November 17, 1966

Functional Objectives for an Array Processor

Introduction

The purpose of the array processor is to significantly reduce the time required to perform a set of arithmetic operations on each set of elements of the input arrays.

The basic arithmetic capabilities of the array processor must include multiplication and addition and these operations should be performed in short form (32 bit) floating-point notation. In addition to floating-point input, the processor should be capable of accepting halfword (16 bits) fixed point input in either 2's complement or signed/true format. A conversion from floating-point to fixed point should also be available under special control.

Among the basic array operations that must be performed by the array processor are convolution, correlation, matrix multiplication, and matrix addition or subtraction. The processor should also be capable of expansion by means of an optional feature to include recursive filtering. Other optional features that may be desirable include the ability to do the fast Fourier transform and third and fourth order correlations. Mathematical descriptions of the basic array operations and recursive filtering are given in a following section.

System Configuration

The array processor must be capable of attachment to System/360, models 44, 65, 75, and 85. Initial discussions with SDD indicate that attachment should be made directly to the channel bus (64 bits) on M65, M75 and M85 and to an RPQ channel bus (32 bits) on the M44. On the three larger systems the array processor should not be allowed to seriously interfere with the operation of the channels nor should it be allowed to block more than about half of the memory accesses by the CPU. On the M44, the array processor should not seriously interfere with the channels but could be permitted to block CPU access to memory whenever necessary.

In order to be compatible with and take advantage of the basic architecture of System/360 it has been proposed in the initial discussions with SDD that the execution of operations in the array processor be initiated by a Start

I/O (SIO) instruction and make use of the Channel Address Word (CAW) and the Channel Command Word (CCW) to supply the necessary control information to the array processor. This sequence of operations is illustrated by the following diagram.

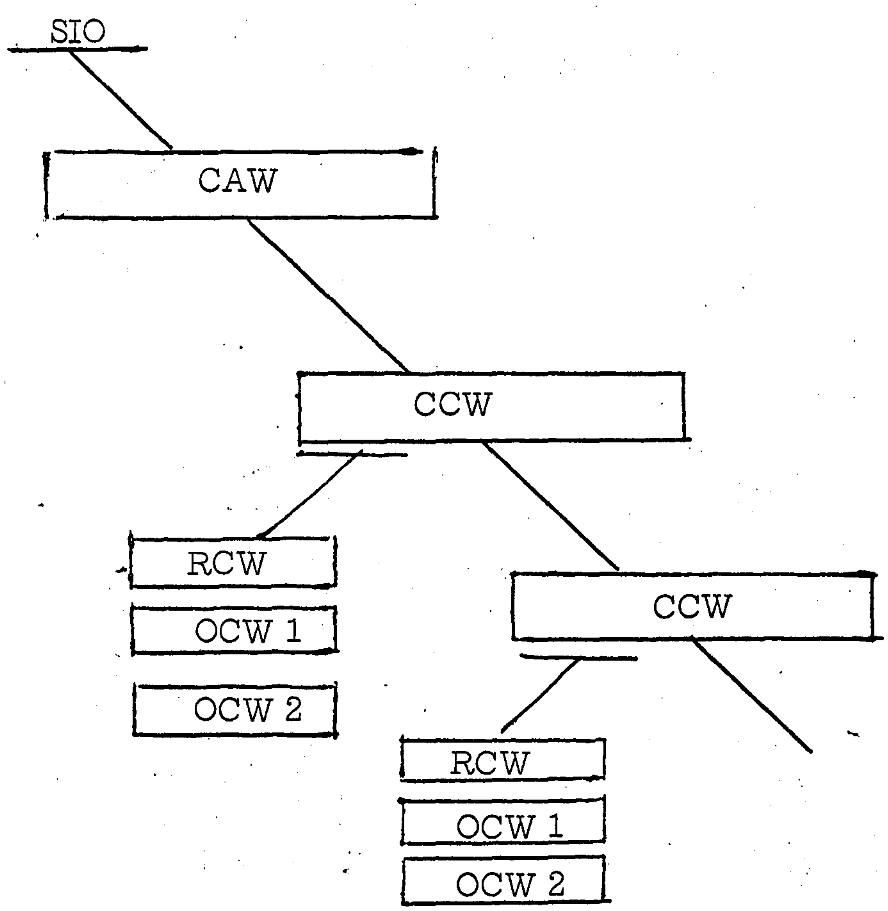

The OCW and RCW are defined as follows:

Operand Control Word (OCW)

Control word format for the two operands.
One control word required for each operand.

Format	Starting Address	Length	Index
0 7	8 31	32 47	48 63
8	24	16	16

Format
1. Fixed or Float
2. If fixed, signed or 2's complement
3. Algebraic or absolute value

4. Forward or reverse indexing
5. Data string or one word constant
6.
7.
8.

Result Control Word (RCW)

Operation code and result control word
All results in 32 bit floating-point format, hexadecimal notation, except Fl. Pt. to Fxd. Pt. Convert and Move operations.

OP	Starting Address	Length	Index
0 7	8 31	32 47	48 63

Since the operands are obtained from main memory and the results are returned to main memory, the starting address of an array is the location in main memory of the first element of the array. The length is the number of elements contained in the array. The indix is the memory spacing between successive elements of the array. Thus, logically adjacent elements in an array need not be stored in adjacent memory locations.

Arithmetic Operations

The following arithmetic operations comprise the basic set required for the processing of petroleum seismic exploration data.

Convolution or Correlation

$$y(j) = y'(j) + \sum_{i=1}^{m} u(i)\, x(i+j) \qquad j = 0, n$$

$y'(j)$ is a prior value to which the summation is added.

Partial Matrix Multiplication

One row of the first matrix by all columns of the second matrix

$$y(j) = \sum_{i=1}^{m} u(i)\ x\ (i+mj) \qquad j = 0, n$$

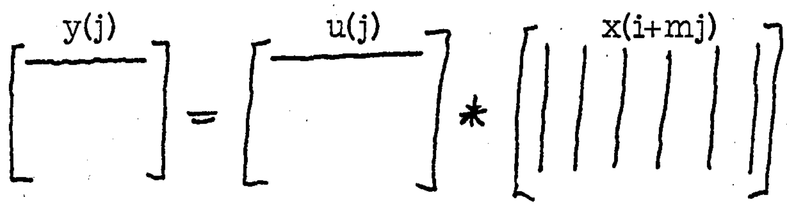

Inner Product

$$y(j) = y'(j) + \sum_{i+1}^{m} u(i)\ x\ (i) \qquad j = 0$$

This is a special case of convolution where the operand lengths are equal and the resultant length is one.

The following three equations are closely related to the Inner Product.

Sum squared array

$$y(j) = y'(j) + \sum_{i=1}^{m} u(i)u(i) \qquad j = 0$$

Sum array

$$y\ (j) = y'(j) + \sum_{i=1}^{m} u\ (i) \qquad j= 0$$

Sum absolute array

$$y(j) = y'(j) + \sum_{i=1}^{m} \left| u(i) \right| \qquad j = 0$$

"Convolving Addition"

$$y(j) = y'(j) + \sum_{i=1}^{m} \left| u(i) \pm x\ (i+j) \right| \qquad j = 0, n$$

This operation is similar to convolution except that addition (subtraction) is performed on the two operands rather than multiplication.

Vector or String Multiply

$$y(j) = y'(j) + u(j)\,x(j) \qquad j = 0, n$$

Signed, Squared Array

$$y(j) = u(j) \times |u(j)| \qquad j = 1, n$$

Vector Sum or Difference

$$y(j) = u(j) \pm x(j) \qquad j = 0, n$$

Scatter Move

$$y(j) = u(j) \qquad j = 0, n$$

In all of the above arithmetic operations that include the prior value $y'(j)$ in this description, it is necessary that the operation be executable with and without the inclusion of $y'(j)$.

Fix to Float

This operation is to be done automatically on all arithmetic operations that use halfword fixed point variables as input.

Float to Fix

This is a special case of vector add where $x(j)$ is specified as a constant to be added to all $u(j)$ with $y(j)$ being stored as a fixed point halfword, signed/true or two's complement (high order 15 bits of resultant fraction)

The following set of equations represent the second-order difference equation used in implementing the recursive filter.

$$y(k) = y'(k) + a_0 u(k) + a_1 x_1(k) + a_2 x_2(k)$$

$$x_1(k+1) = x_2(k)$$

$$x_2(k+1) = u(k) + b_1 x_1(k) + b_2 x_2(k)$$

$$k = 0, n$$

$$x_1(o) = X1 \qquad\qquad x_2(o) = X2$$

$u(k)$ is the input array.
$y(k)$ is the output array.

a_o, a_1, a_2, b_1, b_2 are constants.

X1 and X2 are initial values.

$x_1(n)$ and $x_2(n)$ should replace X1 and X2 at the completion of the operation

$x_1(k)$ and $x_2(k)$, for $0 \quad k \quad n$, are intermediate results and are not part of the output.

a_o, a_1, a_2, b_1, b_2, X1, X2 can be stored sequentially in any order suitable to the array processor.

DPD

Scientific Marketing Newsletter

Special Issue No. 77
August 1967

To be used by IBM
personnel only

THE 2938 ARRAY PROCESSOR

Large-scale processing requirements are characteristic of the scientific marketplace. Satisfaction of such requirements often elicits considerable ingenuity. The invention of logarithms, for example, was an early, ingenious technique for satisfying the computational demands of astronomy and navigation by carrying out the multiplication and division operations through addition and subtraction.

The invention of the digital computer itself was a response to the stimuli of the scientific marketplace. (See The Analytical Engine: Computers - Past, Present and Future by Jeremy Bernstein, Random House, New York, 1964). Subsequent inventions have increased its effectiveness, decreased its cost, and made it easier to communicate with and to operate.

More recent projects aimed at bringing growing scientific computing requirements under control are programs such as the System/360 Remote Access Computing System (RAX) and the System/360 Attached Support Processor System (ASP), algorithms such as the Fast Fourier Transform (see Special Issue No. 74), languages such as PL/I (see Program Announcements P67-63 and P67-71 for the recent DOS/360 and TOS/360 availability announcements), and whole libraries of computational techniques such as the System/360 and 1130 Scientific Subroutine Packages.

Some requirements are so large, however, that special devices are needed to satisfy them, if a suitable cost/performance ratio is to be maintained. This issue is about one such device, the 2938 Array Processor.

SCI-67-032
8-7-67

Scientific Marketing Programs
IBM DPD Headquarters
White Plains, New York

THE 2938 ARRAY PROCESSOR

The 2938 Array Processor is a peripheral processor for attachment to System/360 Models 44, 65, and 75. It performs a set of arithmetic and data format conversion operations on one, two, or three arrays of input data to produce an array of output results. The 2938 attaches to the system in the same manner as an I/O channel. Thus, it is activated by the Start I/O, Test I/O, Halt I/O and Test Channel instructions and utilizes Channel Address, Status, and Command words with the usual System/360 significance. Once operation of the 2938 has been initiated by a Start I/O instruction, it fetches its input data from storage, performs the desired operation, and returns the results to storage without further instruction from the CPU program. Upon completion of its operations the 2938 presents an interrupt request to the CPU. Thus, the 2938 appears to the CPU to operate in essentially the same manner as a high speed channel.

The arithmetic unit of the 2938 is a multiplication and addition unit which performs the operation $Z = U*X \pm Y$. The unit operates on System/360 short-operand floating-point data only. However, the Array Processor has the capability of converting between halfword fixed-point and short-operand floating-point data. The organization of the arithmetic unit is similar to that of an assembly line in that it is possible to have several partially complete results in process through the unit at one time. When the data are supplied to the arithmetic unit at its maximum input rate, the combined multiply and add time is effectively 200 nanoseconds, whereas the total time for a particular set of U, X, and Y to pass through the arithmetic unit to form the resultant Z is approximately 800 nanoseconds. The overall rate at which the Array Processor can perform a given operation is dependent not only on the basic speed of the arithmetic unit but also upon the rate at which data can be transmitted to and from the arithmetic unit.

The 2938 obtains its initial input data from and returns its final results to the main processor storage. It also has two sets of 32 fullword registers of its own in which data or intermediate results can be retained. These registers are capable of transmitting data to or from the arithmetic unit at the maximum rate of that unit. Since the 2938 is attached to the channel bus of the CPU, the rate at which the 2938 can get from or put data into processor storage is dependent upon the rate at which the channel bus operates and the amount of service required from this bus by the channels attached to the system. An accurate set of rules for estimating the data rates for the three different CPU's and for various channel activities is not available at this time. A preliminary estimate of the maximum data rate, assuming no channel activity on the system, might be one doubleword every 3.0 microseconds on the System/360 Models 65 and 75 and one fullword every 2.0 microseconds on the System/360 Model 44. Thus, if the particular mathematical operation being performed can make effective use of the two sets of registers in the Array Processor, the average multiply and add time for this operation can be expected to approach the effective multiply and add time of the arithmetic unit. On the other extreme, where little use can be made of the registers, the average multiply and add time will be a function of the rate at which the data can be obtained from processor storage.

The 2938 provides twelve array operations as standard features. These operations perform array moves, format conversions and vector or matrix arithmetic operations. The fourth-order difference equation operation is available as a special feature. Control of the operations performed by the Array Processor is directed by a microprogram in its read-only storage.

SECTION II - MARKET ENVIRONMENT

A. ~~Assumptions~~ Introduction

There is a growing trend among computer vendors to satisfy the needs of large scientific customers with special purpose devices. The use of these devices, particularly convolvers, has had a long history in seismic petroleum exploration. There are several seismic computing systems in competitive product lines. IBM's 2938 Array Processor has greatly enhanced our position in the seismic market. The competitive large general purpose computers are consistently outperformed by the 2938 Array Processor. The Scientific Multi-Processor enables IBM to substantially increment the performance of other high-compute applications at a small cost increment to the user. This product enables IBM to compete effectively in the large scientific market by outperforming the other fellow with the Scientific Multiprocessor while he must propose larger, more expensive general purpose computers.

A study of the scientific market was completed by DPD HQ Scientific Marketing early this year. This market represents 21% of IBM's revenue. It is growing at an annual rate of 24% in contrast to 18% for all of commercial. While the scientific market will triple by 1972, IBM's share of the market will decline from 56% to 48% by 1972.

Further analysis revealed some more telling numbers. This market is comprised of 80% job shops and 20% high-compute users. Of the 1968 revenue, the high-compute sector accounted for approximately 9 million of the 44 million points. By 1972, competition will have saturated 84% of the long job high-compute market.

B. Objectives and Strategies

The major objective of the Scientific Multi-Processor is to gain re-entry into the high-compute market. Both product design and market strategies are oriented around this theme. The SMP has some applicability to the scientific job shop but it must be marketed carefully because performance will fluctuate widely with specific job mixes. There is little applicability for the product outside what are traditionally termed scientific or engineering applications. To meet the major objective, market strategies have been formulated. First of these was attachment to a wide range of CPU's. Contemplated are the A48, the 553, the Model 85, the 146, FS-4, and FS-3. Implicit in this strategy is the willingness to tolerate highly variable performance from system to system. Performance is directly contingent upon memory speeds. The slower the memory, the lower the performance of the SMP. The memory speeds of the A48 and 553 make them susceptible to degradation.

Also implicit in this wide range of attachments is the capability of SMP to transmigrate. As the scientific user migrates (e.g., from an 85 to a 148 to FS-4), it is anticipated that he would retain the same SMP. The attachment strategy extends the product life of SMP into the FS period. Upon this strategy rests our assumption of 50 month average rental life.

A second market strategy was to design and implement an architecture which is easily expandable. Expandability is considered in two fashions. First, we must provide an easy method of adding tailored algorithms as RPQ's. Our 2938 experience indicated that this sector of the market is prone to making modifications to the standard algorithm repertoire. Writable control store is one attractive implementation alternative for RPQ's. The second aspect of expandability addresses itself to the computing needs in the 1975-1980 time frame. By defining the scientific high-compute market needs accurately, it is anticipated that further expansion of this architecture would accomplish suitable performance rates for the scientific user of 1975-1980. Whether the design is implemented as a multi-processor or as additional function within future CPU's is immaterial. The important element is expansion of a properly defined market need.

A third market strategy was to provide IBM with a product which competes with other vendor devices designed along these general architectural lines. The most serious threat is the CDC STAR computer which is a complex computer system whose performance is estimated to be in the 100 million instruction per second range. Purchase price has been quoted at approximately $10,000,000 (200,000 points). Other vendors have also been active. SDS has recently contracted with Mauchly Associates to provide a 2938-type device in the process industry. They also have submitted an RPQ for a 2938 for attachment to a SIGMA 5. Raytheon recently announced its Array Transform Processor which is another special purpose CPU. Remington Rand has committed an array processor product to a major oil company. For more information on these competitive products, please see the Competitive Section of the Forecast Assumptions. It is safe to say that we are not the only vendor to have identified the inherent superiority of this approach in the high-compute scientific market.

Considered either from the aspect of performance or more attractively from the aspect of price/performance, we effectively compete with competitive hardware in the applications described below, except for STAR. To compete with STAR, IBM will require a higher speed SMP. It is important to note that architecture and market definition do not change. They are designed to be expandable and we have a 200,000 point umbrella in the STAR's.

III. Product Description

The intent of the product description section of the Scientific Multi-Processor Forecast Assumptions is to define Scientific Multi-Processor (SMP) and to provide marketing implications of each design improvement. Secondly, throughput in particular applications areas is projected based on the 2938 Array Processor performance in specific applications which are being performed today, and the SMP performance (in particular algorithms) is extrapolated from 2938 measured times compared to CPU timings. Thirdly, this section will discuss competitive activity in these market sectors.

A. Introduction

The Scientific Multi-Processor is a peripheral high-speed numerical algorithm processor. The device is partly a follow-on to and extension of the 2938 Array Processor in a conceptual sense, but differs radically in the design and range of applicability. The 2938 A.P. is an RPQ device built for seismic data processing in which a set of vector/matrix arithmetic operations and numerical algorithms have been microprogrammed.

The intent of the Scientific Multi-Processor is to extend the concepts originated in the 2938 to impact a number of applications in Industries not directly affected by the 2938 and for which the 2938 is not suited. By correcting the deficiencies of the 2938 Array Processor and adding additional algorithm structures, the SMP will address a significant number of new applications for which high performance is essential on a cross industry basis.

B. Functional Improvements

The 2938 Array Processor and the proposed Scientific Multi-Processor are considerably different, both in hardware and marketability. The 2938 has met approximately 50% of its initial forecast. It has not succeeded in penetrating in significant numbers any industry but Process nor any application but seismic data processing.

There are real and identifiable reasons for the 2938 market limitations.

The 2938 has only single precision arithmetic capability, and the lack of double precision has severely impacted its penetration into the general scientific marketplace by excluding a large number of applications.

It attaches to the CPU through a standard channel interface which severely degrades performance by making all but one of the algorithms implemented data access bound.

The original set of algorithms microprogrammed on the 2938 were directed towards the seismic problems and were in fact conceived by seismic customers. There was no intent to expand these algorithms into general scientific computations. As a result of this, new algorithms were RPQ'ed for the 2938 even within the seismic marketplace.

No provisions were made for the manipulation of sparse matrices, and the device does not have division; or logic capability. The 2938 was an RPQ and not a standard product. The field marketing force has not been exposed to the 2938, as it would be to the SMP as a natural consequence of a planned program. There is little education for the field force outside of the Process Industry, and the 2938 Array Processor is in itself a rather difficult piece of equipment to understand.

The objective of the Scientific Multi-Processor is to correct each of the above deficiences as detailed below.

The Precision Limitation

In the SMP both single and long precision arithmetic (including divide capability) will be available. In addition the accumulation of all sum reductions, particularly those encountered as inner products, will be in extended precision. Long precision (64 bit) arithmetic with extended precision (128 bit) accumulation is a highly desirable feature for all large scale scientific computation - particularly in the solution of partial differential equations encountered in the nuclear marketplace (PDQ and S_N type codes), in reservoir and weather modeling and indeed in all large scale linear algebra problems such as linear programming, the solution of systems of linear equations (structural and network analysis) and in eigenvalue problems (flutter and vibration analysis).

The lack of long precision arithmetic has been the greatest single shortcoming of the 2938 Array Processor preventing its general penetration into the scientific marketplace. Its inclusion in the Scientific Multiprocessor will be the most significant feature added to provide a product suited to the scientific marketplace. The inclusion of extended precision accumulation of sum reductions will answer a long standing call by scientific customers for such a feature. It will make a great impact in scientific computation because it provides a means, transparent to the user, to significantly reduce round-off error normally encountered in these computations.

Discussion of specific attachment not included

Attachment to the Host Processor

The Scientific Multi-Processor will attach directly to the memory of the host processor. It will not attach as a channel. Attachment in this manner will yield an order of magnitude improvement in the data transfer rates to the processor (on the Model 85 today, the 2938 at best can access data at the rate of double word each 3.7 microseconds. In the SMP the transfer rate will be 240 nanoseconds per quadword of data accessed). It should be noted that of the twenty or so numerical algorithms microprogrammed within the 2938 today, only one of these algorithms is compute bound. The others are data access bound. This improvement to the SMP will yield a more closely balanced system and a significantly higher performance level in the algorithms which were previously data access bound. This point is specifically addressed in greater detail ~~below~~. in the performance section below.

$$X_1(k+1) = b_0 u(k) + \sum_{i=1}^{2} b_i X_i(k)$$

$$X_2(k+1) = X1(k)$$

$$X_3(k+1) = b_5 u(k) + \sum_{i=3}^{4} b_i X_i(k)$$

$$X_4(k+1) = X_3(k)$$

In these defining equations each $\underline{u}$ is an element of the input data array; each $\underline{y}$ is an element of the output data array or may be the previous contents of $\underline{y}$; each $\underline{X}_i$ is an intermediate filter characteristic; and both a_0-a_4 and b_0-b_5 are constant coefficients.

ADDITIONAL INFORMATION

Additional information on the 2938 Array Processor exists and is available to you. Because much of it is preliminary in nature and subject to change, however, it has not been included in this newsletter. Your source for this additional information is your Regional Special Equipment Department. The names and telephone extensions of the specific individuals to contact are as follows:

Region	Your Contact	Extension
Eastern	Phil Jung	325
GEM	Al Robuck	7162
Midwestern	Ed Hayes	2032
Western	Jim Lieberknecht	1184

APPLICATIONS OF THE 2938 ARRAY PROCESSOR

Following below are some excerpts from a survey of computing requirements at a large institution. These excerpts typify a number of the application areas in which the 2938 may be profitably used.

1. Tape Digitizing and Signal Processing

   There is a wide-spread requirement within the computing community for a means of processing multi-channel analog tape. The basic requirements are for a high speed ADC, a fast fixed-head file and highly efficient array processing including general vector operations, convolution, recursive digital filtering and Fast Fourier Transforms. Potential users include:

   The Engineering Science Departments

Astronomy
Geophysics
Neurophysiology
Cardio-Vascular Surgery
Radio Physics
Control Systems

2. General Purpose Computing

A large number of computational problems resolve into a sequence of array operations which could be solved on the M44/2938 at a lower cost than on the M67. The basic operations of the 2938 are described in the literature. These operations can be chained to solve the following classes of problems:

a. Solution of Simultaneous Equations. An examination of the Gauss-Seidel algorithm reveals an iterative procedure with each step involving a matrix by vector operation. This can be programmed on the 2938 by chaining and checking for convergence between iterations. A rough estimate indicates that the 2938 could perform at 5-10 times the speed of the M65 on this kernel.

b. Partial Differential Equations. Many problems in thermodynamics, hydrodynamics and magnetohydrodynamics are described by a procedure which involves nested convolutions. Preliminary estimates indicate that the 2938 attached to a M65 will improve the performance on the weather code by a factor of 5.

Research is presently being carried out to determine the suitability of a M44/2938 to this type of problem and also the performance of a recursive filtering algorithm (which is essentially integration) on certain classes of partial differential equations.

c. Fourier Transforms. Large Fourier Transforms are becoming increasingly popular in the scientific computing community, especially since the introduction of the Cooley-Tukey algorithm. It can be expected that demand forthis type of service will increase as the unit cost of computation is reduced.

3. Spectral Analysis and Signal Processing

a. Spectral Analysis. Input for this application is high resolution analog tape (200 KC) which has been recorded at some radio telescope. The essence of the application is to sample the analog tape at a high data rate (200-500 KB), read the discrete time series into memory, perform a Fourier Transform on this record, process the Fourier Transform to produce a power spectrum record, and output the record. The output may go to a plotter, a CRT or to another computer for further processing.

The volume of the work contemplated is dependent primarily on the unit cost for calculation in the Fourier Transform since the radio telescopes are a virtually infinite data source. The comparative execution times for a 1000 point complex transform have been estimated as follows:

M44	.90 sec
M65	.33 sec

b. Signal Processing. The input for this application is again analog tape recorded at a radio telescope site. The essence of the processing is to sample the tape with an ADC, read the data to memory, process the records in memory with some algorithm (typically convolution) to produce filtered output records, and output the data. Output may go to digital tape, to analog tape or to another computer.

It is likely that there will be a need for a peripheral device to perform recursive digital filtering. If such an algorithm were available and applicable, data volumes could be increased (over convolution) by a factor of 10-100.

4. Control Systems

Computations in control theory are generally performed in conjunction with an analog computer. The analog computer is programmed to simulate some physical process such as an oil refinery. Most physical processes can be represented by a set of, in general, non-linear differential equations and the solution of these equations is particularly suited to an analog computer. It general the process, i.e., the analog computer has a multi-dimensional input, a multi-dimensional output, and a single valued profit function which is some combination of the inputs and outputs. The inputs are grouped into a "control vector" and the outputs are grouped into an "output vector". The function of the digital system is to compute a control vector, perturb the analog system with this control, measure the response at the output, then iterate in such a way that the profit variable converges to some maximum value.

The computation is a statistical procedure which involves matrix operations on the output vector. Functionally, the computer must sample the output, multiplex each element of the output vector into an ADC, read the data to memory, perform the matrix operations, then write the updated control vector to an array of DAC's. A key parameter of this computation is the "order" of the system being optimized. The "order" is the number of differential equations which represent the system and is approximately equal to the "order" of the matrices involved in the digital computation. The speed with which an analog computer solves a set of differential equations is independent of theorder of the system but the time required to perform the digital computation is proportional to N**2 if N is the order of the system.

The entire system becomes rapidly bound by the floating-point multiply time of the digital computer as N increases. At present 2nd and 3rd order systems are being solved, but high-order systems would be studied if the matrix multiplication bottleneck could be cleared.

5. Pattern Recognition

Experiments in voice recognition have been carried out for a number of years. A typical computational procedure involves an array of band-pass filters which produce an analog spectral profile of a word syllable in real time. The output of each filter is multiplexed to a high-speed (100 KB) ADC and read to memory. The pattern-recognition computation involves vector and matrix operations which eventually resolve the series of spectral profiles into words. Output may be to any number of standard peripherals or to specially fabricated devices.

In general, these devices must be able to read and/or write randomly to memory at memory speed thus creating a need for a "port to memory".

6. Holography

Holography is a relatively new technology involving laser beams, etc. and is being considered as a future application.

Input would be a high speed scan of an image in some sort of television camera. Since a high resolution scan in 2 dimensions produces voluminous data (1 Megabyte) a long record may be formed on a fixed-head file. Processing consists largely of a 2-dimensional Fourier Transform on the input record. Output would be to an optical plotter or CRT.

This type of holography is suitable for image restoration in aerospace telemetry. Again, the economic feasibility is contingent on a high speed Fourier Transform.

Another application of the 2938 Array Processor, studied elsewhere, is in meteorology. In particular, the suitability of a Model 65 with a 2938 for programs involved in numerical weather prediction has been studied. In a sample problem that has been programmed, it has been determined that all numerical operations involved in the iterative portion can be performed on the 2938. It was estimated that the actual forecast model would run four to five times faster on the 2938-equipped Model 65 than on the largest general-purpose computer now installed in the scientific marketplace, provided that the problem could be structured so as not to be I/O bound.