SPERRY UNIVAC
Series 1100
Scientific Processing

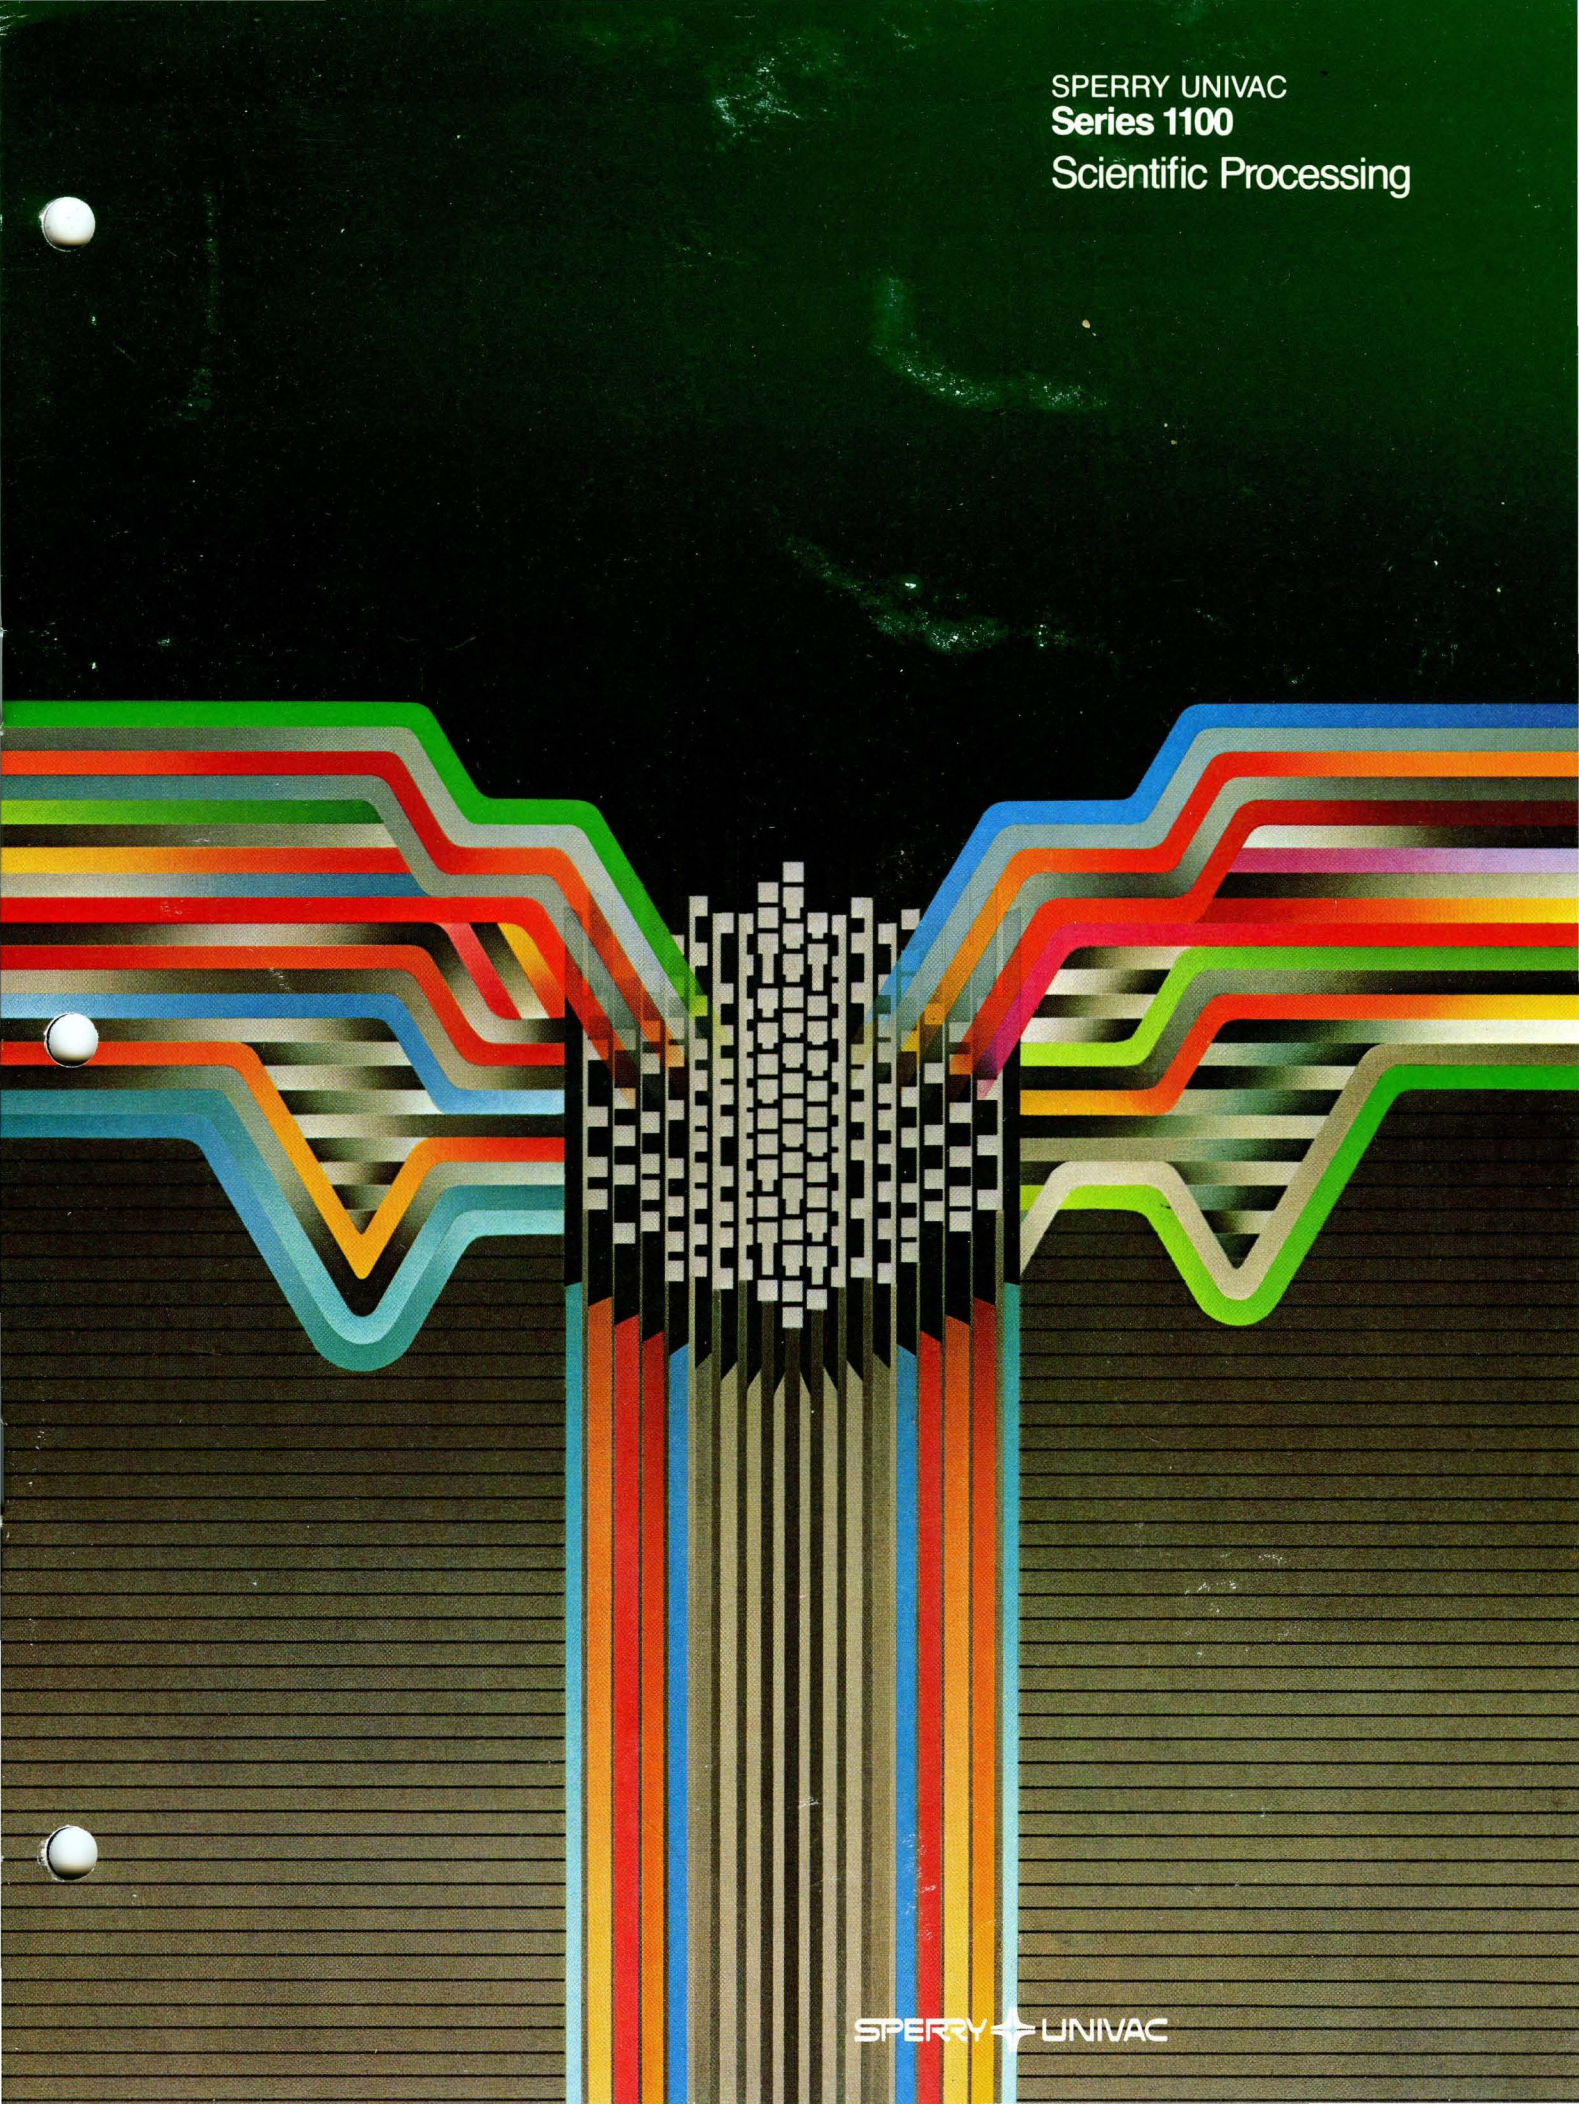

SPERRY UNIVAC

The SPERRY UNIVAC Series 1100 has a comprehensive repertoire of scientific processing capabilities, based on a rich history of scientific processing performance.

Because of that rich history, the Series 1100 has a far-reaching library of scientific applications, including software released both by Sperry Univac and by a multitude of outside vendors. A catalog of this software includes:

- Mathematics
- Statistics
- Operations Research
- Civil Engineering
- Structural Engineering
- Electrical Engineering
- Graphics
- CAD/CAM
- Manufacturing Industry Applications
- Energy

The Series 1100 array processing capabilities can be used in the following scientific application areas:

- Geophysical data analysis
- Reservoir simulations
- Structural analysis
- Nuclear codes
- Electric power flow analysis
- Linear Programming
- Large physical systems modeling

With its scientifically oriented hardware and software—and with its scientific language compilers—the SPERRY UNIVAC Series 1100 returns to state-of-the-art leadership across the full spectrum of scientific data processing.

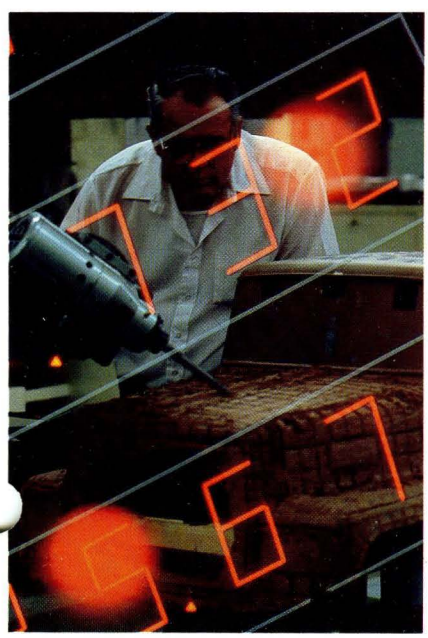

Scientific Accelerators...

To increase throughput in a scientific environment, Sperry Univac offers scientific acceleration modules for the Series 1100 systems.

1100/80 Scientific Accelerator Module (SAM)—The 1100/80 is the largest and most efficient high-volume computer system offered by Sperry Univac. Depending on the processing environment it must accommodate, the 1100/80 can be configured with from one to four central processing units and from one to four input/output units.

An 1100/80 with one central processor (1100/81) has execution power between 1.8 and 2 MIPS (million instructions per second). This means that depending on the mix of jobs, the effective system power of an 1100/84 with four central processors is more than 7 MIPS.

The addition of the Scientific Accelerator Module to each processor can increase single- and double-precision performance from 20 to 35 percent in a predominantly scientific environment.

The SAM is a hardware replacement for the standard 1100/80 processor arithmetic section. It is completely transparent to the system and application software.

Additional throughput increases can be realized with the addition of array processing units.

1100/60 Double Precision Microcode Accelerator—The 1100/60, like the 1100/80, can include up to four input/output units. In a scientific environment, the MIP rate ranges from 0.6 (one processor) to more than 2 for an 1100/64 (four processors). The 1100/60 DP Accelerator increases throughput 20 to 30 percent for scientific runs where double precision dominates.

SCIENTIFIC ARRAY VECTOR PROCESSORS...

In addition to host-oriented accelerators, Sperry Univac offers a number of array processors to increase scientific throughput.

Three array processor options are available for the Series 1100. Two of these—using array processors from Floating Point Systems, Inc. (FPS)—may be attached through the word channel to both the 1100/60 and 1100/80. These two attachments are the FPS-190L and FPS-164. The integrated SPERRY UNIVAC Array Processing System (APS) applies only to 1100/80 models.

SPERRY UNIVAC/FPS SYSTEM FEATURES...

A number of mainframes use FPS to enhance scientific performance, but The SPERRY UNIVAC Series 1100 is unique in that it allows up to six concurrent users to access the FPS-190L. Other systems using the FPS-190L typically permit only one person to use the array processor (AP) at a time, prohibiting the AP from reaching its true potential.

In a typical environment, SPERRY UNIVAC/FPS-190L systems would run three times as much work as a system which permits only one AP user at a time.

The SPERRY UNIVAC/FPS-164 system also utilizes a multi-user approach, allowed by both the FPS-164 and the Series 1100 system. The FPS-164 is a more sophisticated unit than the 190L, containing a microprocessor and its own monitor system called the Multi-user Monitor (MUM). Most non-Sperry Univac implementations of the FPS-164 use a Single User Monitor (SUM). The 1100-based AP control system is a modification of FPS's APEX and allows connection of multiple FPS-164s as well as multi-users under control of one control program.

The SPERRY UNIVAC 1100/FPS-164 is designed for extended-precision (64 bit) accuracy. Sperry Univac's 36 bit, binary normalized, single-precision capability often provides adequate precision where 32-bit hexadecimal normalized floating-point precision is not sufficient.

1100/60 or 1100/80 with FPS-190L—This is a 38-bit, single-precision floating point system which can be attached to single or multiprocessing 1100/60 or 1100/80 systems. It is connected via the word channel of the 1100 system. One or more FPS-190Ls may be used.

The FPS-190L has a peak performance of 12 million floating point operations per second (MFLOPS or megaflops). The FPS-190L operates with the 1100 by means of user or system invoked primitives. Today, more than 30 FPS-190Ls are operating using the Series 1100 systems.

1100/60 or 1100/80 with FPS-164—This system is similar to the FPS-190L discussed above. However, it is a 64-bit rather than a 38-bit system. It is also attached to the Series 1100 through a word channel.

The FPS-164 operates at a peak rate of 12 MFLOPS, producing 64-bit results for an equivalent 15-decimal digit accuracy. The FPS-164 also works with the Series 1100 by means of user or system invoked primitives.

The 1100/80 APS is Sperry Univac's most powerful scientific configuration. The SPERRY UNIVAC 1100/80 APS is an integrated system combining both scalar and vector processing capabilities in a multiprocessing environment. See Figures 1 and 2.

The total 1100/80 APS is under the control of the standard 1100 Operating System. By using combined aspects of parallel and pipeline processing in the vector units of the APS, the user can realize a sustained rate of 80 MFLOPS for certain operations in each of the two possible array processor configurations.

The maximum designed throughput for each 1100/80 APS is 120 MFLOPS. The highest burst rate of the combined vector and scalar processing capabilities of the APS in a maximum 1100/84-APS(2) configuration (four 1100/80 central processors and two array processors) is 245 MFLOPS. See Figure 3.

The high throughputs are achieved by use of four pipelined array arithmetic processors, four control processors a local scratchpad and instructional memories as shown in Figure 4.

The array processing units are directly attached to the 1100/80 real memory (see Figure 3). Thus the APS has direct access to all of the 1100/80 real memory, which can be configured with up to 8 million words. Each APS is interfaced to main memory by means of a high-speed cache memory.

During vector operations by user programs, required data is migrated from main memory to the APS cache. The cache size is 16K words and the cache has a transfer rate of 35 million words per second.

Most importantly, the data bandwidth between the vector units and cache is 40 million words per second. Data is moved into the APS cache in increments of 64-word pages. There are several pre-fetching techniques used to insure that the user's required data is in the cache.

Because of its combined system speed and efficient data accessing methodology, the 1100/80 APS is ideally suited for a broad range of scientific processing applications. It is particularly suited for processing environments characterized by intense numerical calculations, such as seismic processing and oil/gas reservoir simulations.

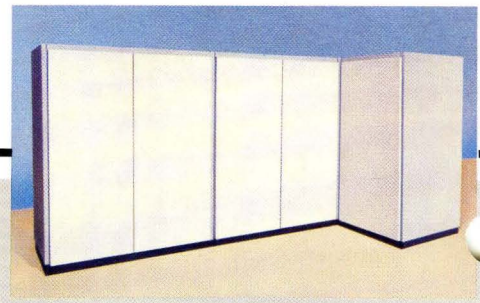

Figure 1. Array Processor Subsystem External Appearance

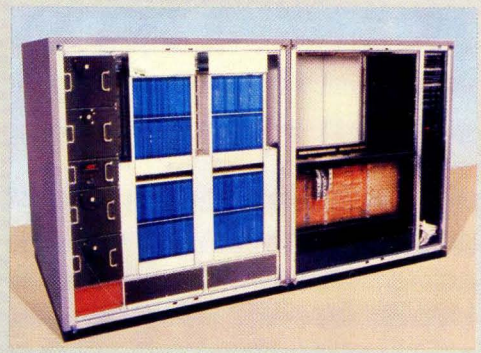

Figure 2. Array Processor Subsystem Internal View

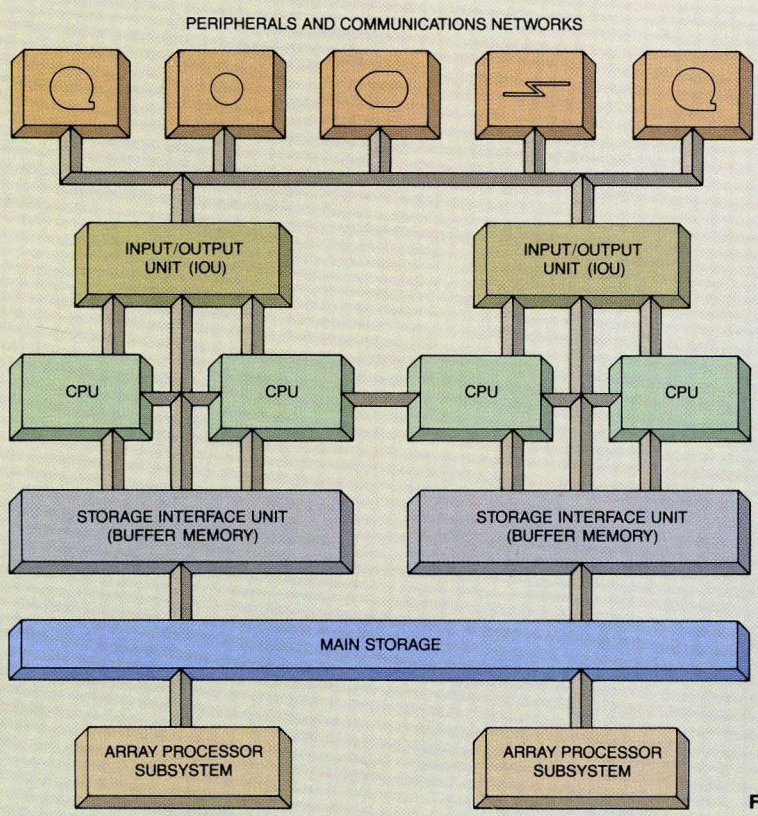

Figure 3. 1100/80 Host System Architecture (Maximum Configuration)

VAST™—VECTOR AND ARRAY SYNTAX TRANSLATOR...

VAST is a software tool designed to make array processors easier to use by making them accessible via standard FORTRAN programs. With VAST, users can write programs in standard FORTRAN without explicit CALLs to APS subroutines.

VAST is a preprocessor which analyzes FORTRAN programs. It will convert certain "DO loops":—where vectorization is possible—into calls to array processor primitives. See Figure 5.

VAST not only invokes primitives, but also can chain groups of primitives (vector operations) with interspersed scalar operations.

User communication with VAST includes both the diagnostics by which VAST sends messages to the programmer and the user directives by which the user directs VAST to take certain actions. These features let the user know exactly what has occurred and control the translation process as desired.

The initial release of VAST supports the SPERRY UNIVAC 1100/80 APS. Forthcoming versions of VAST will support the FPS-190L and FPS-164 when attached to SPERRY UNIVAC Series 1100 systems.

VAST gives the array processor user several advantages. Primarily, it makes the specialized capabilities of the APS hardware available to the scientific programmer after only minimal training. A detailed understanding of the hardware and access methods is not required.

In addition, long-term costs are reduced since programs are retained in standard FORTRAN rather than in a syntax unique to the hardware. This enhances transportability, reduces the need for conversions and eases debugging.

The VAST package for the 1100/80 APS consists of the VAST precompiler, a run-time chainer, an interpreter, and an extended primitive library for the APS. The VAST precompiler accepts FORTRAN input and produces a vectorized output FORTRAN program suitable for processing with the Series 1100 FORTRAN compiler—along with a list analyzing the circumstances where vectorization could not be performed.

Since VAST is a preprocessor, programs written in FORTRAN may be used for development and maintenance. The user may also revert to source with his own calls, with vector-oriented calls or with any desired combination of the two.

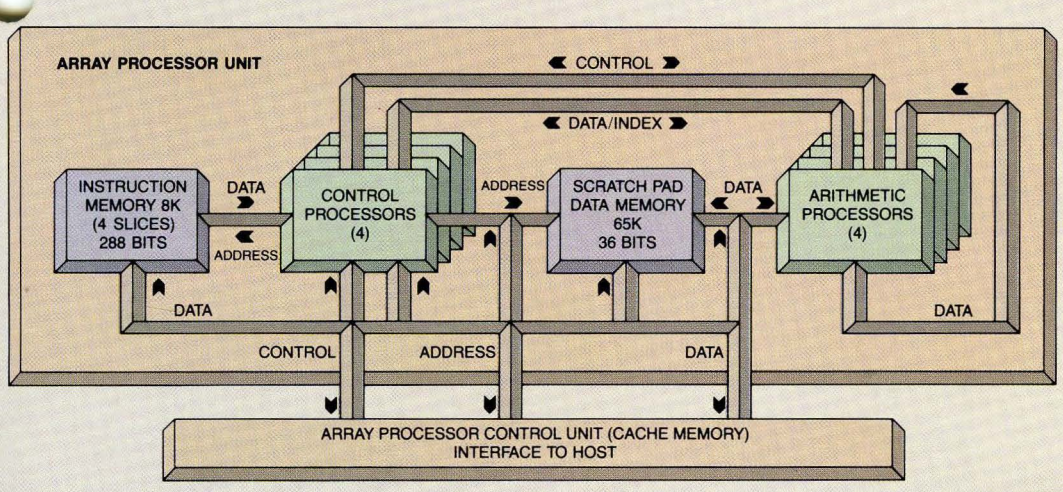

Figure 4. Array Processor Unit Architecture

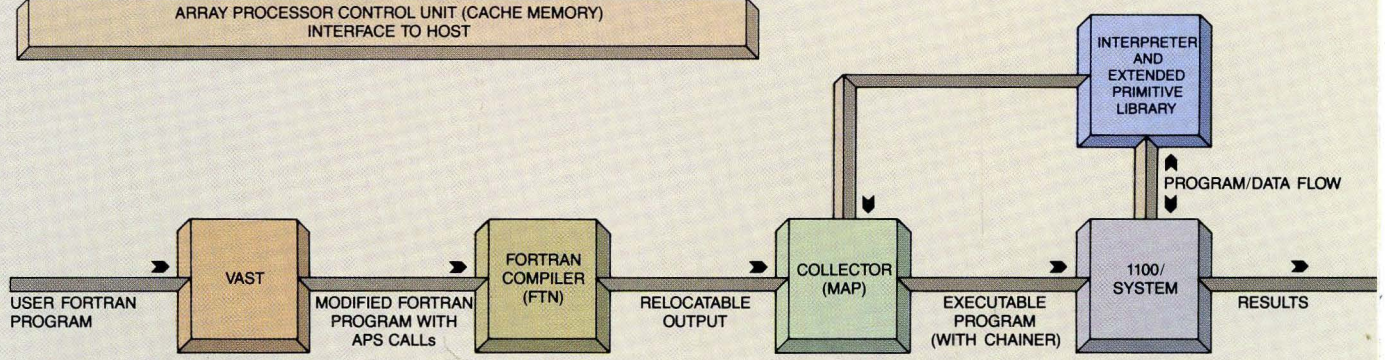

Example:

```
DO 10 I=1,100                    CALL VVT (A, B, C)
  A(I)=B(I) * C(I)               CALL VVP (X, Y, Z)
  X(I)=Y(I)+Z(I)               10 CONTINUE
10 CONTINUE
```

Figure 5. Work Flow for VAST

VAST is a trademark of Pacific-Sierra Research Corp.

All the array processors offered may be invoked by means of explicit calls under the FORTRAN compiler. This capability represents a second way to invoke the array processor.

More than 50 APS primitives beyond the basic algorithm library have been developed and are available (see Figure 6). A VAST extended library is also available.

Name	Operation
FFT CLASS	
FFT	Basic FFT Butterfly
IFFT	Inverse FFT Butterfly
UNSC	Unscramble to Natural Order
FFTN	FFT Normalization
CMCJ	Complex Conjugate
SRCB	Single Real Coefficient Builder
ISRC	Inverse Single Real Coefficient Builder
COEF	Two Real Coefficient Builder
ICOE	Inverse Two Real Coefficient Builder
CONVOLUTION/CORRELATION CLASS	
CONV	Convolving Multiply
CONA	Convolving Addition
CORR	Correlation Multiply
CORA	Correlation Addition
VECTOR REDUCTION CLASS	
VECP	Vector Inner Product
SSSQ	Sum of Signed Squared Vector
SUMR	Sum Reduction
SMSQ	Sum of Squares
VECTOR ELEMENT BY ELEMENT CLASS	
VECA	Vector Add
VECM	Vector Multiply
CPXM	Complex Vector Multiply
VECS	Vector Squared
VESS	Signed Squared Vector
VMAG	Vector Magnitude
VNEG	Vector Negative
VNMG	Vector Negative Magnitude
PARM	Partial Matrix Multiply
LFnn	Logical Function nn where $0 \leq nn \leq 15$
UTILITY CLASS	
SVH	Scan Vector High Value
SVHM	Scan Vector High Magnitude
SVL	Scan Vector Low Value
SVLM	Scan Vector Low Magnitude
EXCV	Exchange Vector
COPY	Copy Vector
CLRM	Clear Memory

Figure 6. Basic Algorithm Library

PRIMITIVES BY USERS...

All array processors let the user write his own microcoded primitive operations. The 1100/80 APS allows primitive development and testing by means of a cross assembler, library builder and APS simulator. This represents a third way of using the array processor.

LARGE ARRAYS USING FORTRAN...

The multibanking capability of the Series 1100 makes the program size of instruction areas virtually unlimited. The data array is developed without the need for user intervention. Paging, when appropriate, is accomplished by the system. In the interests of good memory management, the system develops the page size rather than depending on standard page sizes.

Series 1100 systems carry up to 8 million 36-bit words. In many cases it is advantageous to have large arrays in memory to prevent "thrashing." The user can accomplish this by executing a large array in memory, or by making the pages larger than normal. The Series 1100 system can also restrain certain users from misusing this feature.

A FINAL NOTE ABOUT 36-BIT ACCURACY...

The Series 1100 36-bit word is acceptable in the scientific environment where in many cases the 32-bit word is not.

The SPERRY UNIVAC Series 1100 single precision, 36-bit floating point representation offers two distinct and numerically important advantages versus the hexadecimal 32-bit single precision format, providing:

1. 27-bit precision in the mantissa versus 24 bits.
2. Binary normalization versus hexadecimal normalization.

These two factors result in an eight digit precision for the Series 1100 single precision result versus six digits for the 32-bit format—a 33.3 percent increase in precision.

Direct solution methods are used widely in scientific processing, particularly to invert matrices in such areas as reservoir simulation where large matrix inversion operations are required.

The scientific numerical deficiency of the 32-bit floating point operations means that 32-bit machines must operate in *double precision* for many problems where the Series 1100 can operate in *single precision*. This restriction on 32-bit machines causes *longer execution times* and *double* the memory requirements for data array storage.

The SPERRY UNIVAC Series 1100 systems had their beginnings in the early 1950s. They were created to fill scientific processing needs. Today the Series 1100 remains unexcelled in its scientific processing capabilities—while it has grown in versatility.

Early Series 1100 systems were used strictly for scientific applications. With the 1108, the capabilities of the Series were expanded to include general-purpose processing. Enhancement of those capabilities has continued ever since.

The 1100 Operating System, which made these enhancements possible, was introduced in 1965 on the 1108 system. Since then it has evolved into the single operating system for all members in the 1100 Family.

The exceptionally powerful 1100 OS accommodates the concurrent execution of computer-limited programs, batch-I/O limited programs, interactive timesharing and transaction-oriented jobs.

Not only can the operating system multiprogram mixtures of such jobs and programs, it can do so in a multiprocessing environment. It can also effect multiprocessing with similar tightly coupled systems and with hosts and special processing systems.

The continuing development and improvement of hardware and software in the Series 1100 has always been at the forefront in comprehensiveness and efficiency—and the recent introduction of the Interactive Processing Facility (IPF) represents a new state-of-the-art in distributed timesharing systems.

Because of its many attributes and capabilities, the SPERRY UNIVAC Series 1100 is much more than a family of excellent freestanding scientific and timesharing systems. It is also a comprehensive line of general-purpose systems which can concurrently run timesharing, scientific and commercial applications in a single environment.

The complete capabilities of the Series 1100 are detailed in other Sperry Univac literature. Contact your Sperry Univac representative.


PET 000:00:09:13
DAY HR MIN SEC
LAUNCH PHASE (RLS)
SURGEON
RMS

SPERRY UNIVAC

The computer people who listen.

U7171

8206 Printed in U.S.A.