



Real World Analysis on DDR Cache for the Motorola 7450 Family

A White Paper Regarding Design Choices for L3 Cache Implementation in Processor Upgrade Cards for the Apple PowerMac G4

This white paper is not meant to be an in-depth, highly technical treatise, although there are some technical sections; rather it is written to educate the layman and consumer with regards to performance options of their Apple PowerMac G4 computer system. Is DDR performance real, or is it just marketing hype?

BACKGROUND

Single Data-Rate

For the last few years, the memory technology found in virtually all consumer desktop computers has been 'SDR.' SDR stands for Single Data-Rate although that term was not commonly used until later technology (DDR) became available. SDR architecture means that for every tick of the main memory bus clock, a **single** datum (word) is transferred between the memory and the central processing unit (CPU). Like everything else in the computer world, the main memory bus clock speeds have steadily evolved to the point where it was time for a somewhat more radical approach. Enter DDR.

Double Data-Rate

In the recent past desktop computer systems incorporating DDR technology have entered the market. As its name implies, Double Data-Rate (DDR) topology doubles the information transferred per unit time. Simply put, for every tick of the main memory bus clock, **two** words of data are transferred.

*A "clock" signal is composed of rising edges and falling edge, much like the teeth of a saw blade. SDR technology clocks (transfers) data only on the rising edge of this wave, whereas DDR technology clocks data on **both** edges.*

Although utilizing DDR type of main memory over SDR brings theoretical improvement to any systems performance, there are mitigating factors. Just because some part of the system can perform twice as fast does not necessarily mean that overall throughput is affected the same way. Many factors influence how much of a benefit is realized at the system level – system architecture, system controller design, what type of application, what operating system, etc. There are many other technical reasons which are beyond the scope of this paper.

Cache Memory

Caches are small, very high-speed memory arrays used for temporarily storing the most recently used instructions and data. The CPU spends a very high percentage of the time using the same instructions and data over and over again. Typically, the CPU will find the data it needs a large percentage of the time in one of the caches, thus they are extremely effective at raising performance levels.

Most current CPUs have a number of caches. On the latest PowerMacs, there are usually 3 levels of cache. These are located internally on the CPU chip itself, as well as externally. Only the CPU manufacturers themselves have control over the design and the performance of the internal caches. However, external, or 'backside' cache design and performance have some variation.

Caches are referred to by their different "levels." It is not feasible to make a single, large ultra-fast cache, so designers divide up this function. The levels are differentiated by their comparative size, speed, and ordinal position on the system's bus. The smallest, fastest cache is termed 'Level 1.' The 'Level 2' cache is sometimes slower (but not in every case), but is larger than L1. (The L2 cache on the G4/7455 is 256K in size, running at processor speed.) Any other caches are labeled progressively in the same manner. They all work together creating an extremely efficient mechanism for enhancing system throughput.

Apple

The first appearance of DDR technology in Apple computers was not in main memory, but in the form of the backside or 'L3' cache interfaced to the Motorola PowerPC G4 7450 CPU. The early Apple processor design utilized 1MB of DDR cache (the PowerMac G4 733 using the Digital Audio motherboard.) Other designs use 2MB of DDR cache for L3, although some of the recently introduced Apple offerings have moved back down to 1MB of L3 DDR.

DDR technology for the main memory has only very recently become available in the latest offerings from Apple. Earlier this month (August 2002), Apple has announced [several new desktop systems](#) with DDR technology for the main memory, as well as DDR type backside cache. The Apple [Xserve](#) also sports DDR main memory.

However, much confusion exists regarding the importance of DDR technology in the current Macintosh platforms. The early [benchmarks](#) from independent sources indicate that the use of DDR SDRAM in the Macintosh platform provides virtually no performance increase at all! This is good news for Macintosh users that want to preserve their investment in previous generation G4 systems, as those systems are not as far behind in technology as it may appear.

So, the question is, if there is virtually no difference in performance between DDR main memory and SDR, what about differences between DDR L3 and SDR L3 caches?

L3 Cache Design: SDR vs DDR

At first, it seems like implementing a processor card design with DDR technology for the L3 cache is a 'no-brainer,' as DDR appears to be twice as fast as SDR. Not only that, it would appear that to be able to effectively compare performance-wise against the latest models from Apple, it would almost be a requirement to use DDR in order to overcome some of the inherent design limitations of the older PowerMac G4s (bus speeds of 100 MHz instead of 133 for example; not to mention the perceived marketing value of the term DDR, which by its very name implies a doubling of performance.)

As one might expect, DDR is more expensive technology to implement than SDR—on the order of four times the device cost of SDR in the current marketplace (perhaps less of a cost gap for a company with the relationships and buying power of Apple.) So, DDR certainly must offer quite a substantial benefit before it would make sense to implement. PowerLogix has expended considerable effort designing the latest PowerForce G4 Series 100 and 133 cards and this is one of the major issues we faced: would DDR performance be commensurate with the increase in the product cost? Upon closer analysis of the 7450 architecture and inherent compromises of specific DDR implementations, as well as thorough testing and benchmarking, we made some interesting discoveries.

ANALYSIS

L1 and L2

The L1 and L2 caches are very effective themselves, doing a large share of the work supplying data to the processor. For the majority of accesses, the required data simply is supplied from L1 or L2 caches directly. (Typically, a cache 'hit rate' of L1+L2 approaches 90-98%; see reference below.)

Main Memory

In any instance when data is not in any of the caches (i.e., a 'cache miss') the processor must access main memory. Main memory is always **much** slower than cache memory and has very **large** latency. Because of this huge disparity in performance between fetching data from main memory and getting it from the L3, the actual data throughput (no matter how fast the L3 may be) is substantially diminished. To put it another way, it almost doesn't matter how fast the L3 is as long as it exists. To confirm this hypothesis, we performed tests which compare two speeds of cache of the same type; for example, 167 MHz SDR vs 333 MHz SDR.

The 7450 Cache Interface

The design of the Motorola 7450 Family's L3 cache interface transfers data to/from the L3 cache in short bursts of four words. When an access begins, there is a small, initial delay ('latency'), then the actual burst of data is transferred in four separate, back-to-back 'beats.' For SDR cache, these 'beats' are timed with the rising edge of the L3 clock. For DDR cache, these beats are timed on both rising and falling edge of the L3 clock, thus the

data arrives twice as fast at the L3 interface. At first glance, this would seem to indicate 2x faster data throughput. But what is the real potential throughput gain?

Theoretical Gain

To determine this, we compared two existing cards that are otherwise identical except for the L3 cache. A production Apple CPU card* (G4/7455 @ 1GHz, with DDR cache clocked at 250MHz) is observed to have an L3 cache latency of 5 L3 clocks or **20ns**, and total transfer time for all four data beats of 6.5 L3 clocks, or **26ns**.

A PowerLogix G4 Series 100/133 card (G4/7455 @ 1GHz, with SDR cache clocked at 250MHz) is observed to have an L3 cache latency of 4.5 clocks or **18ns**, and total transfer time for all four data beats of 7.5 L3 clocks, or **30ns**. *This equates to a 15% **theoretical speed increase** under entirely artificial conditions* (i.e., 100% cache 'miss' in L1&L2, and 100% cache 'hit' in L3.) Note this 15% theoretical figure occurs **only** in the condition where main memory **never** is accessed, which is clearly something that will never happen in the real-world.

When either of these cards is required to get the data from main memory (i.e, a complete L1/L2/L3 cache 'miss'), it requires a minimum of **90ns** latency, and **120ns** for total transfer time for all four beats. So, in comparison to main memory access time, the net 4ns 'time advantage' of DDR over SDR for a transfer of four data beats is a whopping 3.3% (4ns/120ns .) In addition, in terms of latency, SDR has 10% less latency or faster access time. This is especially important for the most critically needed *first data-beat*. This means under some special conditions SDR can be faster than DDR.

In addition, with DDR cache, the 7450's L3 interface *does not forward (use) the first beat's data until the second beat's data arrives*. It also *does not use the third beat's data until the fourth beat's data arrives*. This negates some of the advantage of the DDR by not forwarding data to the processor as fast as it is actually being delivered by the DDR devices. When in SDR mode, data is forwarded with no further delay.

One other point of comparison: identical cache type, with a variation only in speed. For example, compare an SDR L3 cache at 167MHz to an SDR L3 cache at 333MHz. This is a *permanent, fixed-under-all-conditions* 2x multiplication in L3 cache speed, yet the benchmarks show typically less than a 3% differential. So, clearly, even if DDR were also fixed at 2x over SDR (all other factors being equal), the net throughput results could not result in a significant performance increase.

Combining all of these factors, and comparing real-world results, we find that although some data can be delivered faster from a DDR L3 cache than SDR L3, in practice, the DDR advantage is practically non-existent.

Careful analysis of the design and operation of the of the 7450 L3 cache interface shows this, but most convincingly, it is evidenced by the benchmarks and real-world applications.

Benchmarks

Benchmarks below include a multiple action Photoshop macro (provided by Barefeats.com), and three [Cinebench 2000](#) benchmarks. This is by no means an exhaustive list of benchmarks. However, considering the results of the speed-doubling 167MHz SDR vs 333MHz SDR test, it stands to reason that it is highly unlikely that other benchmarks will show any practical speed difference between DDR and SDR, considering that DDR is theoretically (under ideal conditions) at most only 15% faster than SDR.

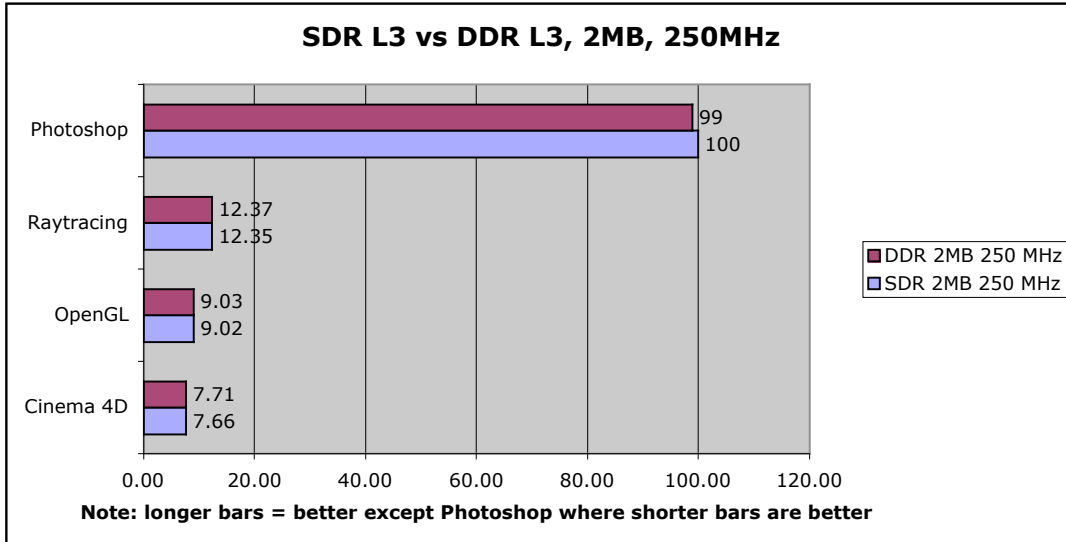


Figure One. Comparison of SDR and DDR L3 caches running at the same speed. Note virtually identical performance.




Figure Two. Comparison of the fastest DDR cache vs the slowest SDR shows only 2-3% difference. Note this difference is almost the same as the fastest SDR vs slowest SDR.




Figure Three. Comparison of two SDR caches with one running fully twice the speed of the other results in minimal performance gain (2.9% in PhotoShop, 1% in the others.)

		Cinebench			
		Shading Cinema 4D	Shading OpenGL	Raytracing	Photoshop
SDR 1MB	333 MHz	7.24	8.78	12.19	101
SDR 1MB	286 MHz	7.23	8.74	12.16	101
SDR 1MB	250 MHz	7.22	8.73	12.14	103
SDR 1MB	167 Mhz	7.16	8.68	12.06	104
SDR 2MB	333 MHz	7.73	9.10	12.40	99
SDR 2MB	286 MHz	7.65	9.07	12.37	100
SDR 2MB	250 MHz	7.66	9.02	12.35	100
SDR 2MB	167 Mhz	7.60	8.94	12.19	102
DDR 1MB	333 MHz	7.29	8.81	12.19	102
DDR 1MB	286 MHz	7.22	8.74	12.16	102
DDR 1MB	250 MHz	7.21	8.73	12.14	102
DDR 1MB	167 Mhz	7.19	8.67	12.09	103
DDR 2MB	333 MHz	7.77	9.12	12.42	99
DDR 2MB	286 MHz	7.73	9.06	12.39	99
DDR 2MB	250 MHz	7.71	9.03	12.37	99
DDR 2MB	167 Mhz	7.62	9.00	12.32	100

Figure Four. Raw benchmark data. (Note: Cinebench results use Maxon's arbitrary benchmark unit called a 'CB.' The PhotoShop benchmark units are in seconds.)

Percentage difference between benchmarks (absolute value of difference)		Cinebench			
		Shading Cinema 4D	Shading OpenGL	Raytracing	Photoshop
1MB DDR vs 1MB SDR	333MHz	0.69%	0.34%	0.00%	0.99%
	286MHz	0.14%	0.00%	0.00%	0.99%
	250MHz	0.14%	0.00%	0.00%	0.97%
	167MHz	0.42%	0.12%	0.25%	0.96%
2MB DDR vs 2MB SDR	333MHz	0.52%	0.22%	0.16%	0.00%
	286MHz	1.05%	0.11%	0.16%	1.00%
	250MHz	0.65%	0.11%	0.16%	1.00%
	167MHz	0.26%	0.67%	1.07%	1.96%
2MB DDR vs 1MB DDR	333MHz	6.58%	3.52%	1.89%	2.94%
	286MHz	7.06%	3.66%	1.89%	2.94%
	250MHz	6.93%	3.44%	1.89%	2.94%
	167MHz	5.98%	3.81%	1.90%	2.91%
2MB DDR vs 1MB SDR	333MHz	7.32%	3.87%	1.89%	1.98%
	286MHz	6.92%	3.66%	1.89%	1.98%
	250MHz	6.79%	3.44%	1.89%	3.88%
	167MHz	6.42%	3.69%	2.16%	3.85%
2MB SDR vs 1MB SDR	333MHz	6.77%	3.64%	1.72%	1.98%
	286MHz	5.81%	3.78%	1.73%	0.99%
	250MHz	6.09%	3.32%	1.73%	2.91%
	167MHz	6.15%	3.00%	1.08%	1.92%
167MHz vs 333MHz	DDR	1.97%	1.33%	0.81%	1.00%
	SDR	1.68%	1.76%	1.69%	3.03%
167MHz vs 286MHz	DDR	1.42%	0.66%	0.56%	1.01%
	SDR	0.65%	1.43%	1.46%	2.00%
167MHz vs 250MHz	DDR	1.18%	0.33%	0.41%	1.00%
	SDR	0.78%	0.89%	1.30%	2.00%
250MHz vs 333MHz	DDR	0.77%	0.99%	0.40%	0.00%
	SDR	0.91%	0.88%	0.40%	1.01%
250MHz vs 286MHz	DDR	0.26%	0.33%	0.16%	0.00%
	SDR	0.13%	0.55%	0.16%	0.00%
286MHz vs 333MHz	DDR	0.51%	0.66%	0.24%	0.00%
	SDR	1.03%	0.33%	0.24%	1.01%

Figure Five. Percentage differences between various combinations.

Conclusion

The net result is that there is no advantage in using DDR L3 cache memory over SDR L3 cache memory on the 7450 family processors. There is also very little performance difference between an L3 cache of a given type running twice as fast as another L3 cache of the same type! Clearly it is important to have some sort of L3 cache, but the type and speed of the L3 are almost irrelevant. In fact, the only factor that matters to any degree, is L3 cache size, and even that is not a huge difference (about 7% maximum difference in our benchmark tests, which is measurable but not necessarily noticeable in day to day usage.)

Given the additional cost of DDR cache components, SDR is definitely the proper choice. It would appear that the only real advantage the term ‘DDR’ brings, is as a marketing buzz word.

*Test conditions:

Apple Dual 1GHz processor card with bus ratio set to 10:1 for use in 100 MHz motherboard; 2MB DDR L3 cache (Samsung); second processor disabled by removing Apple CPU Plugins file from System Folder; cache size adjusted via software as needed.

PowerForce G4 Series 100 1GHz processor card with bus ratio set to 10:1, for use in the same motherboard; 2MB SDR L3 cache (Samsung); cache size adjusted via software as needed.

MacOS 9.2.2

Cinebench 2000

724MB RAM

Photoshop 7.0 using [Barefeats.com](http://www.barefeats.com) Photoshop mutiple actions (600MB allocated to Photoshop to minimize disk access)

References:

<http://www.barefeats.com/pmddr.html>

<http://www.hardwarecentral.com/hardwarecentral/reviews/1616/2/>

<http://e-www.motorola.com/brdata/PDFDB/docs/AN1794.pdf>

<http://e-www.motorola.com/brdata/PDFDB/docs/AN2180.pdf>

<http://e-www.motorola.com/brdata/PDFDB/docs/AN2182.pdf>

<http://e-www.motorola.com/brdata/PDFDB/docs/MPC7455EC.pdf>

<http://e-www.motorola.com/brdata/PDFDB/docs/MPC7450UM.pdf>

<http://e-www.motorola.com/collateral/L3SPCALC.xls>

http://www.cs.duke.edu/~kedem/CACHE/Intro_slides.pdf

[http://www.samsungelectronics.com/semiconductors/SRAM/High_Speed/DDR/8M_bit/K7D803671B/ds_k7d8036\(18\)71b.pdf](http://www.samsungelectronics.com/semiconductors/SRAM/High_Speed/DDR/8M_bit/K7D803671B/ds_k7d8036(18)71b.pdf)

[http://www.samsungelectronics.com/semiconductors/SRAM/High_Speed/SPB_n_FT/8M_bit/K7A803608B/ds_k7a8036\(18,32\)08b.pdf](http://www.samsungelectronics.com/semiconductors/SRAM/High_Speed/SPB_n_FT/8M_bit/K7A803608B/ds_k7a8036(18,32)08b.pdf)

©2002 PowerLogix, All rights reserved. All trademarks are property of their respective owners.

PowerLogix, 8701 W. Parmer Lane, Suite 1120, Austin, TX 78729 USA