

Netlist™ Inc.

Low-Power Fully-Buffered DIMM

White Paper

Revision 1.01





Overview

Fully-Buffered DIMM (FBDIMM) based server systems were introduced in 2006 into the marketplace. FBDIMM architecture presents significant increase in memory bandwidth utilizing a high-speed serial architecture. However, FBDIMM modules face significant thermal management challenges due to the higher power consuming components in the DIMM, primarily the Advanced Memory Buffer (AMB). This poses additional thermal management burden for form-factor constrained systems like the blade and 1U servers. Thermal challenges therefore limited the overall system performance and capacity.

At high temperatures, electronic components reach their physical limits and fail to perform as desired, which results in degraded system reliability. Many systems throttle requests from the memory controller thereby reducing the power dissipated and eventually cooling down the electronics. However, this has a performance impact. In addition, the system's cooling mechanisms (fans, fluid-based, etc.) are boosted in an effort to lower the temperature, which results in louder acoustic conditions.

Why Low Power?

The primary motivation for Netlist's effort in designing low power FBDIMM modules is to mitigate the above challenges presented by FBDIMM. The low power FBDIMM modules potentially enable systems to be configured with more memory capacity, yield higher performance and enable improved acoustic conditions.

Another significant motivation for this effort is the increasing significance of the Energy-Star guidelines from Environmental Protection Agency (EPA) as applicable to Workstations and Servers.

Identifying Power

Primarily the active components of the FBDIMM include the Dual-Data Rate (DDR) DRAM and AMB. Therefore, in order to save power, these components and their usage need to be optimized. An end user only benefits from power savings achieved at a system level. FBDIMMs are vital components of all server and workstation systems. As the individual components save power, so will the system.

The industry standard FBDIMM specification quantifies active and idle current consumed by the AMB based on certain modes of operation. It is challenging to identify a standard methodology for calculating the average power dissipation due to the variety of factors that affect it, including the application, system configuration, and the DIMM slot within a given channel. For the purpose of comparison, the peak operation is identified and extrapolated. For a given FBDIMM, the absolute maximum power consumed is estimated by using the AMB's IDD_Active_1 value, and the DRAM's active and standby currents (IDD7 and IDD2N respectively). For definitions of these parameters, please refer to the JEDEC standard or vendor specification for the AMB and DRAM.

Most commonly available 2GB FBDIMMs currently in the market are of a 2-Rank x4 configuration based on JEDEC design. The Netlist™ Low-Power FBDIMM solutions offer a 4-Rank x8 configuration. **Table 1** compares the power consumption during peak operation of an active module (Note: DRAM and AMB data indicate typical values for DDR2-667). As it is evident from the data, the Netlist™ products are more energy efficient as compared to the 2-Rank modules.

Savings through DRAM

The most simplistic approach to achieving power efficiency on the



NETLIST

FBDIMM is to reduce the number of active components present on the module. Until recently, the chipset and AMB components only supported a maximum of two Ranks on a FBDIMM. However, as memory requirements increase, 4-rank support is available on the latest generation of AMB devices. This enables the use of lower density, x8 DRAM components for a cost-effective solution. The 4-rank configuration has only 25% of the DRAM components in active state (vs. 50% in a 2-Rank configuration). A module can benefit from power savings as long as the total active current of the DRAM components is reduced. Taking the example of a 2GB 2Rx4 FBDIMM, this module is assembled using thirty-six 512Mb DDR2 DRAM. Eighteen devices are active during operation along with an AMB. The remaining eighteen DRAM components are powered in standby mode.

The Netlist™ Low-Power FBDIMM alternative offers FBDIMM modules in a 4-Rank, x8 configuration. Although the module contains the same number of DRAM components, thirty-six, the 4-Rank module will only have nine DRAM chips in the active mode of operation. This reduces the DRAM active power consumption in half. The remaining twenty-seven DRAM devices operate in the standby mode. Depending on the

DRAM vendor and die-revision, the ratio of an active to idle current range is 3-7. Using an average value of 4, this translates into a 30% power savings from DRAM on a 4-Rank module. To evaluate a system-level impact, one must also consider the inactive modules in a system. Similar to the active power calculation, **Table 2** shows the idle power comparison of the FBDIMM modules.

Savings through the AMB

The largest consumer of power in the memory module is the AMB. The prior generation of AMB consumes 6-7 watts of power. Several prominent AMB vendors have optimized the designs for power through enhancements in the interface. The AMB draws power from two domains, V_{DD} (1.8V) for the IO and V_{CC} (1.5V), for the core. The current generation of the AMB devices has optimized the core and I/O design.

Implications of Energy Savings

By implementing Low-Power solutions in the FBDIMM based server systems, the direct benefit is in the energy savings. The benefits of these savings however may be greater than one may imagine.

Table 1 - 2GB FBDIMM Active Power Calculation

Configuration	Component	Current (mA)		Number of Devices	Total Current (mA)	Operating Voltage (V)	Total Power (W)	Module Power (W)
2GB 2Rx4	AMB	ICC_Active_1	2900	1	2900	1.6	4.64	14.42
		IDD_Active_1	1275	1	1275	1.9	2.42	
	DRAM	IDD7	180	18	3240	1.9	6.16	
		IDD2N	35	18	630	1.9	1.20	
2GB 4Rx8	AMB	ICC_Active_1	2370	1	2370	1.6	3.79	8.85
		IDD_Active_1	700	1	700	1.9	1.33	
	DRAM	IDD7	125	9	1125	1.9	2.14	
		IDD2N	31	27	837	1.9	1.59	

Power Savings with 4Rx8 over 2Rx4: 38.61%

Note: Component specifications as indicated by the vendor datasheets; vendors remain confidential under NDA



Table 2 - 2GB FBDIMM Idle Power Calculation

Configuration	Component	Current (mA)	Number of Devices	Total Current (mA)	Operating Voltage (V)	Total Power (W)	Module Power (W)
2GB 2Rx4	AMB	ICC Idle 1	2550	1	2550	1.6	4.08
		IDD Idle 1	675	1	675	1.9	1.28
	DRAM	IDD2Q	18	36	648	1.9	1.23
2GB 4Rx8	AMB	ICC Idle 1	2370	1	2370	1.6	3.79
		IDD Idle 1	130	1	130	1.9	0.25
	DRAM	IDD2Q	23	36	828	1.9	1.57

Power Savings with 4Rx8 over 2Rx4: 14.89%

Note: Component specifications as indicated by the vendor datasheets; vendors remain confidential under NDA

Energy Cost Reduction

Aside from the above stated explanation of performance benefits through low-power options, the server system becomes more attractive through cheaper operating costs. As mentioned earlier, the low-power FBDIMMs run cooler, which in turn reduces the whole system's cooling costs. In an environment that contains a large number of server systems in confined storage space, the ambient temperature needs to be monitored closely. By high level yet conservative standards, it is generalized that a watt of energy is spent to cool the heat generated by a watt of dissipated power. Following this premise, the energy savings by using Netlist™ Low-Power FBDIMMs has twice the impact at a system level. Using the 2GB FBDIMM example, **Table 3** presents the analysis of the total power savings on a typical 4-Channel, 4-Rank system populated with eight DIMMs.

Energy Star Initiative

The future of electronics and computing will be heavily impacted by the conscious efforts of the global attention driven towards energy conservation. The above stated EPA program of ENERGY STAR and other industry initiatives is expected to play a major role in shaping the design of future products. Specifically, the ENERGY STAR initiative introduces stringent standards requirements for computing systems. "Ratio Typical Energy Consumption" (RTEC) is a new metric developed for scalable workstations and servers.

This efficiency index helps measure the value or production of a device relative to the energy used in the total time of operation. Rather than limiting power consumption in certain power modes, the new approach quantifies the overall system behavior over the operational lifetime.

Table 3 - System-level impact of Low-Power FBDIMMs

Configuration	Power Component	Power per Module	Number of Modules	Total Power (W)	System Power (W)
2GB 2Rx4	Active	14.42	4	57.66	84.04
	Idle	6.59	4	26.37	
2GB 4Rx8	Active	8.85	4	35.40	57.85
	Idle	5.61	4	22.45	

Power Savings with 4Rx8 over 2Rx4: 26.19 or 31.16%

Annual Energy Cost Savings (at 10¢/kWh) \$45.76

Note: Assumed savings of a watt in cooling cost savings for every watt of power savings



NETLIST

Although usage models vary widely, many systems are majority of the time, in the standby or idle states. The energy consumption is much greater in the active state, but over time, the energy consumed by the idle states collectively, may outweigh that of the active state. The RTEC measure helps highlight these systems allowing designers to focus on areas of greater power consumption that produce little output.

The Netlist Edge ...

Using the Netlist™ Low-Power FBDIMM solutions enable the system designers to take a step towards designing systems that are power, cost and performance optimized. Netlist™ delivers low-power designs using industry-leading technologies, optimized component usage, superior thermal solutions and high quality manufacturing, at reduced costs.

Netlist FBDIMM Offerings:

Part. No.	Capacity	Speed	Performance	Refresh Rate*	CL
Quad-Rank x8 DDR2 Fully-Buffered DIMM					
NMD517A21207FD64	4GB	800 MT/s	PC2-6400	7.8µs	5
NMD517A21207FD53	4GB	667 MT/s	PC2-5300	7.8µs	5
NMD257A26407FD64	2GB	800 MT/s	PC2-6400	7.8µs	5
NMD257A26407FD53	2GB	667 MT/s	PC2-5300	7.8µs	5
Quad-Rank x4 DDR2 Fully-Buffered DIMM					
TBD [^]	8GB	667 MT/s	PC2-5300	7.8µs	5
Dual-Rank x8 DDR2 Fully-Buffered DIMM					
NMD257A21207FD64	2GB	800 MT/s	PC2-6400	7.8µs	5
NMD257A21207FD53	2GB	667 MT/s	PC2-5300	7.8µs	5

Notes: * Refresh rate – 7.8µs for Tcase < 85°C; 3.9 µs for Tcase 85°C - 95°C

[^] Product under development, check for availability

For additional information or orders: Contact (949) 435-0025 or sales@netlist.com

Corporate Headquarters:

Netlist, Inc.
51 Discovery, Suite 150
Irvine, CA – 92618
Tel. (949) 435-0025
Fax. (949) 435-0031
www.netlist.com

Netlist Inc., 2007

All other names used are trademarks of their respective owners.

About Netlist, Inc.

Netlist designs and manufactures high-performance memory subsystems for the server and high-performance computing and communications markets. The Company's memory subsystems are developed for applications in which high speed, high-capacity memory, functionality, small form factor, and heat dissipation are key requirements. These applications include tower servers, rack-mounted servers, blade servers, high performance computing clusters, engineering workstations, and telecommunication switches. Netlist maintains its headquarters and manufacturing facilities in Irvine, California, and in Suzhou, China.