

*The Series 10000*

*Personal Supercomputer*

*Apollo Computer Inc.*

*I n s i d e  
a N e w  
A r c h i t e c t u r e*



## *The Personal Supercomputer: A Performance Milestone*



△  
*The Series 10000  
brings supercomputer  
performance to the  
office, which means  
advanced applica-  
tions like fluid dynam-  
ics can now be done  
on the desktop.*


From the start, Apollo's goal was to create a new class of computer: one that combined supercomputer performance with the comfort, convenience, and versatility of an interactive graphics workstation. To meet this ambitious goal, our engineers started from scratch, using the newest technologies and best ideas from today's supercomputers and advanced graphics systems to design a completely new architecture.

The result is Apollo's Series 10000 Personal Supercomputer™ based on our new *PRISM*™ architecture. Fully compatible with existing Apollo systems, the Series 10000 is the first Apollo system to boast the groundbreaking *PRISM*

(Parallel Reduced Instruction Set Multiprocessor) architecture. And it marks the beginning of a new generation of workstations.

From top to bottom, the Series 10000 Personal Supercomputer incorporates today's leading-edge thinking and technologies. For example, it features multiprocessors and 64-bit data paths taken from supercomputer designs. Its groundbreaking compilers are based on sophisticated data flow techniques developed first at major research institutes.





Its central processing units feature high-speed integer processors *and* independent, coequal floating point processors providing parallel instruction execution. And its tightly integrated graphics capabilities include techniques usually confined to specialized display systems.

All this makes the Series 10000 a milestone for the computer industry. But equally important is the unprecedented performance it brings to the office environment.

#### *Holistic Engineering is the Key*

Many new machine designs promise great theoretical leaps in performance, but because they don't meet the challenges of overall system performance and throughput, they fail to deliver. A promised gain in CPU power, for instance, is often held back by the limitations of the compiler. Or the expectation of high-speed throughput is constrained by insufficient bus bandwidth.

At Apollo, we took a *holistic* approach to building the Series 10000, which means every component was carefully designed to work with every other component. By taking a holistic approach and designing from scratch, our engineers were unfettered by traditional limitations. As a result, not only does the Series 10000 feature the fastest microprocessors, its system architecture, operating system, and compilers have all been developed together, with a great deal of cross-influence, to deliver unprecedented *system* performance.

This means, more specifically, that the central processors, with independent integer processors (IPs) and floating point

units (FPUs), are able to handle the most intense computational applications. That the compilers, using sophisticated data flow techniques, are able to easily convert and map existing applications to the *PRISM* architecture ensuring fast response. That the system architecture, based on supercomputer technology, features wide datapaths throughout and easily handles a massive flow of information to as many as four processors. And that high-speed, semicustom CMOS and custom ECL are used throughout so that performance is never held back by logic speed.

#### *The Series 10000 at a Glance*

The Series 10000, the newest and highest-performance family of Apollo workstations, brings supercomputer-class performance to the office environment. It has the power and versatility to handle the full gamut of the most demanding technical applications, from compute-intensive floating-point and integer programs to applications presently residing on mainframes and supercomputers.

The Series 10000 is a superior computational resource for individuals, workgroups, departments, or entire corporations. Here are some highlights:

- *Multiprocessor Design.* Apollo's symmetrical multiprocessor design, with up to four CPUs, lets users match their application environment to processing performance.
- *High-performance RISC CPUs.* Apollo has incorporated key




*Interior views of the left and right sides of the Series 10000 reflect the system's highly efficient design and easy accessibility.*



- A** *Gull wings*
- B** *5-slot VME card cage*
- C** *Power supply and CPUs*
- D** *5 1/4-inch disk drives*





RISC features such as single-cycle LOAD/STORES, fixed-length instructions, and delayed branching to achieve maximum performance from today's advanced semiconductor technology.

- *Superior Floating Point Performance.* Each CPU includes an independent FPU that delivers a performance level approaching that of supercomputers.
- *Large Dual Caches.* Separate 64KB data and 128KB instruction caches dramatically boost bandwidth by continuously feeding the processors and delivering maximum application throughput.
- *64-bit Data Paths.* This is the first workstation to incorporate 64-bit data paths and registers, which minimizes machine cycles per instruction and makes it possible for double-precision IEEE 754 floating point operations to be processed in a single cycle.
- *Fast System Bus.* Communications from CPU to main memory and other processors is over a high-speed 150 megabyte/second, 64-bit bus, which delivers sustained throughput for complex applications.
- *Parallel Instruction Execution.* Wide data paths allow simultaneous IP and FPU instruction dispatch and parallel execution of up to three operations. This lets floating point and integer operations occur within one cycle and increases system response.
- *Data Flow Compilers.* Sophisticated data flow techniques used in the compilers direct the instruction flow to take full advantage of local parallelism in the hardware.
- *Large Main Memory.* Up to 128 megabytes of main memory is

16-way interleaved to achieve the rapid I/O speeds required in multiprocessor designs.


- *Shared Virtual Memory.* Shared virtual memory gives programs a common view of virtual memory and allows transparent migration of operations from one processor to another, as well as sharing of data among processes running on different processors.
- *Industry-standard Buses.* The Series 10000 supports both the IBM PC AT®-compatible bus and VME bus for links to all standard peripherals.
- *High-performance Mass Storage.* Up to four 5¼-inch ESDI fast actuator disk drives providing as much as 3 gigabytes of local storage can be configured in the system, and disk striping is supported for maximum data throughput and fast data access.
- *Efficient, Quiet Packaging.* The new power supply technology boasts an 87% efficiency rating, while the system produces a total sound output below 55db, which means quiet, cool operation.

The Series 10000 is the beginning of a new generation of Apollo systems. It combines supercomputer performance with workstation convenience and, thereby, sets a new standard for office computing. And it offers an unrivaled growth path as Apollo builds on the new *PRISM* technology in the years to come.

◁  
The Series 10000 exploits the most advanced technology: shown to the left is a 1.5 micron semi-custom CMOS VLSI processor chip.



## *Customized Performance: Multiprocessing*



*The Series 10000 supports up to four central processors—each with an independent integer processor, FPU, dual caches, and memory management unit.*


*The starting point for the Series 10000 high-power design is a highly efficient multiprocessor architecture. The system can support as many as four independent CPUs—each one with its own floating point processor, integer unit, cache memory, and memory management unit. To ensure that users can fully benefit from a multiprocessor configuration, each processor is entirely independent and can execute an instruction stream either by itself or concurrently with the other CPUs without user intervention. This makes it easy for existing applications to exploit the multiprocessor architecture.*

### *How It Works*

In today's typical multitasking O/S environment, tasks are continuously waiting for CPU processing time. The Series 10000 solves this problem by speeding up each process with creative and fast CPU design, and then applying multiple CPUs to increase overall system throughput. Actual multiprocessor task management is handled by the operating system. Shared O/S code ensures that each free processor selects the next highest priority process from a common ready-process queue.

In cases where explicit interprocess communication is required, the Series 10000 uses no special synchronizing methods. As a result, users can run explicitly dependent programs using standard system calls.





The most important of the synchronizing methods is *system lock*, which provides exclusion by giving only one processor at a time access to protected data. Holding a lock stops other processors from seizing the data, but does not interfere with any memory system activity. Releasing the lock flushes all changes and updates memory for data coherency. Moreover, it doesn't slow DMA (Direct Memory Access) or non-interlocked processor reads or writes. For programs requiring explicit coordination across multiple processors, the system lock may be added or released in user mode.

#### *Extensive Memory Support*


To make sure that all active processors have the necessary bandwidth for interactive response, the Series 10000 boasts an unusually extensive support system. Foremost among its features is the high bandwidth supplied by the pipelined main memory array and access path. There are also cache coherency guarantees, which are extended to the instruction, data, and translation caches. And the instruction set as well as the memory architecture are specifically designed for symmetric multiprocessing.

To provide complete processor independence, the Series 10000 incorporates extensive CPU symmetry. For example, all processors have a common view of memory and share a common access path to the I/O. And an interrupt may be directed to any processor. As a result, processes can migrate from one processor to another without significant penalty.

#### *Near-linear Multiprocessor Performance*

The ideal goal of any multiprocessor design is a direct linear relationship between system performance and the number of processors online. Apollo's Series 10000 achieves near-linear performance for today's typical application workloads. But deviations can be caused by effects ranging from memory paging to a shortage of tasks necessary to keep all processors busy.

The Series 10000 also anticipates the trend toward the development of multiprocess-based applications. As more and more of these applications become available, they will be able to take full advantage of the system's multiprocessor performance.




*The fully customized ECL technology used in the floating point processors makes possible supercomputer-class FPU performance.*



*Adding processors to the Series 10000 can result in near-linear performance increases.*



## Balance of Power: The CPU



*The heart of the Series 10000 is its extremely powerful and uniquely designed central processing unit which incorporates a high degree of local parallelism as well as advanced RISC (Reduced Instruction Set Computer) concepts.*

The RISC instruction set is implemented directly in hardware as opposed to the microcode used in traditional CISC (Complex Instruction Set Computer) computers. This eliminates a level of interpretation and reduces overhead per instruction, so execution is faster. It also makes possible parallelism through pipelining and the simultaneous execution of serial instructions.

Among the RISC concepts the system implements is *fixed length instructions*, which is a simplified instruction format that boosts pipeline flow. *Delayed branching*, which executes instructions after a branch until the branch target is fetched, eliminates delays caused by waiting for pipeline refilling. And *single-cycle execution*, which ensures virtually all instructions are executed in one machine cycle.



- A** Integer processor
- B** Floating point register file
- C** Floating point ALU
- D** Floating point multiplier
- E** Dual caches
- F** Memory management unit



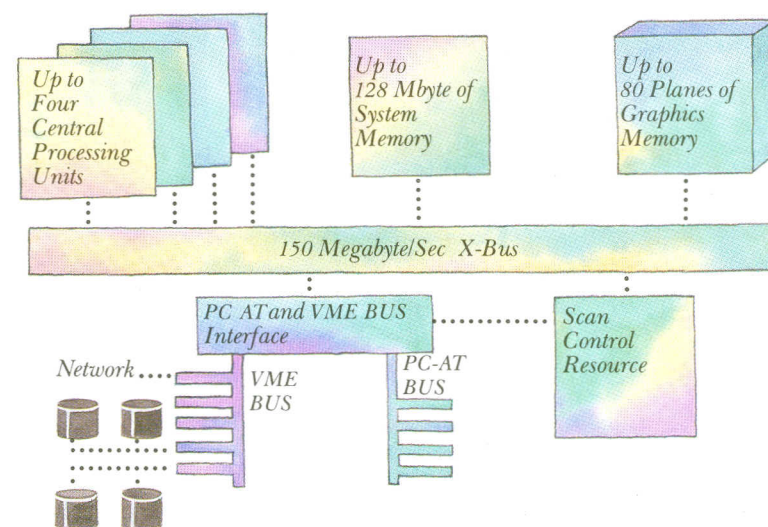
Even more significant and unlike traditional approaches, the system combines both an integer processing unit (IP) and floating point unit (FPU) in every central processing unit. Although the two are independent of each other, they're tightly integrated and designed to work in unison to reduce set-up time and improve response.

In essence, the floating point unit is a peer of the integer processor, and because of this the traditional overhead problems associated with floating point coprocessors have been eliminated. The IP is a 1.5 micron semicustom CMOS, VLSI RISC-based design, while the FPU features a semicustom CMOS register file combined with an independent ALU and multiplier that exploit custom ECL technology for unprecedented high-speed floating point computations. This allows single-cycle execution of all the instructions in the instruction set except for floating point divides, floating point square roots, and integer divides.

Each CPU includes dual caches, one for instructions and the other for data. With a dedicated instruction cache, both the IP and FPU can receive 64-bit instructions that are executed simultaneously every cycle. The independent data cache eliminates conflicts between instruction fetches and loads and stores, making possible single-cycle executions of loads and stores.

### High-speed IP

The system's integer processing unit features an interlocked, multistaged pipeline that



◀ *System Architecture. The Series 10000's well-balanced architecture means all system components are designed to work together to deliver high system performance.*

ensures single-cycle execution of virtually all instructions. It also boasts a special power shifter that works with the integer ALU in performing such operations as insert and extract. With this shifter, the IP can shift the contents of a 32-bit value 0-31 places in one cycle and merge the results with the ALU output. And to help with scaled indexing and operations like small integer multiplication, there's a preshifter residing between the register file and the ALU. In addition, the IP has 32 32-bit registers for storing intermediate operands. The result of all this is a very fast and efficient integer unit.

The integer pipeline supports synchronous instructions and is based on a single-phase clocking strategy, with each macro instruction consisting of one, and only one, execution cycle. All integer processor control and data paths are built around the instruction pipeline and each

instruction pipeline stage comprises one machine cycle.

Integer pipeline depth is a function of the current instruction. The instruction set includes the following forms: register-to-register with a pipeline depth of 4; load and store, each with a 5 depth; macro branches with a depth of 2 (either taken or not taken); and call instructions with depths of 4.

### Low-latency Floating Point

In place of the traditional "coprocessor" approach, the Series 10000 has tightly integrated, low-latency, high-performance floating point processors. As the FPU is a logical equivalent of the IP, floating point instructions can be dispatched in parallel with integer operations, doubling the

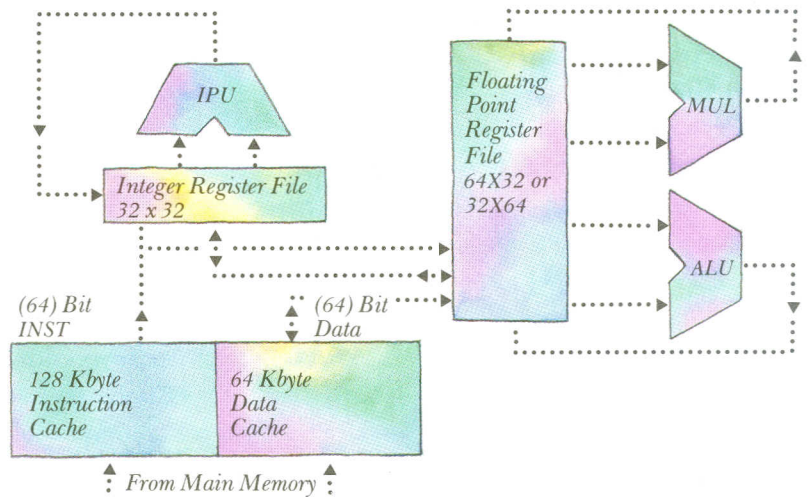


throughput compared to conventional techniques. In addition, the floating point unit has both ALU and multiplier, which can operate in parallel using compound instructions.

This local parallelism lets each CPU execute up to three operations in a single cycle. For example, an integer operation plus a floating point multiply and add. This means a more efficient use of both system and cache memory, and unprecedented performance from a single CPU. When multiple processors are added to the system, the result is super-computer-class performance at workstation prices.



*CPU Subsystem. Each of the up to four CPUs boasts independent integer and floating point units, plus dual instruction/data caches.*

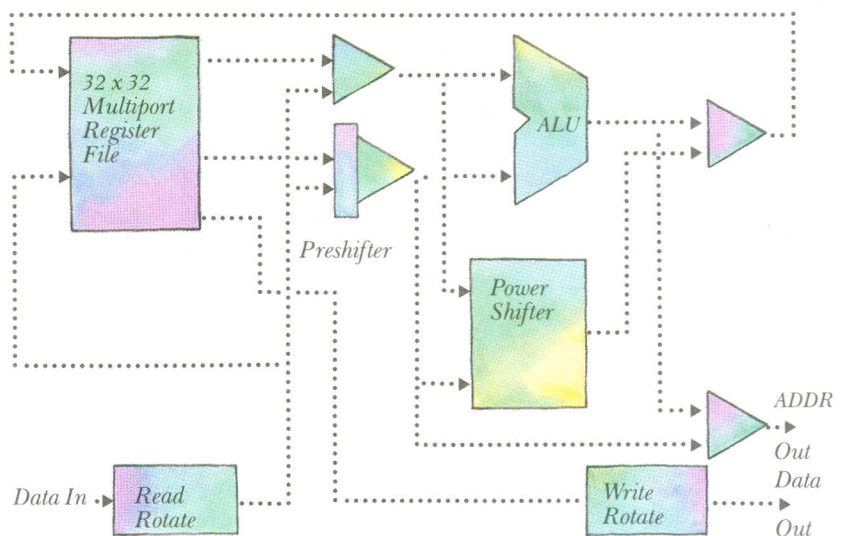


### Large Floating Point Register File

The Series 10000 is specifically designed to deliver superior scalar processing, but the inclusion of a large floating point register file also makes it outstanding for processing large arrays commonly associated with vector processors. The register file can be treated as 32 64-bit registers or 64 32-bit registers and this unusual size and versatility make the system uniquely qualified for such applications. In addition, the large floating point register file is multiported, with a sustained bandwidth in excess of 1 gigabyte per second. This means more data can be kept within the register file, reducing memory access that can impede high-speed processing.



*The IP features a large register file and multistaged pipeline, as well as special shifters for integer processing.*



### New Hybrid Cache Design

Besides being specifically developed to allow multiple operations to be executed in a single instruction cycle, the Series 10000's caches incorporate the

best features of both fully physical and fully virtual caches. The result is a new *virtually indexed, physically tagged write-through cache* that lets cache RAM access proceed entirely overlapped in time with any virtual-to-physical address translation. This overlap reduces memory access pipeline depth, guaranteeing single-cycle execution and eliminating the translation penalty typical with physically indexed designs.

The physical tags allow cache validation across processors. The addressing scheme, based on the virtual address coupled with the physical tag, maintains coherency and allows data to be shared by multiple processors.

The Series 10000 features separate 128K byte instruction and 64K byte data direct-mapped caches, which means that instruction fetch operations can be fully overlapped with memory operand access. The dual-cache design provides a wider bandwidth to the CPU for improved throughput. And the data cache has been specifically designed to complete loads and stores in a single cycle. As a result, load or store instructions require no more processor execution time than a register-to-register instruction (1 cycle).

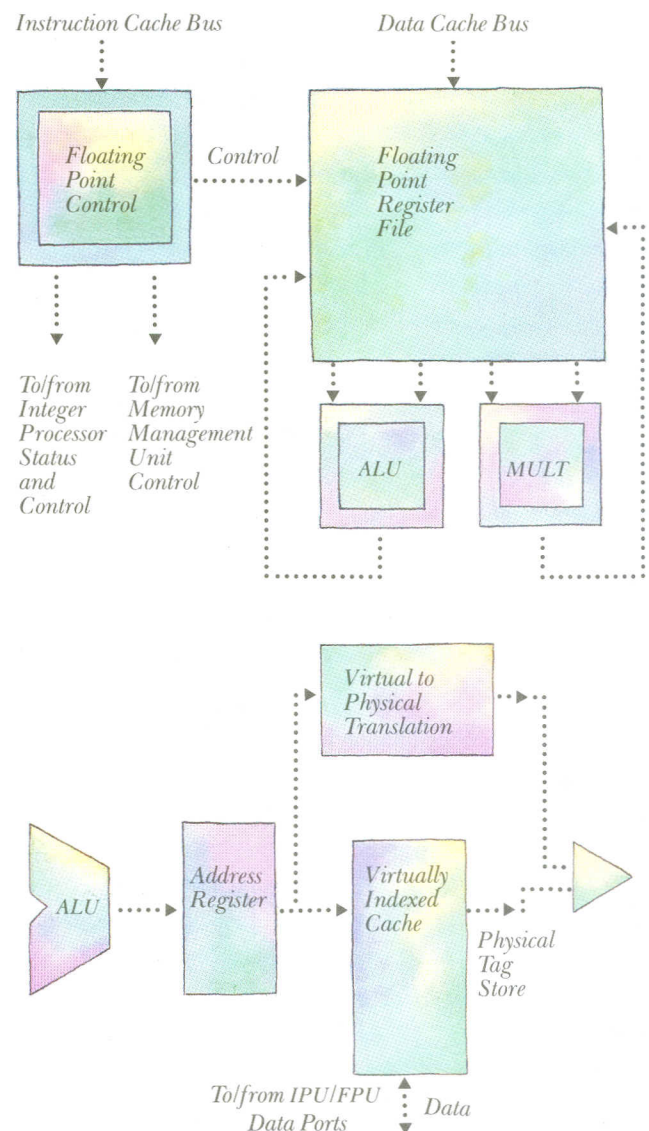
In addition, both the instruction and data caches are a full 64 bits wide. This unique feature allows integer and floating point instructions to be dispatched in parallel, as opposed to traditional serial dispatching. The power of the 64-bit data cache is fully realized during double-precision operations because floating point values can be transferred to or from memory every machine cycle.

This means double-precision operations can be executed at nearly the same speed as single-precision math.

### *Multiprocessor Cache Coherency*

The cache design is also important to the Series 10000's multiprocessor architecture. The caches have been carefully constructed to ensure cache data coherency in a true multiprocessing environment. The write-through cache design ensures that all data modifications are echoed through to main memory, allowing old copies to be purged from processor caches and thereby making process migration transparent. This approach is used so a processor only needs to access main memory to begin a new task—as opposed to reading the task out of the processor cache, which is common to some multiprocessor architectures.

Cache validity is continuously monitored by specialized VLSI ASIC devices. Validity checking is performed on a duplicate of the physical tag store, thus avoiding cycle stealing. This technique eliminates the interference frequently experienced with other validation methods, and only the cache write “hit” will impact CPU performance.



△  
The Series 10000 features a high-performance, low latency FPU, with a separate ALU and multiplier, and a new hybrid cache design.



## Data Flow Breakthrough: The Compiler

<i>Typical Compiler</i>	<i>New DN10000 Compiler</i>
Parser	Parser
Optimizer	Optimizer
Instruction Selection	Instruction Selection
	Flowgraph Analysis
	Dataflow Analysis
Optimizer	Optimizer
	Scheduler
	Register Allocation
	Optimizer
Object Code Generation	Object Code Generation

Everyone recognizes that a new computer design is in reality only as strong as its compilers. In building the Series 10000 compilers, not only have Apollo engineers rewritten the object code generator, a requirement for RISC-based architectures, but they have also incorporated sophisticated data flow concepts developed at some of the nation's leading research institutes. These improvements have been included in all the standard Apollo compilers.

By paying careful attention to compiler technology and using data flow techniques to map and schedule instruction processing, our engineers have developed an advanced compiler that fully harnesses the latent power of the Series 10000.

### *More Powerful, Still Compatible*

As you can see in the diagram, ordinary compilers have five basic stages. The Series 10000's data flow compiler has five additional sections, which enables it to exploit the inherent parallelism of the architecture. Despite these additions, existing source code for Apollo systems is not affected and present applications are fully compatible with the new compiler.

Data flow techniques have been applied to the local parallel architecture of the Series 10000, resulting in extremely efficient object execution modules. The fact that the CPU performs a multiply, add, and integer operation in the same instruction cycle reduces the inner loop by 3-to-1 compared to conventional sequential architectures.



Five new stages have been added to Apollo compilers to exploit the local parallelism of the new CPU design.

*Data Flow Scheduling*

The heart of Apollo's data flow compiler is the *scheduler*. Its function is to "map" the order in which instructions are executed to the available hardware, and thereby achieve maximum productivity from each machine cycle. It works by transforming the source code into a "data flow graph", finding the critical execution path, and then scheduling the critical path on to the hardware. Because the scheduler can observe instruction latencies and identify the exact instant each instruction should occur, it can make the most efficient use possible of processor elements.

Note in the example on the right that in all steps of the object code, more than one operation was executed in a single cycle. In addition, because data paths are all 64-bit and can be treated as 32-bit integers, the scheduler is able to perform double-integer (32-bit) loads. In the third instruction, three operations are performed simultaneously in one cycle—integer load, single-precision multiply, single-precision add—resulting in reduced execution time, and therefore a reduction in the memory required to evaluate this type of expression. All together, 14 instructions were executed in 7 instruction cycles in the example to the right.

Here's an example of data flow scheduling a ninth order polynomial:


$$4X^9 + 92X^7 + 580X^5 + 824X^3 + 336X$$

It can be restated by the programmer as:


$$Z = X * X$$

$$ZSQ = Z * Z$$

$$Result = (ZSQ + 14.0 * Z + 12) * 4.0 * X * (ZSQ + 9.0 * Z + 7.0)$$



The data flow graph is first created (above) and then scheduled (below) onto the hardware, making note of instructional latencies.

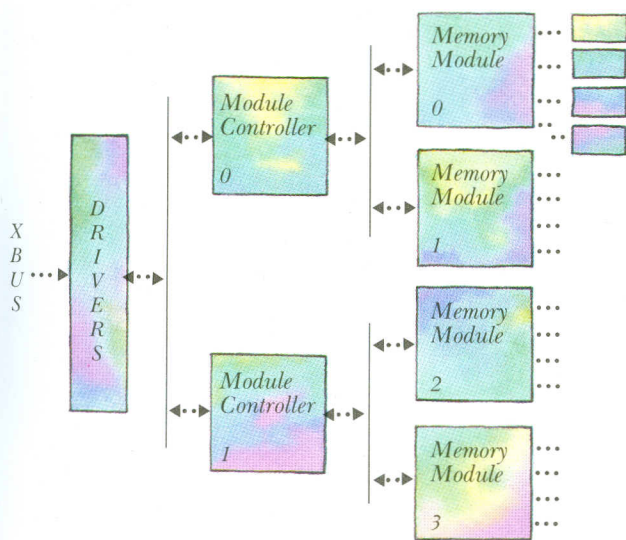


And this is object code produced from the scheduled graph:

Integer Pipeline	Multiplier Pipeline	ALU Pipeline
LOADS C(12.0, 14.0), (S0, S1)	MULS X, X, Z	
LOADS C(4.0, 7.0), (S8, S9)	MULS Z, Z, ZSQ	
LOADS C(9.0), S11	MULS S1, Z, S1	ADDS ZSQ, S0
	MULS X, S8, S5	ADDS S9, ZSQ
	MULS S11, Z, S6	ADDS S0, S1
	MULS S5, S1, S7	ADDS ZSQ, S6
	MULS S6, S7, RESULT	



## High Bandwidth: Main Memory



*Apollo's holistic approach means that, unlike other vendors, our engineers have paid just as much attention to the memory, mass storage, and datapaths as to the CPUs. As a result, the Series 10000 solves one of the key problems in multiprocessor systems: the capacity of the memory bandwidth to handle massive information flow. That's why Apollo engineers made sure the Series 10000 could easily meet the demands of today's high throughput applications.*

Its advanced memory architecture features a parallel design that easily supports heavy I/O traffic over Apollo's high-bandwidth X-bus, as well as quick simultaneous access from each of the multiple processors feeding the bus. The parallel design eliminates the bottlenecks long associated with traditional memory architectures. In addition, the main memory incorporates unique intelligence for managing write buffers, and for increasing bandwidth by using static column RAMs. The benefit of all this is uninterrupted high-speed data flow to and from the main CPUs.

### *Designed for Speed*

The Series 10000 parallel memory is organized into four modules of 100-nanosecond static column CMOS DRAMs. Individual modules boast 8 or 16 megabyte capacity and are physically located on independent daughter boards. Up to four daughter boards can be mounted on each mother board and the entire



*Memory Subsystem.  
The Series 10000 features 1-megabyte DRAM technology to support a 16-way interleaved, 128-megabyte memory.*

set of installed modules may be configured as a contiguous memory space regardless of the mix of daughter boards used or what sites are populated.

Modules are further divided into four interleaved, independently controlled memory banks. This innovative design ensures a bandwidth easily capable of absorbing the full 150-megabyte/second capacity of the databus that links system memory to the multiprocessors.

In addition, simultaneous single *reads* and *writes* are permitted in the independent memory banks with the same module making possible a high bandwidth even on systems with small memories.

Special intelligence is incorporated in the system—via dedicated semicustom CMOS circuits—for controlling the read/write queues of the modules. This intelligence recognizes read operations and automatically advances them to the top of the queue for immediate attention from memory. The ASIC arrays also continuously monitor memory-write addresses to take advantage of static column bursts whenever possible, thus opportunities for fast writes are exploited continuously.

For read multiples and all writes, responses are issued every cycle by ping-ponging between the two bank pairs on a module and exploiting static column cycles. As 8 bytes are transferred on each response once the burst begins, the peak bandwidth of 150 megabytes/second is maintained constantly.

To increase reliability, each memory bank has error correcting code (ECC) for 32-bit data words, allowing correction of

single-bit errors and detection of double-bit errors (SEC-DED).


The Series 10000's main memory is configured with 1-megabit DRAMs, but it's also designed to accept 4-megabit DRAMs as they become available, making future storage capacity as high as 512 megabytes.

#### *High-speed Disk Striping*

The Series 10000 uses sophisticated disk striping to provide the large bandwidth required by a well-balanced system. Disk striping enables files to span multiple drives, allowing high bandwidth access to a file via multiple controllers. It's supported by new high-speed ESDI fast actuator drives with 15-megabit/second transfer rates.

Series 10000 users can stripe two or four drives, with the total data transfer rate depending on how many disk controllers are striped. For example, two striped disks double the data transfer rate of disk-bound applications. Users select how many and which disks are to be striped at initialization. This way disk-intensive files can be put on striped drives and the remaining drives can be used for general purpose applications.

All these innovations are part of the Series 10000's holistic design and they add to the massive flow of information from memory and mass storage at very high speeds. This flow of information can keep the multiple processors constantly busy.




*The system's high bandwidth main memory provides shared virtual memory to all CPUs.*



## *Future Views: Graphics*



*This ray-traced DNA molecule is an example of the Series 10000's advanced graphics applications.*



*The second major implementation of the PRISM architecture is a dramatically new 3-D graphics system which perfectly fits the Series 10000 design. Because graphics are key to today's most advanced applications, the graphics system has unequalled capacity to rapidly generate the full range of engineering and design representations, from 3-D smooth-shaded renderings and complex 3-D wireframe animation to realistic representations of complex computer-aided design structures. In addition, it furthers the state of the graphics art by incorporating features like alpha buffering of translucent objects and texture mapping. These extensive capabilities mean that the Series 10000 will be able to meet the challenges of new graphics applications as they develop in the coming years.*

To accomplish all this, Apollo's approach differs significantly from those of other workstation and graphics terminal makers. That's why the Series 10000 isn't dependent on long, rigid hardware geometry pipelines that are tuned to specific functions while they impede others. Instead, its short and simplified hardware geometry only handles the per-pixel drawing functions used by a wide variety of graphics techniques. And these functions are implemented in blazing-fast pure hardware rather than in



more commonly used micro-code, which eliminates an entire level of interpretation.

All higher-level graphics tasks are handled directly—and efficiently—by the CPU. The fact is, the CPUs deliver transformations equal to, and often exceeding, those of rigid geometry pipelines. What's more, the central processor is specifically designed to process geometries rapidly while remaining a general purpose, user-programmable central processing unit. To ensure users have plenty of computational capacity, the system can be configured with up to four supercomputer-class central processing units.

### High-speed Graphics Subsystem

The Series 10000's Graphics Subsystem features an ultrahigh-speed RISC drawing engine tightly coupled to a deep, heavily interleaved frame buffer. A small set of pixel synthesis operations are implemented directly in VLSI, operating at cycle times twice as fast as those of the CPU.

The subsystem is linked to the CPU via the X-bus, which provides a bandwidth well in excess of the requirements of any drawing operations. It includes a FIFO (first-in, first-out) buffer to smooth burst transmissions from the CPU. And it's mapped as a virtual device to provide the illusion of exclusive ownership of the graphics engine to each program.

### Versatile Frame Buffer

The Frame Buffer Memory is available in two configurations, 40-bit and 80-bit planes. Each offers unprecedented flexibility in assigning groups of planes for application requirements.

For example, the 40-plane version can be configured under software control for a wide variety of images. Among these are true-color images (8 bits red, 8 bits green, 8 bits blue) with a 16-bit Z-buffer for hidden surface removal. When configured for 80 planes, the flexibility increases even more—for instance, allowing double-buffering of true color, 24-bit images, with a full 32-bit Z-buffer.



Each window can have a lookup table and pixel format, and each can be independently double-buffered. As a result, multiple applications can, for the first time, simultaneously and effectively share the screen.

The new graphics implementation means that the Series 10000 can easily handle sophisticated 3-D graphics applications as well as compute-intensive applications. It also gives users access to the raw graphics power that's tightly coupled to the user's program.




### Frame Buffer Configuration Versatility.

Series 10000 graphics make possible a wide range of display formats in each window.





## *Mainframe Reliability: Scan Path Technology*



*Scan path control devices throughout the Series 10000 ensure unprecedented quality and reliability in a workstation.*

*In order to ensure an extremely high degree of reliability in the Series 10000, Apollo has incorporated scan path technology throughout the system. Pioneered by IBM in the 1970s for mainframes, scan path technology provides a built-in system for testing dense VLSI arrays.*

With up to 256 pins and 40,000 gates per VLSI chip, each Series 10000 gate array presents a formidable verification problem. Scan path technology makes it possible to "look" inside each VLSI chip. By shifting the input through the chip, applying an input function, and then shifting to the output, testers can determine its state and functionality.

### *A Testing Chain*

Because all of the Series 10000's semicustom chips have scan path capability, they can be linked together in series to form a single shift chain. This allows built-in test circuitry to analyze functionality with a nearly 100% level of confidence.


The Series 10000's scan path system is controlled by a dedicated 68020 processor located on the service processor board. Each board also has its own Scan Control Resource (SCR) which can isolate failures to the chip level. Communication between the 68020 chip and the other boards is over a private bus so that scan path testing can be done with minimum reliance on the rest of the system.

### *Scan Path Benefits*

The service processor architecture makes possible a number of additional benefits, among which are remote diagnosis and remote debugging across the network, initialization conditioning, and the ability to determine the serial number, revision, and type of specific chips.

Of course, above all, scan path technology gives the Series 10000 an extremely high level of availability and overall quality. It's used throughout the manufacturing process to insure quality at the component, board, and system level—as well as ensuring the quality of the entire product.

*Advanced Ergonomics:  
Personal Packaging*



*The Series 10000 combines supercomputer-class performance with the convenience and comfort of a personal computer.*



*In addition to performance, what makes the Series 10000 unique was Apollo's ability to put its power and speed into a package that fits conveniently and comfortably in an office environment. That meant paying particular attention to the things that many vendors might consider unimportant, but which can make a significant difference to the user.*

To begin with, the dimensions of the Series 10000 are 24½ inches x 20 inches x 29 inches. So not only does it easily fit into the office, it can slide under the standard office table for added convenience. And along with its streamlined dimensions, the Series 10000 is designed so that it can be directly abutted to a wall.

The Series 10000 is also designed for extremely quiet operation. In fact, the system—including both monitor and disk—is virtually noiseless. Among the many things that help achieve this are air-movers to direct air in a straight line, front-to-back, without bending, which dramatically reduces the noise level. The Series 10000 meets all major standards, including the German Environmental Standard (VDI 5058) and AINSI HFS100. The product has been tested in accordance with standards such as ECMA-74, ISO/DIS 7779 Acoustics, ISO/DIS 3744 Acoustics, and ANSI S1.29.

In addition, the Series 10000 incorporates all the features that have made Apollo workstations the standard for ergonomic design, such as a high-resolution screen, 68Hz non-interlaced native graphics, easy-to-use keyboard, and tilt-swivel monitor.




All in all, the Series 10000 ergonomics make it as easy and comfortable to use as a personal computer.

#### *Three-part Packaging*

Series 10000 system packaging is designed in three distinct parts. The lower section houses up to four fast actuator disk drives, as well as the interface for incoming AC power. The center section contains the high-frequency DC-to-DC switching power supply, as well as the main computational processing units. The top section includes the industry-standard I/O card cages for the VME and PC AT expansion boards and the removable media drives for ¼-inch cartridge and 5¼-inch floppy disk drives. Users have easy access to the VME and PC AT chassis to facilitate option installation.

#### *Controlled Cooling System*

The Series 10000's microprocessor-controlled cooling system varies air speed as a function of operating temperature. Air flow volumes exceed 200 linear feet per minute and air temperature is monitored at the main cage card. The VME and PC AT peripheral cages and mass storage subsystems are monitored by a single sensor within the top section of the package. The sensors also detect gross variance



from the recommended operating temperature and shut down the system when activated, protecting the system in case of failure.

Because the Series 10000 is designed to fit comfortably in an office environment, the system's acoustical output does not exceed 55db.

#### *Advanced Power System*

A new computer such as the Series 10000, with such enormous growth potential, requires an advanced power system to fuel it. The Series 10000 is powered by a high-efficiency modular switching power supply. Composed of 150-watt power blocks, the system can be easily upgraded in any office environment. This modular approach is very economical because the power supply can be matched to the actual Series 10000 configuration being used and not to the maximum system configuration, as is usually the case.

#### *Packaging Specifications*

##### *Series 10000 Operating Environment*

Heat Dissipation—5934 BTUs  
(fully configured)

Noise Output—55db (fully  
configured @ 30°C)

Temperature—15°C to 32°C  
(59°F to 90°F)

Humidity—20% to 80% relative  
(non-condensing)

Wet Bulb—25°C max (77°F)

Ceiling—8000 feet (2438 meters)


Vibration—5-22Hz, 0.25 G peak

Shock—5G Peak, 10msec  $\frac{1}{2}$   
sinewave



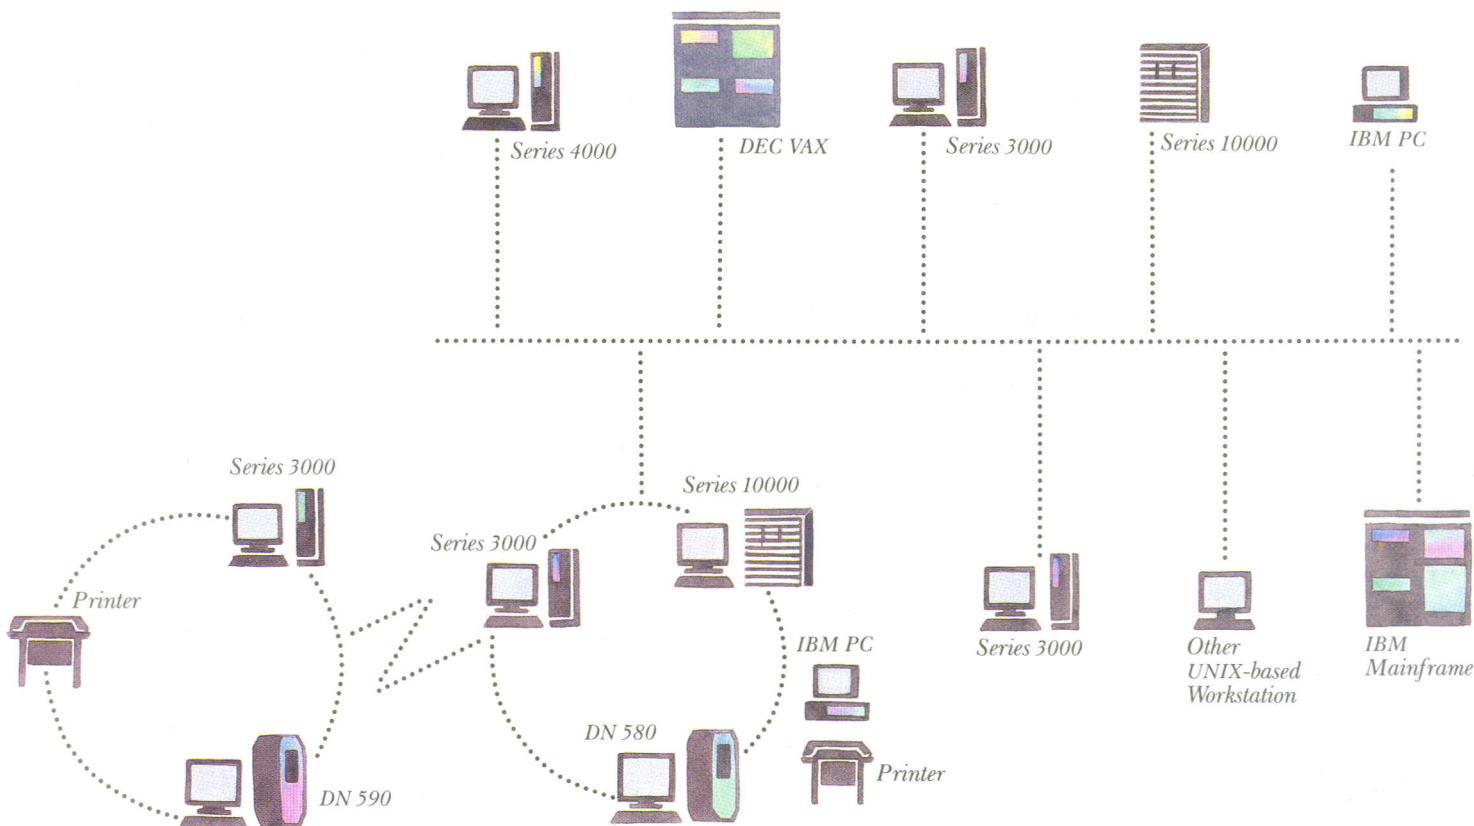
*From the earliest design stages, shown in these 3-D wire-frame representations, the Series 10000 was conceived as a new class of workstations.*





△  
*High-efficiency, modular power supply lets users match power needs to system configuration.*

# *Network Independence: Advanced Communications*



△  
Apollo networks unite the diversity of computers used today into a single transparent system.





*The Series 10000 anticipates the growth of advanced fiber optics networks in the workstation environment.*

*The Series 10000 can dramatically increase the available power of virtually any network. Because, as a fully compatible member of Apollo's Domain<sup>®</sup> family, it can be linked to a wide variety of media—including Apollo Token Ring and Ethernet<sup>®</sup>. And as new standards and technologies develop, such as FDDI (Fiber Optics Digital Data Interconnect), they are being added to Apollo's networking capabilities.*

*Apollo's layered network architecture and communications products make it quick and easy to link the Series 10000 and other Apollo systems, as well as entire Domain networks, to each other and to other networks. These products make it possible for existing networks to take full advantage of the Series 10000.*

*Domain/Bridge.<sup>™</sup> A family of high-speed communications bridges that provide links between Domain rings over a choice of common carriers, including T1, Ethernet, microwave, coaxial cable, and local area network media.*

*Domain TCP/IP.* Provides shared access to existing Ethernet facilities, and supports X.25 transfer files.

*Domain/SNA.<sup>™</sup> Lets users share information and resources with the IBM SNA network and with IBM and IBM-compatible mainframes. And supports 3270 and 3770, and IBM's HASP.*

*Domain/LU 6.2.<sup>™</sup> Extends Apollo's networking capabilities to include IBM Advanced Program-to-Program Communications (APPC) and PU2.1 for peer-to-peer SNA connections.*

*Domain 5080 Emulator.* This emulator lets any Apollo workstation fully simulate the IBM 5080 graphics terminal, and run


any of the hundreds of 5080 applications in a window, while allowing other applications to run simultaneously in other windows. In addition, data from the 5080 window can be merged with files in other windows.

*Domain/PCI.<sup>™</sup> Links users of IBM and IBM-compatible PCs to the distributed Apollo networking environment.*


*Domain/Access.<sup>™</sup> Gateway that supports file transfer, file management, and remote login to DEC VAX/VMS<sup>®</sup> systems, and provides technical professionals access to information stored on these systems.*

*VT100<sup>®</sup> Terminal Emulator.*

Emulates DEC's terminal and supports it through RS236C ports and the Domain Ethernet gateway.



## *The Series 10000: Created for Those Who Demand Performance*



Just eight years ago, Apollo revolutionized the computer industry by introducing the technical workstation. Our workstation changed the way people used computers and how they thought about computing.

Now Apollo is doing it again with the Series 10000 Personal Supercomputer and the new *PRISM* architecture. The Series 10000 brings unprecedented power into the office environment. With multiple independent CPUs, it can easily handle the full gamut of technical applications, whether in traditional areas such as mechanical CAD and electronic design, or in emerging fields like computer-aided molecular design, scientific data visualization, and financial modeling.

The Series 10000 also marks the beginning of a new generation of Apollo computers. Apollo's goal has always been to use advanced workstation technology to improve the productivity of professionals and their organizations—and the Series 10000, with its groundbreaking architectural design, makes it possible for us to do that better than ever before.

### *With all the Apollo Extras*

The Series 10000 may be a new class of computer, but it's available with all the features that make Apollo the leader in workstation technology. Including your choice of native UNIX® 4.3 or System V environments,

Aegis™ NCS (Network Computing System™), and over 1500 applications from more than 750 of the world's leading solution suppliers.

The Series 10000 with its *PRISM* architecture means that you can now get the state-of-the-art in workstation technology and the backing of an established leader in the industry from the same place. Apollo.

Now's the time to find out more about the computer that marks a milestone in our industry. Contact your local Apollo sales representative for more information.

Apollo and Domain are registered trademarks of Apollo Computer Inc. Series 10000, Personal Supercomputer, Domain/Bridge, Domain/LU 6.2, Domain/SNA, Domain/PCI, Domain/Access, Aegis, and Network Computing System are trademarks of Apollo Computer Inc. IBM and PC AT are registered trademarks of International Business Machines Corporation. Ethernet is a registered trademark of Xerox Corporation. DEC, VAX, VMS, and VT100 are registered trademarks of Digital Equipment Corporation. UNIX is a registered trademark of AT&T.

Copyright © 1988, Apollo Computer Inc., Chelmsford, MA 01824.

Screen image on page one courtesy of Intelligent Light, Inc. and United Technologies Research Center, copyright 1987—Intelligent Light, Inc.

Screen image on page fifteen courtesy of National Institute of Health.

The materials contained herein are summary in nature, subject to change, and intended for general information only. Details and specifications concerning the use and operation of Apollo products are available in the applicable manuals, available through local sales representatives.



*By taking a holistic design approach, Apollo has created a new class of workstations, the Series 10000 Personal Supercomputer family.*



*Corporate Headquarters:*

Apollo Computer Inc.  
330 Billerica Road  
Chelmsford, MA 01824  
617-256-6600  
TWX: 710-444-8017  
CABLE: APOLLOCO

*Canadian Headquarters:*

Apollo Computer Inc.  
1530 Markham Road  
Suite 130, Scarborough  
Ontario, Canada  
M1B 3G4  
416-297-0700  
FAX: 416-297-1020

*International Headquarters:*

Apollo Computer, S.A.  
108, Avenue Louis-Casai  
P.O. Box 409  
1215 Geneva, Switzerland  
(41-22) 98 57 88  
TWX: 286 18 ch  
FAX: (41-22) 98 58 79

apollo