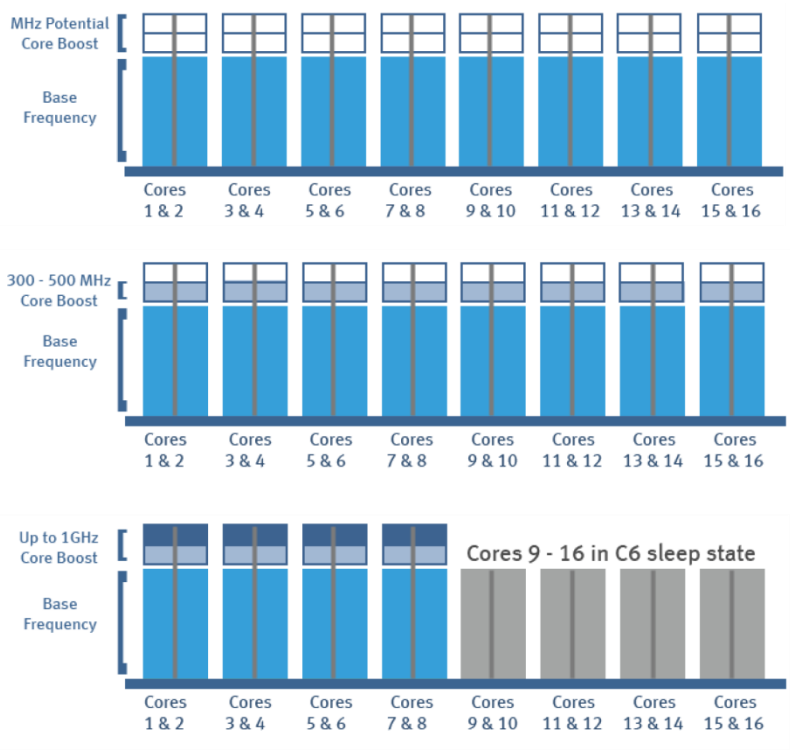

Processor Block Diagram for 1P Mainboards

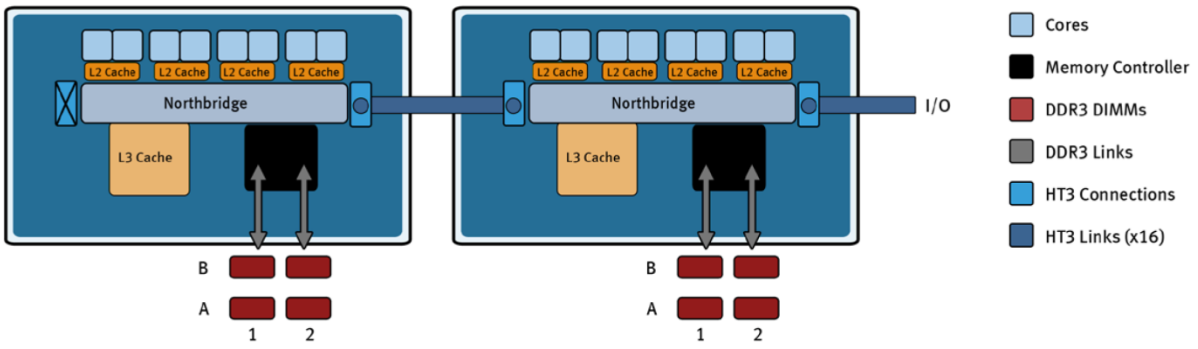

Processor Block Diagram for 2P Mainboards

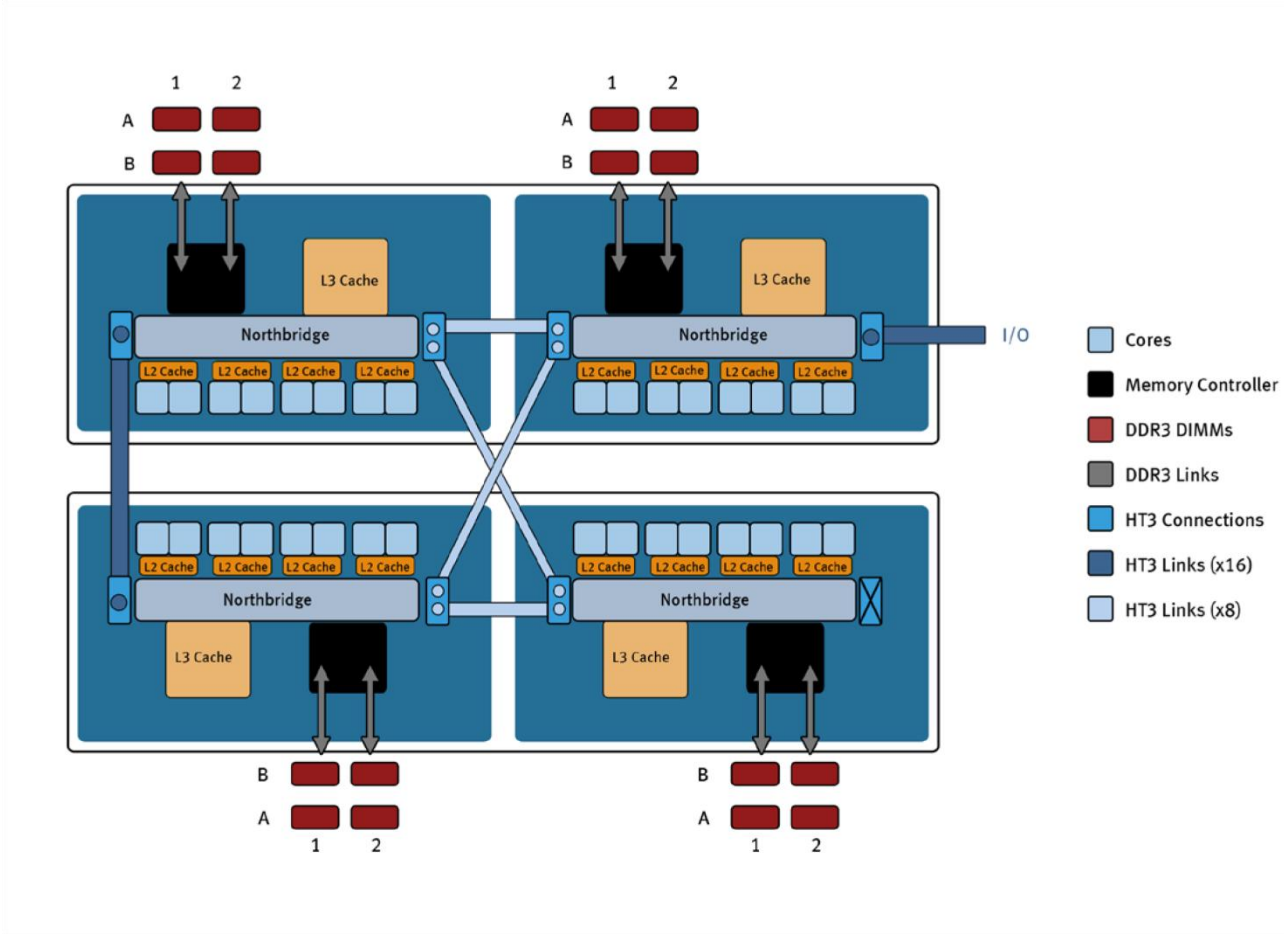

Processor Block Diagram for 4P Mainboards

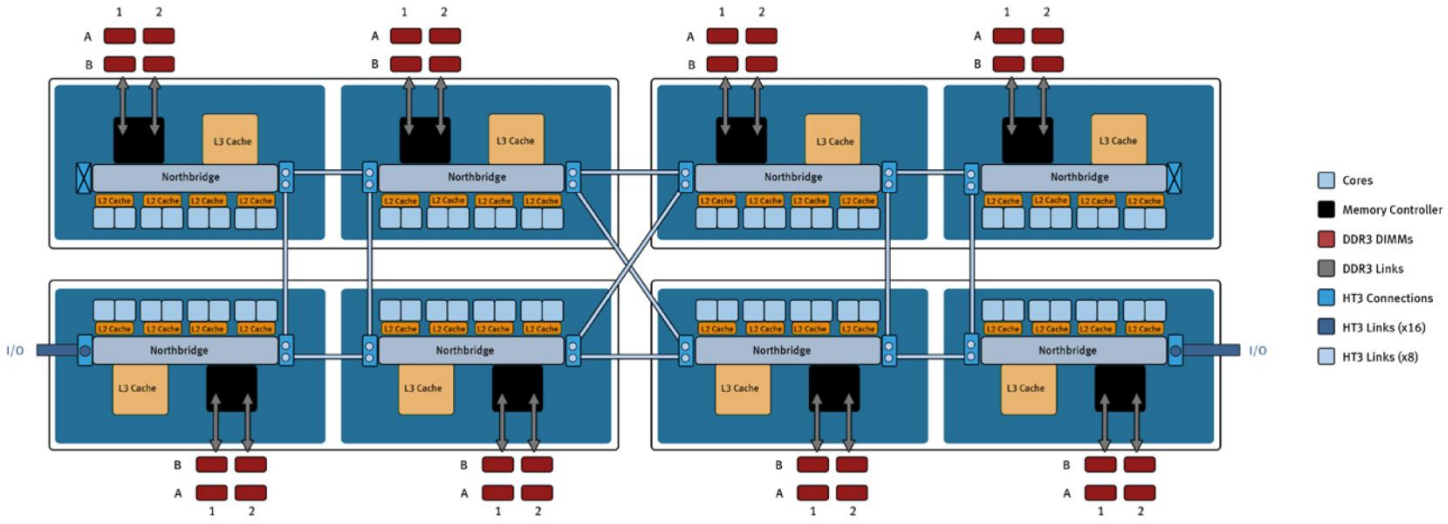

HT3 Connections for 2P Mainboards

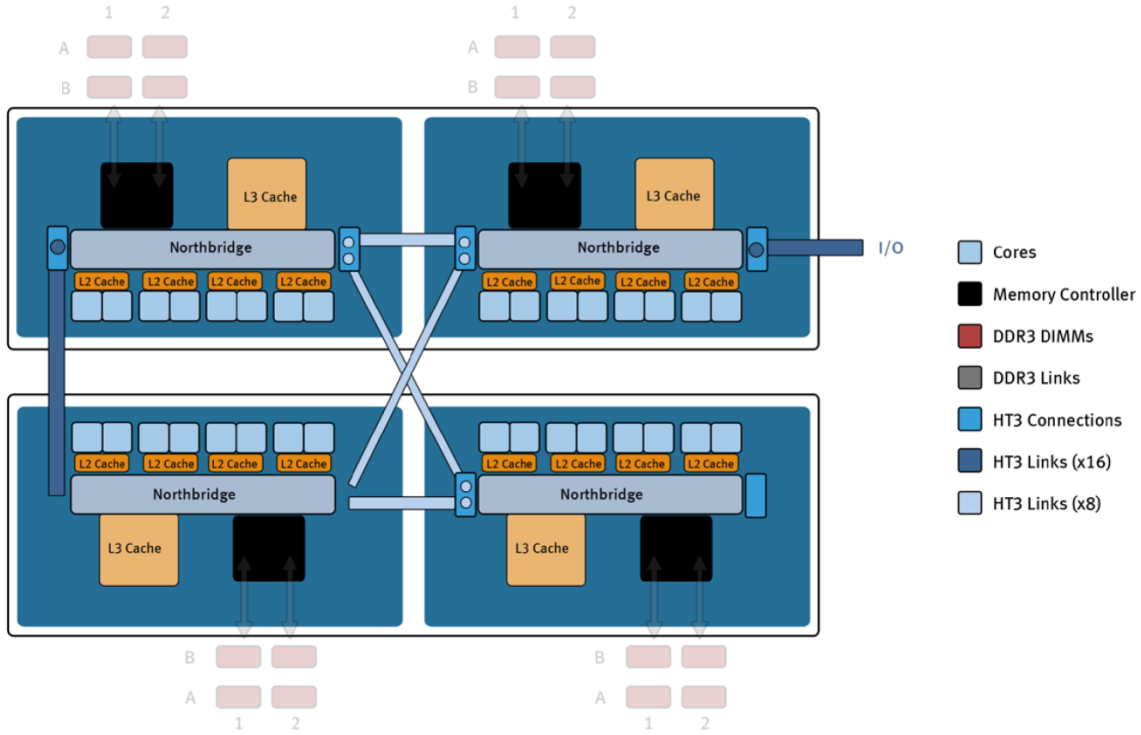

HT3 Connections for 4P Mainboards

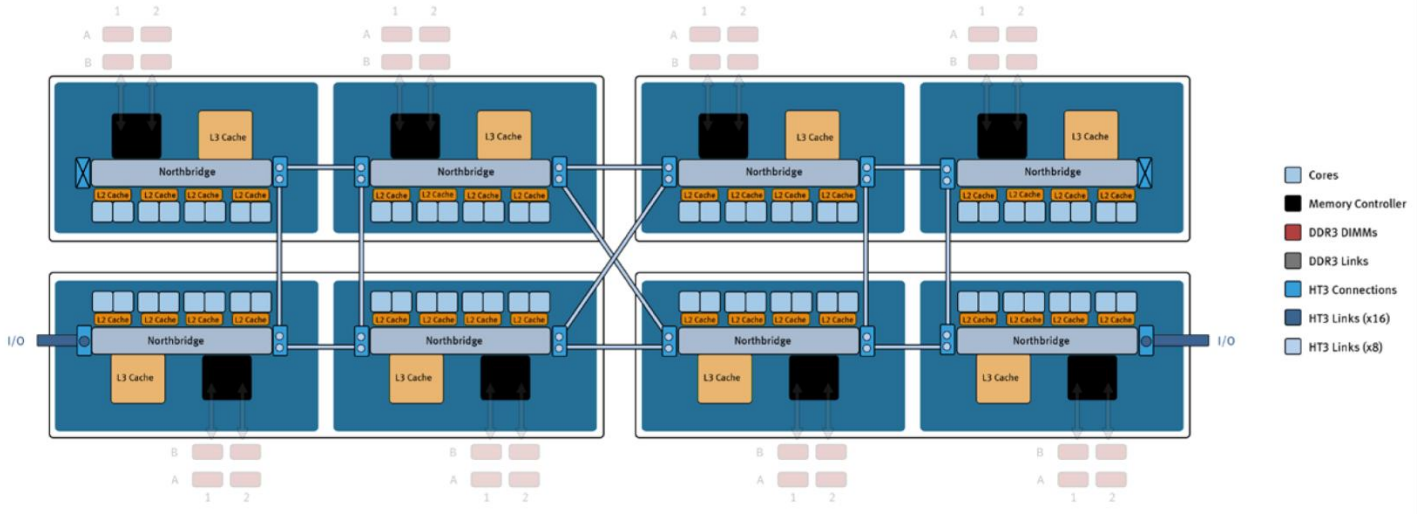

Memory Population Guidelines: 1 DIMM per Channel

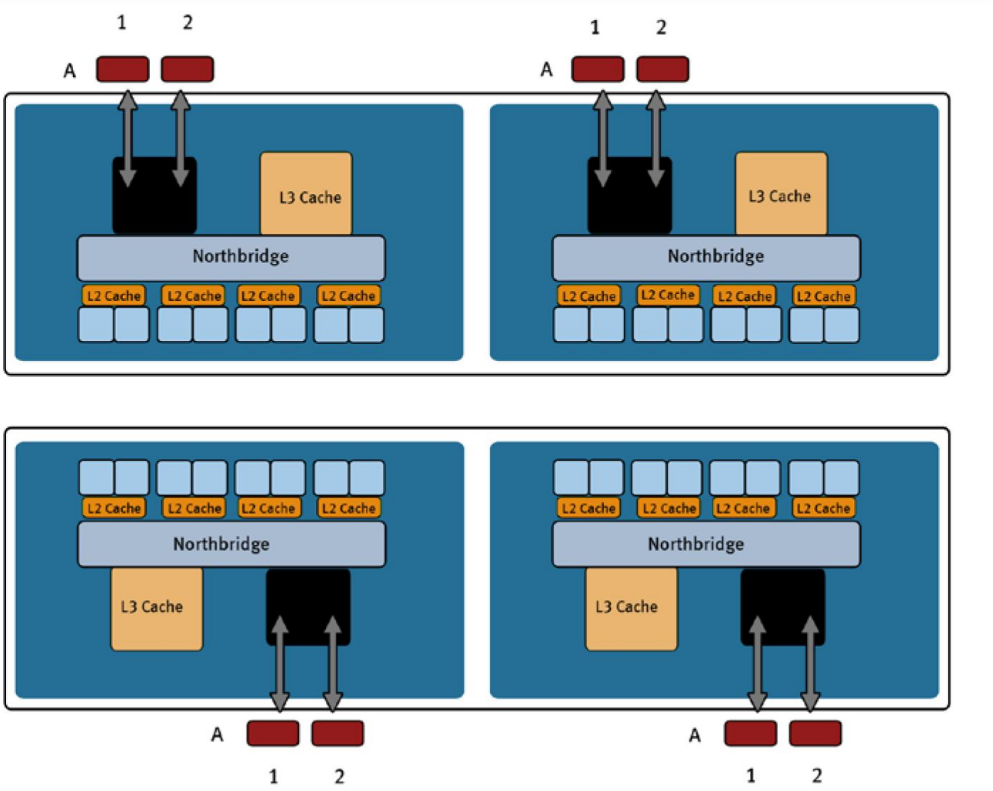

Mainboards with 1 DIMM per Channel

	DIMM Bank A	Max MHz, 1.5V DIMMs	Max MHz, 1.35V DIMMs	Max GB/Channel
UDIMM	1R or 2R	1600	1600	4GB
RDIMM	1R or 2R	1600	1600	16GB
	4R	1333	1333	16GB

Notes:

AMD supports single, dual, and quad rank (1R, 2R, and 4R). There is no distinction between memory population scenarios for 1P, 2P, and 4P mainboards.

UDIMM: Unbuffered DIMM

RDIMM: Registered DIMM

Memory Population Guidelines: 2 DIMMs per Channel

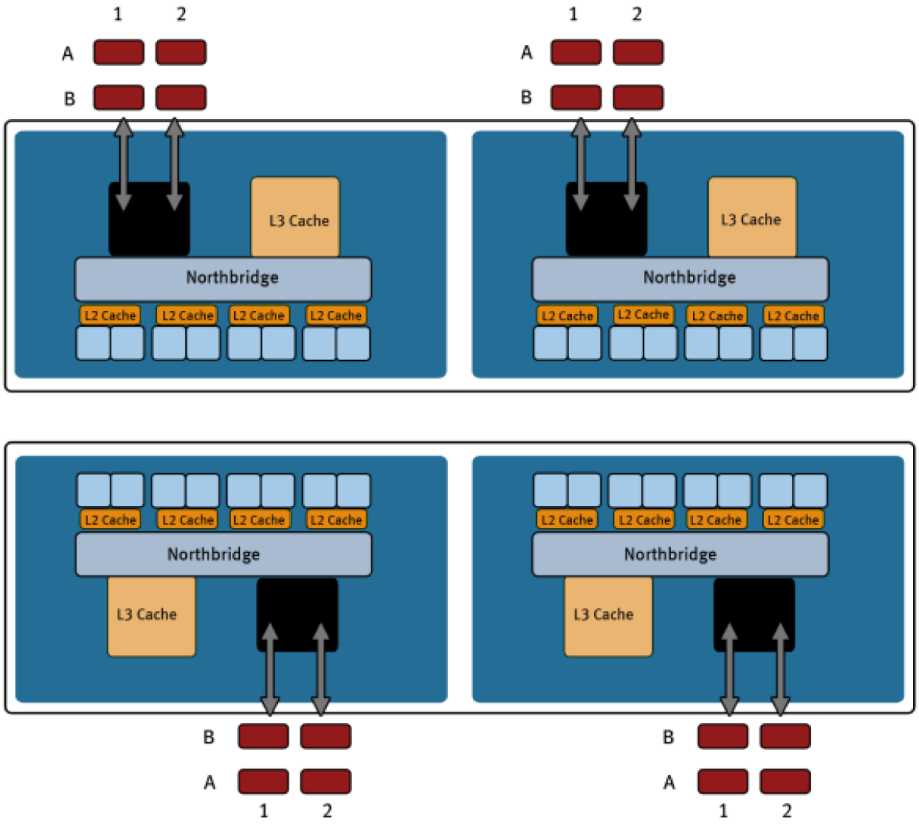

Mainboards with 2 DIMMs per Channel

	DIMM Bank A	DIMM Bank B	Max MHz, 1.5V DIMMs	Max MHz, 1.35V DIMMs	Max GB/Channel
UDIMM	1R or 2R	Empty	1600	1333	8GB
	1R	1R	1600		
	2R	2R	1333		
RDIMM	1R or 2R	Empty	1600	1333	16GB
	1R	1R			
	2R	2R			
	4R	Empty	1333	1066	16GB
	4R	4R	1066	800	32GB

Notes:
AMD supports single, dual, and quad rank (1R, 2R, and 4R). There is no distinction between memory population scenarios for 1P, 2P, and 4P mainboards.
UDIMM: Unbuffered DIMM
RDIMM: Registered DIMM

Memory Population Guidelines: 3 DIMMs per Channel

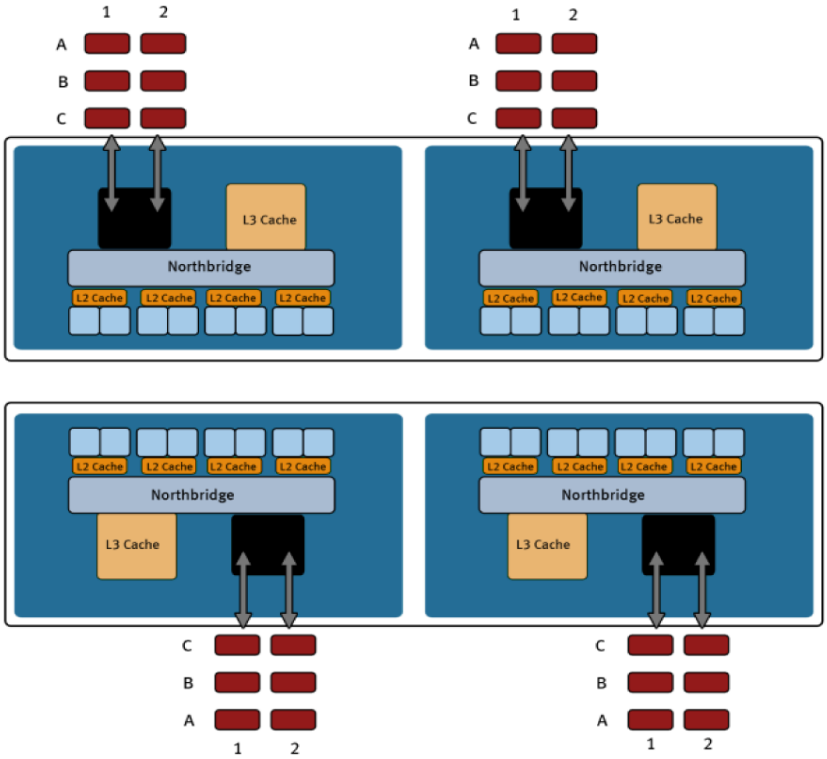

Mainboards with 3 DIMMs per Channel

	DIMM Bank A	DIMM Bank B	DIMM Bank C	Max MHz, 1.5V DIMMs	Max MHz, 1.35V DIMMs	Max GB/Channel
UDIMM	1R or 2R	Empty	Empty	1600	1333	8GB
	1R	Empty	1R	1333		
	2R	Empty	2R			
RDIMM	1R or 2R	Empty	Empty	1600	1333	16GB
	1R	Empty	1R	1333		
	2R	Empty	2R			
	1R	1R	1R	1066	1066	12GB
	Empty	4R	Empty			16GB
	2R	2R	2R		800	24GB
	1R or 2R	4R	Empty	800		32GB
	1R or 2R	4R	1R or 2R			

Notes: AMD supports single, dual, and quad rank (1R, 2R, and 4R). There is no distinction between memory population scenarios for 1P, 2P, and 4P mainboards.

UDIMM: Unbuffered DIMM

RDIMM: Registered DIMM

Rackmount Servers

1U Servers

- AR700 *
- AF700 *
- MM1640 **
- MX1630 ***

2U Servers

- MM2680 **
- MX2660 ***

3U Servers

- MM3680 **

4U Servers

- MM4680 **
- MM4690 **
- MX4680 ***

* 1P Motherboard
** 2P Motherboard
*** 4P Motherboard

Contact Ironsystems Sales

For answers regarding processor selection,
memory matching, or other questions you may have,
contactSales@ironsystems.com Toll Free: 800-943-IRON(4766)

For more infornation visit: **http://www.ironsystems.com**

Iron Systems, Inc.
540 Dado Street, San Jose, CA 95131, USA
Phone: 1-408-943-8000 (Local) 1-800-921-IRON
Fax: 1-408-943-8222
Email: *info@ironsystems.com*
Website: *http//www.ironsystems.com*