



PowerPC G5

The World's First 64-Bit Desktop Processor

White Paper

June 2003

Contents

Page 3	Introduction
Page 4	The World's First 64-Bit Desktop Processor An Exponential Leap in Computing Power Memory Addressing up to 18 Exabytes High-Precision Calculations in a Single Clock Cycle Clock Speeds up to 2GHz Industry-Leading 1GHz Frontside Bus Full Support for Symmetric Multiprocessing Native Compatibility with 32-Bit Application Code
Page 7	Next-Generation PowerPC Architecture Ultrafast Access to Data and Instructions Highly Parallel Execution Core Aggressive Queuing and Register Renaming Optimized 128-Bit Velocity Engine Two Double-Precision Floating-Point Units Two Integer Units Two Load/Store Units Condition Register Three-Component Branch Prediction Logic State-of-the-Art Process Technology from IBM
Page 11	Industry-Leading Performance SPEC CPU2000
Page 14	Technical Specifications

Introduction

Key Features

- 64-bit architecture, capable of addressing 18 exabytes of memory
- Clock speeds up to 2GHz
- 1GHz frontside bus for throughput of up to 8 GBps per processor
- Dual independent 1GHz frontside buses in dual processor systems
- Superscalar execution core supporting up to 215 in-flight instructions
- Velocity Engine for accelerated single-instruction, multiple-data (SIMD) processing
- Two double-precision floating-point units for high-speed advanced computation
- Massive three-component branch prediction logic to increase processing efficiency
- Native compatibility with existing 32-bit application code
- State-of-the-art process technology from IBM

The revolutionary PowerPC G5 changes everything you know about personal computing. Suddenly, the next generation of high-performance applications for design and graphics, media production, and scientific research is possible and practical on the desktop. That's because the PowerPC G5 brings a 64-bit architecture to the Mac platform—ushering in an exciting new era in personal computing.

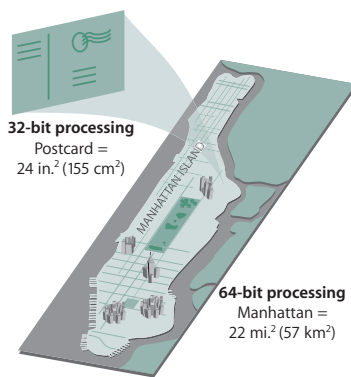
The introduction of the PowerPC G5 is a product of Apple's partnership with IBM, leveraging the most advanced chip design and manufacturing expertise in the world. The results are phenomenal: 130-nanometer fabrication technology, 2GHz clock speeds, and an all-new PowerPC architecture. Together, they put enormous computing power within the reach of personal computer users:

- The ability to address huge amounts of memory provides ultrafast data access, boosting performance for 2D imaging, 3D design, and video rendering tasks.
- A high-bandwidth execution core with 12 functional units improves performance by executing multiple instructions per cycle in parallel.
- An optimized 128-bit Velocity Engine cranks through image editing tasks, high-definition video transitions, and complex scientific analysis.
- Two double-precision floating-point units accelerate 64-bit calculations for 3D visualization, research simulations, and multitrack audio creation.

The G5 processor is making its debut in the Power Mac G5, the world's fastest personal computer according to industry-standard SPEC benchmarks.¹ Now Power Mac users can tackle projects never before possible on a desktop system—and blaze through their work faster than ever. In fact, the Power Mac G5 runs Adobe Photoshop more than two times faster than the fastest Pentium 4-based system.² Best of all, the PowerPC G5 runs 32-bit code—including existing Mac OS X applications—natively, so the transition to 64-bit power is absolutely seamless.

Welcome to the PowerPC G5, the world's first 64-bit desktop processor and the heart of the new Power Mac.

The World's First 64-Bit Desktop Processor



4.3 billion times bigger

To grasp the exponential leap from 32-bit to 64-bit processing, imagine equating the range of numbers a processor can express with a two-dimensional area. A 32-bit processor can express a range of integers equal to the size of a postcard, while a 64-bit processor can express a range of integers larger than the island of Manhattan.

The PowerPC G5 marks the arrival of 64-bit performance to the personal computer market. With 64-bit-wide data paths and registers, this groundbreaking new processor can address vast amounts of main memory and handle multiple large integer and floating-point math calculations in a single clock cycle.

An Exponential Leap in Computing Power

The label “32-bit” or “64-bit” characterizes the width of a microprocessor’s data stream, which is a function of the sizes of its registers and the internal data paths that feed the registers. A 64-bit processor moves data and instructions along 64-bit-wide data paths—compared with the 32-bit-wide paths on 32-bit processors, such as the Pentium 4. In addition, 64-bit processors have wide registers that can store 64-bit numbers as well as 32-bit numbers.

The leap from 32-bit to 64-bit processing represents an exponential advance in computing power. With 32-bit registers, a processor has a dynamic range of 2^{32} , or 4.3 billion—which means it can express integers from 0 to 4.3 billion. With 64-bit registers, the dynamic range catapults to 2^{64} , or 18 billion billion—4.3 billion times larger than the range of a 32-bit processor.

Memory Addressing up to 18 Exabytes

The move to 64-bit processing results in a similarly dramatic leap in the amount of memory supported. A memory address is a special kind of integer, and each address points to one byte in memory. Since memory addresses are computed in 64-bit registers capable of expressing 18 billion billion integers, the PowerPC G5 can theoretically address 18 exabytes (18 billion billion bytes) of virtual memory.


In practice, memory addressing is defined by the physical address space of the processor. The PowerPC G5, with 42 bits of physical address space, supports a colossal 2^{42} bytes, or 4 terabytes, of system memory. Although it’s not currently feasible to purchase 4 terabytes of RAM, very large quantities of memory enable a desktop system to contain a gigantic 3D model, a complex scientific simulation, or a sequence of HD video entirely in RAM—drastically reducing the time to access, modify, and render the data.

High-Precision Calculations in a Single Clock Cycle

With 64-bit-wide data paths and registers, the PowerPC G5 can execute instructions on 64 bits of data in a single clock cycle—making it possible to perform huge integer calculations and highly precise floating-point mathematics. In contrast, a 32-bit processor would have to split up any data larger than 32 bits and process it over multiple clock cycles. The advanced calculation capability of the PowerPC G5 is critical in state-of-the-art applications for scientific simulations, 3D modeling, and video effects.

Clock Speeds up to 2GHz

The PowerPC G5 features a scalable design that enables it to run at clock speeds up to 2GHz. This represents a 600MHz jump, the largest in PowerPC history, over the fastest G4 processor at 1.4GHz. What's more, this breakthrough clock speed boosts performance across the board, accelerating everything from gaming frame rates to digital audio effects.



The bidirectional frontside bus allows data to travel to and from the PowerPC G5 processor at the same time. In dual processor systems, each PowerPC G5 has its own dedicated interface to maximize throughput—compared with dual Xeon-based systems, which force the processors to share a single bus.

Industry-Leading 1GHz Frontside Bus

To harness the power of the G5 processor, a 64-bit Double Data Rate (DDR) frontside bus maximizes throughput between the processor and the rest of the system. Unlike conventional processor interfaces, which carry data in only one direction at a time, the PowerPC G5 features two dedicated unidirectional 32-bit data paths (64 bits total): One travels into the processor and one travels from the processor, with no wait time while the processor and the system controller negotiate which will use the bus or while the bus switches direction. In addition, the data streams integrate clock signals along with the data—allowing the frontside bus to work at speeds up to 1GHz for an astounding 8 GBps of total bandwidth.

In dual PowerPC G5 systems, each processor has its own discrete 1GHz frontside bus. The result is a maximum aggregate bandwidth of 16 GBps on dual 2GHz Power Mac G5 systems, well over twice the 6.4-GBps maximum throughput of Pentium 4- or Xeon-based systems. In addition to providing fast access to main memory, this high-performance frontside bus architecture enables each PowerPC G5 to discover and access data in the other processor's caches—a process called intervention. Cache intervention is made possible by cache coherency, which ensures that the processor always fetches the correct data, even if it has been modified and is in L2 cache.

Full Support for Symmetric Multiprocessing

The PowerPC G5 is designed for symmetric multiprocessing—enabling multiple applications to run independently on different processors or a single multithreaded application to perform multiple tasks simultaneously. For example, while performing an edit, Final Cut Pro can decode two pieces of source video, one on each processor, at the same time.


With dual independent frontside buses and built-in cache coherency, a dual processor system manages priorities between the two processors for maximum efficiency. And since Mac OS X was built from the ground up for symmetric multiprocessing, no special optimization is required for software to take advantage of this powerful capability.

Native Compatibility with 32-Bit Application Code

On other platforms, switching to a 64-bit computer requires migrating to a 64-bit operating system (and purchasing 64-bit applications) or running a 32-bit operating system in a slow emulation mode. With the PowerPC G5, the transition to a 64-bit system is seamless: Current 32-bit code—such as existing Mac OS X and Classic applications—runs natively at processor speed, with no interruptions to your workflow and no additional investment in software.

This easy compatibility is possible because the PowerPC architecture, unlike competing instruction sets, was designed from the beginning to run both 32-bit and 64-bit application code. And because the PowerPC G5 uses the same Velocity Engine instruction set introduced in the PowerPC G4, applications that have been optimized for Velocity Engine will immediately run faster on the new processor. What's more, as applications are optimized and as Mac OS X is further enhanced for the PowerPC G5, performance gains will be even greater.

Next-Generation PowerPC Architecture



The PowerPC G5 is fabricated using state-of-the-art 130-nanometer circuitry with more than 1130 feet of ultrathin wiring—nearly 800 times thinner than a human hair.

The PowerPC G5 is a highly parallel implementation of the PowerPC architecture, capable of handling multiple assorted tasks at the same time. It's based on the execution core of IBM's 64-bit POWER4 processor—recipient of the *Microprocessor Report's* 2001 Analyst's Choice Award for Best Workstation/Server Processor, which recognizes excellence in semiconductor technology innovation, design, and implementation. With two double-precision floating-point units, advanced branch prediction logic, and a high-bandwidth frontside bus, the POWER4 drives IBM's top-of-the-line pSeries 690 servers.


Apple collaborated with IBM to leverage this industry-leading design for the next generation of personal computing. The development of the PowerPC G5 builds on previous PowerPC designs, combining an optimized Velocity Engine with a superscalar, superpipelined execution core that supports up to 215 simultaneous in-flight instructions. And all this power is fabricated in IBM's state-of-the-art 130-nanometer process technology using high-performance silicon-on-insulator (SOI) transistors and copper interconnects.

Ultrafast Access to Data and Instructions

The PowerPC G5 features processing innovations that optimize the flow of data and instructions—making it ideal for media streaming, HD video editing, real-time effects, audio synthesis, image processing, 3D rendering, numerical analysis, and physical modeling.

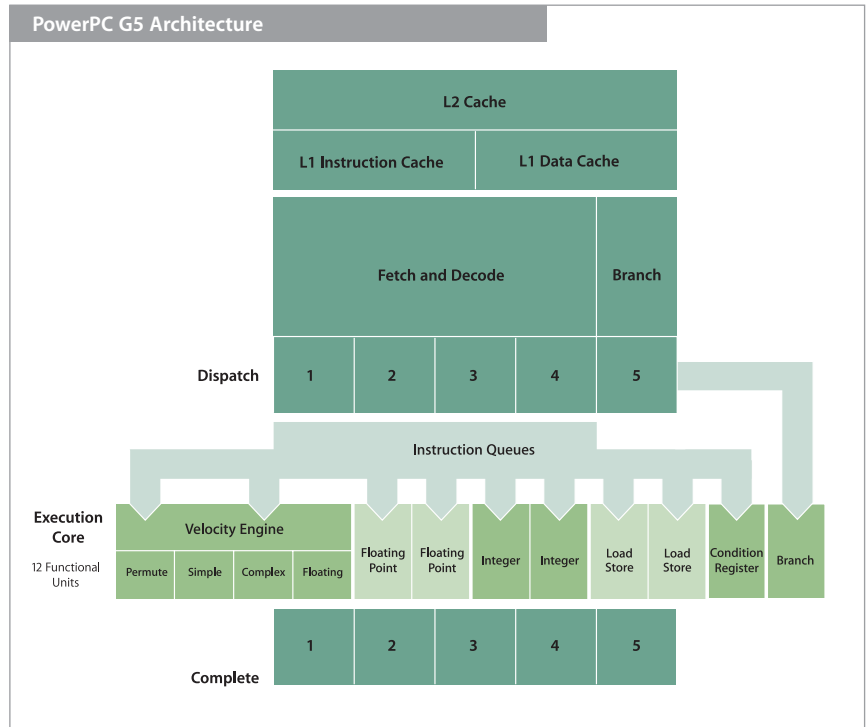
It starts with 512K of L2 cache that provides the execution core with ultrafast access to data and instructions—up to 32 GBps. Instructions are prefetched from the L2 cache into a large, direct-mapped 64K L1 cache at up to 64 GBps. As they are accessed from the L1 cache, up to eight instructions are fetched per clock cycle. Next, instructions are decoded and divided into smaller, faster-executing operations. In addition, 32K of L1 data cache can prefetch up to eight active data streams simultaneously.

This efficient preparation maximizes processing speed as instructions are dispatched into the execution core and data is loaded into the registers.



Up to 215 in-flight instructions

A wide architecture with 12 discrete processing units enables the PowerPC G5 to contain up to 215 in-flight instructions simultaneously—71 percent more than the 126 instructions in a Pentium 4.



Highly Parallel Execution Core


At the heart of the PowerPC G5 is an entirely new superscalar, superpipelined execution core, composed of 12 functional units that execute different types of instructions concurrently for massive data throughput. Before instructions are dispatched into the functional units, they are arranged into groups of up to five. Within the core alone, the PowerPC G5 can track up to 20 groups at a time, or 100 individual instructions. This efficient group-tracking scheme enables the PowerPC G5 to manage an unusually large number of instructions “in flight”: 20 instructions in each of the five dispatch groups, in addition to 100-plus instructions in the various fetch, decode, and queue stages.

Aggressive Queuing and Register Renaming

The PowerPC G5 includes extensive instruction queueing and renameable data registers that maximize the utilization of each functional unit. Once an instruction group is dispatched into the execution core, it is broken into individual instructions, which proceed to the appropriate functional unit. Each functional unit has its own dedicated queue, where multiple instructions are sequenced for processing.

At the same time, the load/store units load the associated data from L1 cache into the data registers behind the units that will be processing the data. To improve processing efficiency, the PowerPC G5 features a large number of rename registers that act as proxies, or placeholders, until the appropriate data arrives for execution. The instruction is held in queue, allowing other operations to take place until the actual data is loaded into the registers.

When operations on the data are complete, the PowerPC G5 recombines the instructions into the original groups of five, and the load/store units store the data in cache or main memory to complete processing.



The Velocity Engine can manipulate 128 bits of data at a time, up to four times faster than the general processing units in 32-bit processors.

Optimized 128-Bit Velocity Engine

The PowerPC G5 uses a dual-pipelined Velocity Engine optimized with two independent queues and dedicated 128-bit registers and data paths for efficient instruction and data flow. This 128-bit vector processing unit accelerates data manipulation by applying a single instruction to multiple data at the same time, known as SIMD processing. Originally implemented in the PowerPC G4, the Velocity Engine in the PowerPC G5 uses the same set of 162 instructions, enabling it to run—and accelerate—existing Mac OS X applications that have been optimized for the Velocity Engine.

Vector processing is useful for transforming large sets of data, such as manipulating an image or rendering a video effect. For example, when a designer uses a filter to apply a motion blur to an image, each pixel of the image must be changed according to the same set of instructions—a highly repetitive processing task. Each Velocity Engine pipeline speeds up this task by processing up to 128 bits of data, in four 32-bit integers, eight 16-bit integers, sixteen 8-bit integers, or four 32-bit single-precision floating-point values, all in a single clock cycle.

Two Double-Precision Floating-Point Units

Today's powerful applications demand both precision and performance. That's why the PowerPC G5 has two double-precision floating-point units, enabling it to complete at least two 64-bit mathematical calculations per clock cycle. In fact, each floating-point unit can perform both an add and a multiply with a single instruction. This dramatically accelerates highly complex computations that are critical in research simulations and in many of the applications used to manipulate or render 3D graphics and video content.

Weather prediction is one example of a highly iterative computing task made possible by floating-point math. Large-scale models simulate weather patterns over time by measuring multiple influences, such as atmospheric pressure and airflow, at various instants and recalculating the model every minute. The PowerPC G5 provides the precision and performance to deliver accurate results within a useful timeframe.

Two Integer Units

Integer units perform simple integer mathematics—such as add, subtract, and compare—which are commonly used in many basic computer functions, as well as in imaging, video, and audio applications. The PowerPC G5 has two integer units capable of a broad range of simple and complex instructions involving both 32-bit and 64-bit data. What's more, they take full advantage of the processor's 64-bit registers and data paths to complete 64-bit integer calculations in a single pass.

Two Load/Store Units

Load/store units manage data as it is processed, loading it into the registers of each functional unit and, after processing, storing the new data in L1 cache, L2 cache, or main memory, as appropriate. The PowerPC G5 is generously equipped with three large sets of registers: A general-purpose register file contains 64-bit registers for integer calculations; a floating-point register file contains 64-bit registers for floating-point calculations; and a vector register file contains 128-bit registers for the Velocity Engine. Each register file holds 32 registers for architected values, as well as 48 rename, or proxy, registers. With two load/store units, the PowerPC G5 is able to keep these registers filled with data for maximum processing efficiency.

Condition Register


This special 32-bit register summarizes the states of the floating-point and integer units. The condition register also indicates the results of comparison operations and provides a means for testing them as branch conditions. By bridging information between the branch unit and other functional units, the condition register improves the flow of data throughout the execution core.

Three-Component Branch Prediction Logic

Advanced processors use branch prediction and speculative operation to increase efficiency. A branch is a question in the processing queue: Which instruction should go next? Branch prediction anticipates the answer; and speculative operation causes that instruction to be executed. If the prediction is correct, the processor works more efficiently—since the speculative operation has executed an instruction before it's required. If the prediction is incorrect, the processor must clear the unneeded branch, as well as any related data and instructions, resulting in an empty space called a pipeline bubble. Pipeline bubbles reduce performance as the processor marks time waiting for the next instruction.

The PowerPC G5 features an innovative three-component branch prediction logic to reduce pipeline bubbles and maximize processor efficiency. The success or failure of each prediction is captured in three 16K branch history tables—local, global, and selector—that are used to improve the accuracy of future branch predictions.

Local branch prediction takes place as individual instructions are fetched into the processor and the types of branches are recorded in the local branch history table. Global branch prediction occurs at the same time: Branches are identified in their processing context, relating to preceding and subsequent operations; and the results are recorded in the global branch history table. The third, “selector” history table identifies which prediction type, local or global, was more accurate in predicting the outcome of each branch. This dynamic local/global/selector branch history scheme can predict branch processes with a high degree of accuracy, allowing the PowerPC G5 to efficiently use every processing cycle.



Copper interconnects improve conductivity and boost processor performance.

State-of-the-Art Process Technology from IBM

The PowerPC G5 is fabricated in one of IBM's world-class semiconductor manufacturing facilities. With industry-leading build, assembly, and test technology, IBM uses a cutting-edge 130-nanometer process with more than 58 million silicon-on-insulator (SOI) transistors and eight layers of copper interconnects.

SOI refers to the placement of a thin layer of insulator between transistors and bulk silicon. When transistors are built on this SOI layer, their capacitance, or the tendency to store an electrical charge, is reduced—resulting in faster operation.

For further performance gains, an additive-copper wiring process replaces the conventional subtractive-aluminum process. As semiconductor wires are made thinner and narrower, aluminum resists the flow of electricity and slows down transmission of electrical signals. Copper wiring is less resistant and results in a 40 percent gain in conductivity, again for faster processor operation.



IBM is a worldwide leader in processor fabrication technologies, with a new \$3 billion, state-of-the-art facility in East Fishkill, New York.

Industry-Leading Performance

Ultrafast clock speeds and a highly parallel 64-bit architecture make the PowerPC G5 ideal for next-generation multimedia, graphics, and scientific applications. Integer and floating-point math calculations are faster than ever thanks to 64-bit-wide registers and data paths.

To demonstrate the performance advantages of this groundbreaking new processor, Apple put the dual 2GHz Power Mac G5 to the test against the top-of-the-line Pentium 4- and Xeon-based systems, the 3GHz Dell Dimension 8300 and the dual 3.06GHz Dell Precision 650, respectively.

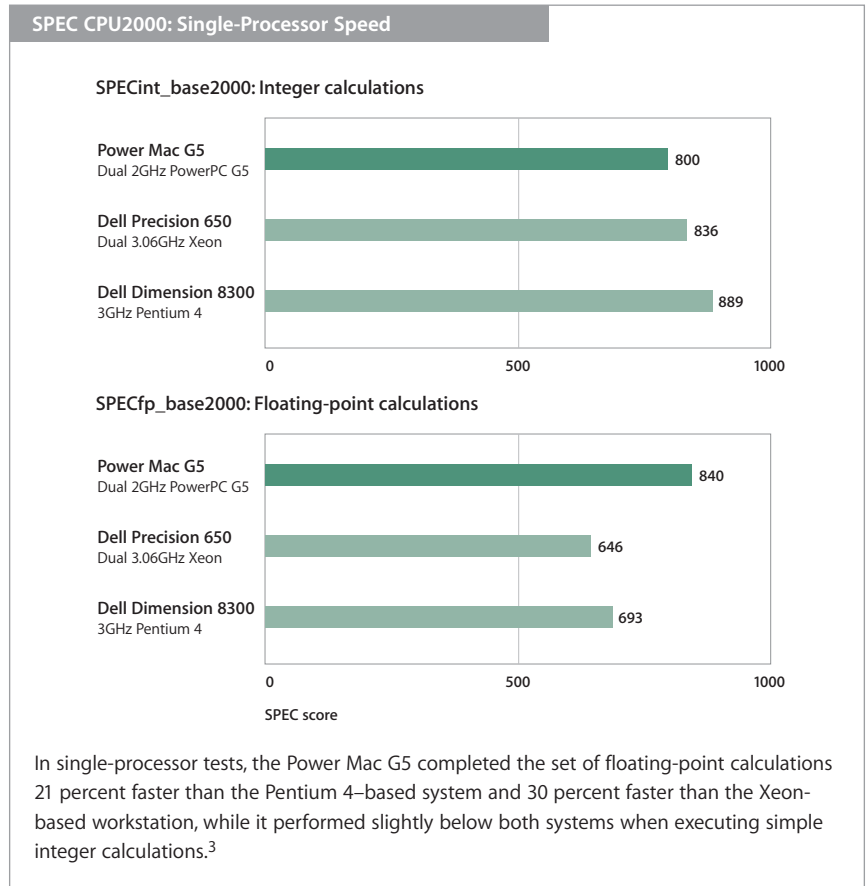
SPEC CPU2000

The Standard Performance Evaluation Corporation (SPEC) CPU2000 benchmark suite is the recognized industry standard for assessing processing performance. SPEC is a nonprofit organization of hardware and software vendors, universities, and consultants. They developed the SPEC CPU2000 benchmarks based on actual end-user applications. These tests depend on processor, memory subsystem, and compiler performance when executing integer and floating-point computations. For more information on the benchmarks, see www.spec.org.

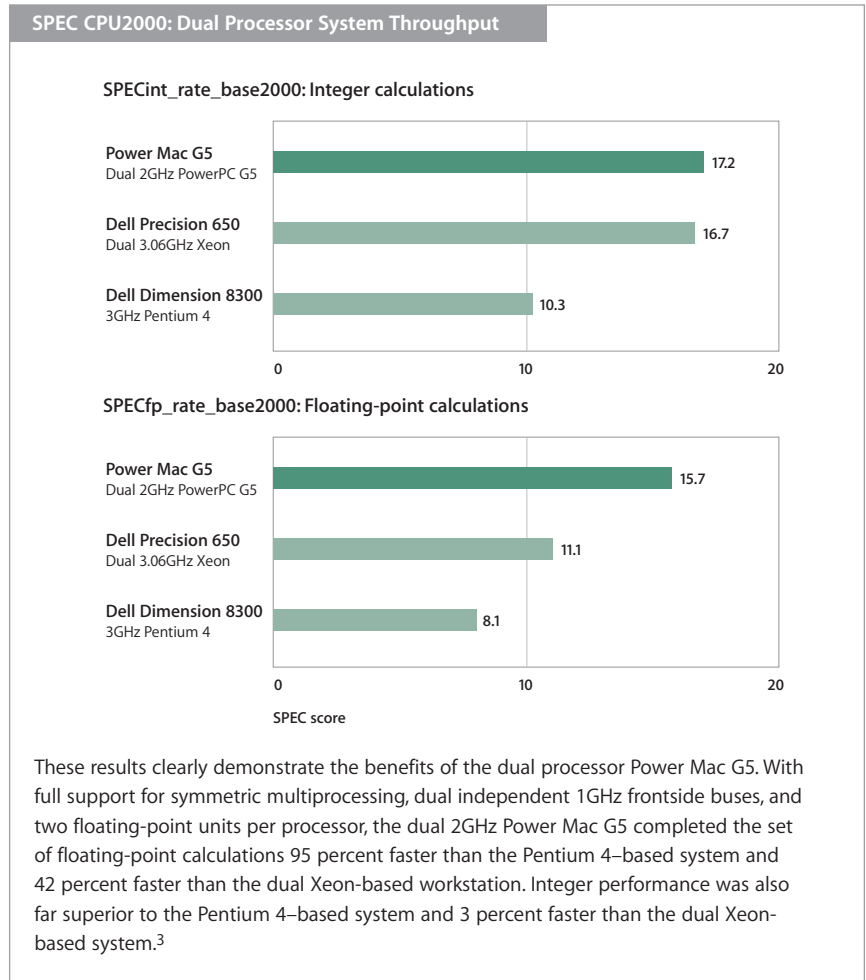
VeriTest, an independent testing service, conducted the tests according to the specifications on the SPEC website. SPEC provides the benchmarks as source code, which VeriTest compiled with GCC, a cross-platform, open source compiler used by programmers around the world. The Power Mac G5 used Mac OS X v10.2.7 (G5), and the Intel-based systems used Red Hat Linux 9.0.

The SPEC CPU2000 benchmark suite comprises four metrics—SPECint_base2000, SPECfp_base2000, SPECint_rate_base2000, and SPECfp_rate_base2000—to represent different performance characteristics of the system.

SPECint_base2000 and SPECfp_base2000 measure the speed of a single task—either an integer calculation or a floating-point calculation—executing on a single processor. Each test measures how long the processor takes to complete the benchmark set of single tasks relative to a SPEC-defined baseline score. SPECint_base2000 is composed of eleven C and one C++ benchmark applications, including a chess program, a data compression utility, and a place-and-route simulator. SPECfp_base2000 consists of six Fortran-77, four Fortran-90, and four C benchmark applications, including shallow-water modeling, neural-network simulation, and computational chemistry.



For comparisons that more accurately demonstrate the performance of a dual processor system, VeriTest used the “SPEC rate” metrics, which recognize multiple processors. With SPECint_rate_base2000 and SPECfp_rate_base2000, the benchmark code is compiled and multiple copies are run concurrently, allowing both processors to work in parallel. SPEC rate tests determine the number of times a system can complete the benchmark per hour, also referred to as system throughput.



For a detailed report of SPEC CPU2000 test results, see www.veritest.com.

Technical Specifications

64-bit PowerPC processor architecture

- Virtual address range: 64 bits, or 18 exabytes
- Physical address range: 42 bits, or 4 terabytes
- Full 64-bit data paths and registers
- Native support for 32-bit application code
- 64K L1 instruction cache; 32K L1 data cache
- 512K internal L2 cache
- Dedicated data flow for dividing one instruction into two internal operations
- Microcoded instructions for up to four internal operations
- Support for up to eight outstanding L1 cache line misses
- Hardware-initiated instruction prefetching from L2 cache
- Hardware- or software-initiated data stream prefetching; support for up to eight active streams
- Maximum core frequency: 2GHz

Frontside bus

- Throughput: up to 8 GBps per processor
- Frequency: 1GHz DDR
- Width: 64-bit (32-bit in, 32-bit out) bidirectional

Wide execution core

- Support for up to 215 in-flight instructions
- In-order dispatch of up to five operations into distributed issue queue structure
- Simultaneous issue of up to 10 out-of-order operations:
 - One Velocity Engine permute operation
 - One Velocity Engine arithmetic logic operation
 - Two floating-point operations
 - Two fixed-point register-to-register operations
 - Two load or store operations
 - One condition register operation
 - One branch operation
- Out-of-order and speculative issue of load operations
- Dual-pipeline 128-bit Velocity Engine for single-instruction, multiple-data (SIMD) processing
- Two independent floating-point units for double-precision calculations

Three-component branch prediction logic

- Speculative superscalar inner core organization
- Fast, selective flush of incorrect speculative instructions and results
- Prediction of up to two branches per cycle
- Support for up to 16 predicted branches in flight
- Prediction hints added to branch instructions
- Prediction support for branch direction and branch addresses

Physical specifications

- 58 million transistors
- 130-nanometer, silicon-on-insulator (SOI) process
- Die size: 118 square millimeters

Comparison of PowerPC G4 and PowerPC G5

	PowerPC G4	PowerPC G5
Architecture	32-bit	64-bit
Addressable memory	4 gigabytes	4 terabytes
Maximum clock speed	1.4GHz	2GHz
Frontside bus	Up to 167MHz shared	Up to 1GHz per processor
In-flight instructions	16	215
Floating-point units	One	Two
Integer units	One	Two
Load/Store units	One	Two
L1 data cache	32K	32K
L1 instruction cache	32K	64K
L2 cache	256K	512K
Branch prediction logic	Local	Local/Global/Selector
Process technology	180-nanometer	130-nanometer
Die size	106 square millimeters	118 square millimeters

For More Information

For more information about the PowerPC G5 processor, visit www.apple.com/g5. For more information about the Power Mac G5, visit www.apple.com/powermac.

¹“World’s fastest” based on SPEC CPU2000 benchmark results and leading professional application performance tests against 3GHz Pentium 4–based Dell Dimension 8300 and dual 3.06GHz Xeon-based Dell Precision Workstation 650. SPEC CPU2000 benchmarks run with GCC 3.3 compiler and independently tested; professional applications tested by Apple, June 2003. ²Tests performed by Apple in June 2003 using preproduction Power Mac G5 units. The Power Mac systems ran a PowerPC G5–optimized version of Photoshop 7.0.1 including optimized AltivecCore, ACE, and BIB Carbon Libraries; the Dell Dimension ran Photoshop 7.0.1. ³Based on SPEC CPU2000 benchmark results against 3GHz Pentium 4–based Dell Dimension 8300 and dual 3.06GHz Xeon-based Dell Precision Workstation 650m, performed by VeriTest, June 2003.