The IBM POWER2 Architecture Implementations

By Sohel R. Saiyed and Jacob Thomas

This article describes some of the important characteristics of IBM POWER2 Architecture implementations used in products announced by IBM in September 1993 and May 1994.

All POWER2 Architecture implementations (shown in Figure 1) are completely binary compatible. Applications written for the POWER Architecture will also run on the POWER2. The major differences among POWER2 implementations are the size of the data cache, the use of a second level (L2 cache), width of memory interface, and width of processor interface (processor-to-data-cache interface).

These characteristics have different performance impacts on the three classes of codes: integer, floating point, and commercial Transaction Processing (TP). Performance discussions in this article are based on the results of SPECint92™, SPECfp92™, LINPACK, and TPC-C™ benchmarks.

POWER2 Architecture Implementations

In September 1993, IBM announced three RISC System/6000 (RS/6000) servers based on the first implementation of the POWER2 Architecture (the *original* POWER2). Although IBM customers received these three servers with enthusiasm, many hoped for and requested even more—a desktop implementation and greater capacity for commercial TP. IBM set a major goal of satisfying these two requests for the May 1994 announcement.

Meeting these two requests was not easy. The stringent packaging and price targets needed for a desktop product meant several design trade-offs. Achieving greater capacity for commercial TP at a reasonable cost required several other design trade-offs. In the end, IBM achieved the results requested by its customers in the form of two new POWER2 Architecture implementations: the *new-desktop* POWER2 and the *new-server* POWER2. Figure 1 summarizes the three implementations of POWER2 Architecture and the corresponding RS/6000 model numbers.

RS/6000 Model	Package	POWER2 Implementation	Announcement Date
58H, 590	Deskside	Original	September 1993
990	Rack	Original	September 1993
59H	Deskside	New Server	May 1994
R20, R24	Rack	New Server	May 1994
380, 390	Desktop	New Desktop	May 1994

Figure 1. POWER2 Architecture implementations

The three implementations are identical at the processor level, making them completely binary compatible. All three can execute up to six instructions (branch, conditional register, two fixed point, and two floating point) per cycle and have the same instruction cache size (32 KB). Figure 2 summarizes the major differences among these implementations:

- **Data cache size:** Depending on the implementation, data cache size can be 64, 128, or 256 KB.
- **L2 cache:** This is standard on Models 59H, R20, and R24, and an option on Model 390. L2 cache is not available on Models 58H, 590,

Sohel R. Saiyed

Jacob Thomas

POWER2 Implementation	RS/6000 Model	Data Cache/ Line Size	L2 Cache Size	Memory Interface	Processor Interface
Original[1]	58H,590,990	256 KB/256 bytes	N/A	8 words	8 words
New-Server	59H,R20,	128 KB/128 bytes	1 MB	4 words	8 words
	R24		2 MB		
New-Desktop	380	64 KB/64 bytes	N/A	2 words	4 words
	390		0-1 MB		

[1]Four or more memory cards assumed. With less than four cards, the effective cache size and memory interface width are 128 KB and 4 words, respectively.

Figure 2. Key aspects of the POWER2 implementations

990, and 380. Existing application programs will generally exploit the L2 cache without recompilation.

- **Memory interface:** Width of the memory data bus could be eight, four, or two words, depending on the implementation. The width of L2 cache to memory interface (on models with L2 cache) is the same as the memory interface width.
- **Processor interface:** Width of the processor to data cache bus is four words on Models 380 and 390, and eight words on all other models.

Characteristics of Different Classes of Codes

Integer code (SPECint92) typically does not access large amounts of data and usually fits in any large data cache. A reduction in data cache size may reduce performance slightly, but adding an L2 cache may compensate for the performance loss in some application codes. For integer codes, a four-word processor interface is sufficient to keep the two fixed-point units busy. Also, a 32 KB instruction cache is usually sufficient for these types of codes.

Typical floating-point code (SPECfp92 and LINPACK benchmarks) accesses large amounts of data. Smaller data caches and the reduced interface widths have a greater impact on floating-point code than on integer code. The L2 cache can compensate for some lost performance, but not all, particularly on the new desktop models. For some floating-point code, the four-word processor interface may not be sufficient to keep the two floating-point units busy. However, the 32 KB instruction cache is usually sufficient for these types of code.

Commercial TP workloads consist of integer code, but their characteristics differ from those of SPECint92. TP code typically contains a large instruction footprint that does not fit in a 32 KB instruction cache. In this environment, adding an L2 cache significantly reduces the penalty for instruction cache misses. Since TP data is primarily 32-bit integer data, a four-word processor interface can feed the two fixed-point units. Therefore, the new-desktop and new-server POWER2 implementations show significant performance gains on TP code. Although this behavior of code classes is typical, the behavior of customer application codes may differ from these examples.

Performance Differences

This section discusses the performance of RS/6000 models selected from the three POWER2 implementations: Model 590 (original POWER2), Model 59H (new-server POWER2), and Model 390 (new-desktop POWER2). Since these models operate at about the same frequency, the performance results (Figure 3) reflect implementation differences. Discussion of these results is limited to the performance of benchmarks SPECint92, SPECfp92, LINPACK, and TPC-C.

The Models 590 and 59H have nearly identical performance on SPECint92, while the SPECint92 result on the Model 390 is about 7% lower. Since the four-word CPU cache interface (4x32 bit) of the Model 390 should be sufficient to feed the two integer units, the most likely sources of the performance decline are the smaller data cache and the memory interface.

On LINPACK SP (100x100 Single-Precision Benchmark), all three models offer the same performance. The LINPACK SP data fits in the 64 KB

data cache. The L2 cache and the memory interface have minimal impact on this benchmark result. The four-word processor interface is sufficient to keep both floating-point units busy.

On LINPACK DP (100x100 Double-Precision Benchmark), the Models 590 and 59H performance results are the same. Since the LINPACK DP data fits in the 128 KB data cache, the differences in cache size and memory interface do not affect performance. The performance of the Model 390 LINPACK DP is less than half that of Model 590. The eight-word processor interface can deliver two quad-word, floating-point storage reference instructions (two load quads or two store quads) per machine cycle. Since the Model 390 four-word interface operates at half the capacity of the eight-word interface on the Model 590, a 50% drop in performance should be expected. Furthermore, the DP data does not fit in the 64 KB data cache, which causes the performance to drop.

The Model 590 offers slightly better performance than the Model 59H on SPECfp92. The reduction in memory interface width and data cache size is not fully compensated by the addition of the L2 cache. The SPECfp92 of the Model 390 is about 20% lower than the Model 59H results. The smaller data cache and narrower interfaces contribute to the decline in performance compared to Models 590 and 59H.

The new-desktop and new-server POWER2 implementations have smaller data cache lines than the original POWER2 implementation. The smaller cache lines can affect the performance of floating-point code with contiguous memory access.

On double-precision, floating-point code (typical of engineering and scientific applications), the original POWER2 implementation (such as the Model 590) delivers better performance than the new-server implementation (such as the Model 59H).

The Models 59H and 390 show significant performance gains over the Model 590 on the TPC-C benchmark. These gains are the result of adding L2 cache, application tuning, and operating system enhancements. The addition of an L2 cache accounts for a large part of the performance gain. Because the L2 cache is significantly larger than the instruction cache, it reduces the effect of instruction cache misses by containing the majority of code paths executed by the transaction processing software. A four-word processor interface is sufficient to feed the two integer units. The performance loss (if any) caused by a reduced memory interface is inconsequential compared to the performance boost from the large L2 cache. By adding the L2 cache, the 59H performance gain over the 590 is estimated to be about 20% for workloads similar to TPC-C.

	Implementations		
RS/6000 model:	Original 590	New-Server 59H	New-Desktop 390
Clock cycle:	66.7 MHz	66.7 MHz	67.0 MHz
Processor interface:	8 words	8 words	4 words
Data cache size:	256 KB	128 KB	64 KB
L2 cache size:	N/A	1 MB	1 MB
Memory interface:	8 words	4 words	2 words
SPECint92	121.4	122.4	114.3
SPECfp92	254.2	250.7	205.3
LINPACK SP	73.1	73.1	73.0
LINPACK DP	132.0	132.0	55.1
tpmC(c/s)	726.1[1]	1122.3	901.5
$/tpmC	1395[1]	968	909

[1]The Model 590 measurement, performed in 1993, does not include recent enhancements in the TPC-C application and AIX 3.2.5.

Figure 3. POWER2 implementation performance comparison

The TPC-C benchmark results confirm that the new-desktop and the new-server POWER2 implementations are more suitable for commercial applications than the original POWER2 implementations.

Finally, the eight-word memory interface of the original POWER2 implementation requires a minimum of four memory cards to achieve full bandwidth potential. With less than four memory cards, the effective memory interface width and the data cache size become four words and 128 KB, respectively.

Sohel R. Saiyed, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Saiyed is an advisory engineer in systems architecture and performance in the RISC System/6000 Division. Mr. Saiyed has an MS in Computer Engineering from Clemson University and a B.Tech. degree in Electrical Engineering from the Indian Institute of Technology in Kharagpur, India.

Jacob Thomas, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Thomas is an advisory programmer in processor performance in the RISC System/6000 Division. Mr. Thomas has an MS in Statistics from Michigan State University and an MS in Physics from Birla Institute of Technology and Science in Pilani, India.