

A HIGH-SPEED DIGITAL FOURIER ANALYZER

A. M. Chwastyk

Precision high-speed spectral analysis is essential for advanced multiple channel sonar and radar processors. The high resolution and accuracy requirements can be satisfied with a general purpose computer. However, to keep the cost of processing large banks of data low, the physical size small, and the data throughput rate high, special purpose processors are necessary. The combination of new developments in digital computing elements and fast Fourier transform algorithms, has made possible the construction of special purpose processors with outstanding performance characteristics.

Introduction

THE FAST FOURIER TRANSFORM ALGORITHM (FFT) has revolutionized the digital processing of waveforms. The algorithm is an efficient method for computing the discrete Fourier transform. The high-speed signal analyzer built at the Laboratory is based on a shift register implementation of a fast Fourier algorithm coupled with a single flow-through arithmetic unit. Speed is obtained by using an organization that performs the rapid data reordering required for executing the algorithm. 4096 complex input data samples can be transformed in less than 12 ms, making it one of the fastest systems in operation.

In addition to the Fourier transform hardware, the analyzer also includes analog to digital converters, a core memory for multiplexing data, and a postprocessing unit. Postprocessing is used to increase the quality of spectra estimation in statistically random short term spectra. The unit can be used either as an integrator or a first order recursive filter. It is implemented in floating point word format to maintain dynamic range. Hardware is kept to a minimum by establishing minimum and maximum ratios of the two signals required in the

recursive filter computation. Time required to generate the power spectra from the Fourier coefficients and update 4096 post detection filter bins located in the core memory is 6 ms.

The analyzer is significant for its high processing speed at moderate clock rates and the advanced data handling concepts used in its design. It is being used for seismic, vibrational, radar doppler, and sonar analysis.

The Discrete Fourier Transform

Since the fast Fourier transform is a method for computing the discrete Fourier transform (DFT), it is appropriate to briefly review some of the properties of the DFT. The Fourier transform is used to determine the frequency content of temporal functions. The Fourier transform, $X(f)$, of a continuous function, $x(t)$, can be described by the following equation:

$$X(f) = \int_{-\infty}^{+\infty} x(t) \exp(-j2\pi ft) dt \quad (1)$$

In order to implement this equation digitally, the continuous input, $x(t)$, must be sampled in time.

Using Shannon's sampling theorem, the rate, f_s , at which the sampling occurs must be at least twice the highest frequency component of the input function. This yields an infinite set of samples, x_n , whose transform as calculated by Eq. (1) no longer contains magnitude and phase information at frequencies $f_s/2$ or above. The sampling process can be represented mathematically as the multiplication of the continuous input by the sampling function. The transform of the sampled input can be found by convoluting the transform of the original input. Since the sampling function is a train of impulses at a sampling rate f_s , the convolution results in a series of repeats of the signal spectrum at multiples of f_s . For input functions that have frequency components higher than $f_s/2$, the input signal is band limited so that signals beyond $f_s/2$, i.e., the Nyquist rate, are highly attenuated. For some coherent applications, as in the case of some pulse doppler radars, ambiguities are allowed to exist with absolute frequency (doppler) resolved by means of a change in the f_s or by range rate measurements.

A second consideration in the practical implementation of Eq. (1) requires limiting the observation time. The effect of this time window is to limit spectral line resolution. For a given time interval, T , the Fourier transform of a unit amplitude signal is

$$X(f) = \int_{-T/2}^{+T/2} \exp(-j2\pi ft) dt = \frac{T \sin \pi f T}{\pi f T} \equiv T \operatorname{sinc}(\pi f T). \quad (2)$$

This function has a 4 db bandwidth of approximately $1/T$ with a characteristic sinc x response with nulls every $1/T$ Hz. A periodic sampling of the input waveform at N points in the observation time extends it into a sampled periodic function with the transform yielding a discrete set of N coefficients. For this discrete finite transform, Eq. (1) can be written in the following series form

$$X(m\Delta f) = \sum_{n=0}^{N-1} x_n \exp\left(\frac{-j2\pi n m\Delta f}{N}\right), \quad (3)$$

where m is the Fourier coefficient or frequency bin index and $\Delta f = 1/T = f_s/N$. The transform is actually a sampled Fourier series where evaluations at $m\Delta f$ correspond to a set of filters, each with the sinc x characteristic of Eq. (2). As in the

series, the coefficients indicate the amplitude and phase at each frequency in x_n . The power at a given frequency is obtained by summing the squares of the real and imaginary parts of the complex frequency coefficients.

For input signals periodic in the time window, i.e., having frequency components that are exact multiples of $1/T$, all the energy will appear in its proper frequency bin with zero crosstalk into other bins since the center frequency of any bin corresponds to zeroes of the sinc x characteristic of all the other bins. For nonperiodic signals, i.e., the normal case, leakage of energy into adjacent frequency bins will occur. The worst case leakage occurs for signals whose frequency lies midway between two adjacent frequency bins. Most of its energy is distributed between the two closest bins, 3.9 db down from maximum. Filters on either side of the center two will be 13.5 db down from maximum with others falling off at a 6 db per octave rate. Because of the discrete nature of the transform, aliasing will tend to decrease the rate of attenuation of leakage. A band limiting function called Hanning weighting is often used to decrease this leakage at the expense of a decrease in resolution in the main lobe. This can be achieved either in the frequency domain by smoothing the Fourier coefficients or in the time domain by multiplying the input data x_n by $(1/2 - 1/2 \cos 2\pi n/N)$. Main lobe broadening of 1.62 occurs with a maximum sidelobe at -31.4 db and a fall off of approximately 18 db per octave.

The properties of the DFT are in agreement with corresponding properties of the Fourier integral transform. The inverse transform can be executed using the same basic operations used in the forward transform. By using the DFT for frequency analysis one obtains the effect of a bank of identical filters, all based on the same input data block. Also, since x_n can be complex, it finds unique application in coherent radar and sonar systems. With complex inputs, the DFT can be used for obtaining N frequency filters distributed from $-f_s/2$ to $+f_s/2$ rather than the $N/2$ filters covering the range 0 to $+f_s/2$ for normal inputs.

The Fast Fourier Transform Algorithm

The requirement for fine resolution ($1/T$) over wide bandwidths ($f_s/2$) limits the applicability of direct implementation of Eq. (3). Since N complex computations (four multiples plus additions)

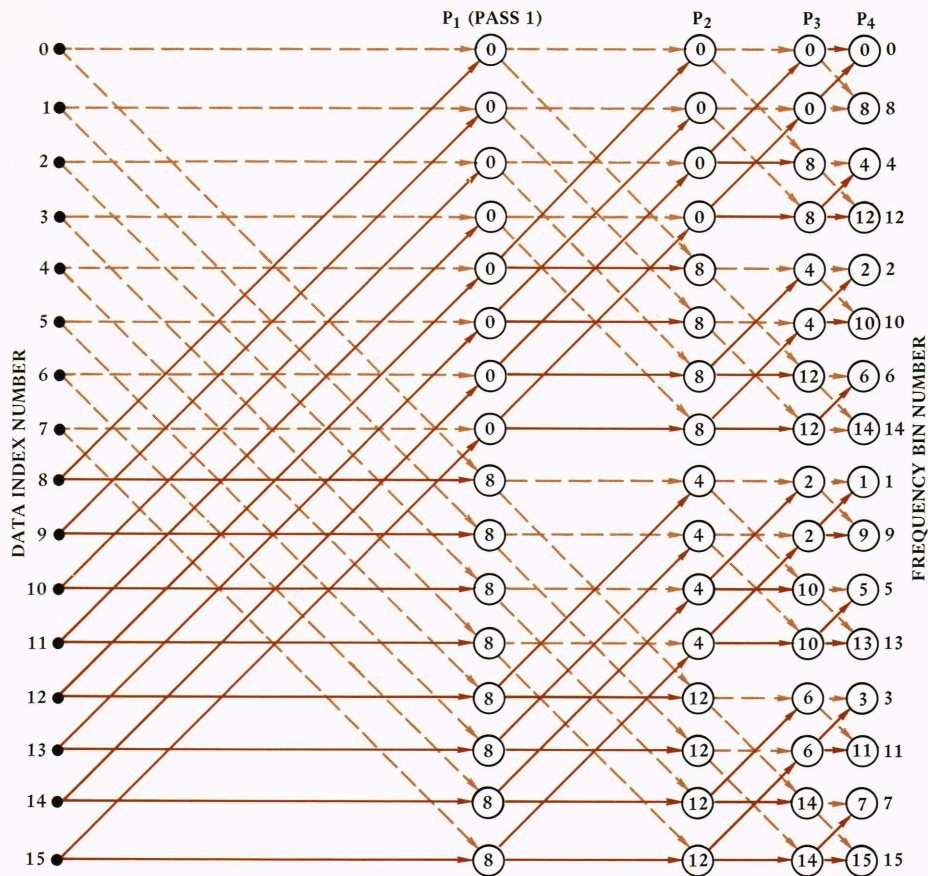


Fig. 1—Flow graph for naturally ordered time samples.

are required for each filter, a total of N^2 operations are required for the coherent systems. For real-time systems, processing rates rapidly become excessive for large values of N . The Cooley-Tukey algorithm,¹ limits the basic number of complex operations to $1/2 N \log_2 N$. The algorithm eliminates the redundant operations which are performed in the direct implementation of the DFT. The process essentially utilizes a series of two point transform computations where foldover ambiguities are systematically reduced at each stage until the full sampling frequency is reached. This requires using the sampled data points two at a time to generate N intermediate answers. These in turn are combined into mutually exclusive subsets and transformed so that occurrence times are accounted for in another pass through the data. This sequential combining to progressively larger weighted sums of the data results in the forming of the discrete Fourier transform.

There are several good derivations of the transform in the literature.^{2,3} For our purposes, considerable insight can be obtained from the study of a graph showing the computational procedure. Also, the mechanics of its general usage will aid in understanding the FFT implemented. A flow graph for naturally ordered time samples is shown in Fig. 1. For those not familiar with flow graph notation, it consists of nodes and directed branches. Each node represents a variable which is the weighted sum of the variables of all the branch nodes on the left that terminate on that node. In the flow graph, the left hand column of small circles corresponds to the input data points with index numbers 0 through 15 for the $N = 16$ array used for the example. In a random access system, this would correspond to relative memory address. The nodes of the remaining columns represent an operation of $A + B \exp(-j2\pi Z/N)$

¹ J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," *Math. of Comput.*, **19**, Apr. 1965, 297-301.

² W. T. Cochran, et al., "What is the Fast Fourier Transform?," *Proc. of the IEEE*, **55**, Oct. 1967, 1664-1677.

³ G. D. Bergland, "A Guided Tour of the Fast Fourier Transform," *IEEE Spectrum*, **6**, July 1969, 41-52.

where the A input can be determined by following the dotted line to the left of the node and B by following the solid line. The Z factor is the rotational value listed in the nodes of the flow graph. A pass is defined as a completion of all operations in a column. The intermediate answers obtained are indexed 0 through 15 to correspond to the input data set.

To minimize computations and memory access cycles when implementing the flow graph, it is desirable to rotate the B inputs only once per two point transform. This is done by

$$\begin{aligned} A' &= A + B \exp(-j2\pi Z/N) \\ B' &= A - B \exp(-j2\pi Z/N) \end{aligned} \quad (4)$$

where Z is now modulo $N/2$ and the additional 180° rotation is supplied by the negative sign.

For example, in pass 1, data points at locations 0 and 8 are read from memory and combined as in the equation above. The computed values (A' , B') are then returned to memory locations 0 and 8 respectively. This is repeated for point pairs (1,9), and (2,10); . . . until (7,15) are combined and returned to their respective memory locations to complete the pass. The next pass utilizes data from memory positions (0,4), (1,5), (2,6), (3,7), (8,12), (9,13), (10,14), (11,15) always returning the computed pairs to corresponding memory positions as they are combined. This process is repeated in similar fashion until all passes are complete.

A study of the flow graph reveals patterns that can be utilized in programming the algorithm. For instance, each pass requires only the data generated from the preceding pass. That is, data can be processed in place, two words at a time. Second, if data are processed in order of arrival, data pairs for executing Eq. (4) come from locations $N/2$ apart on the first pass, $N/4$ on the second or $N/2^p$ for the general case, where p is the pass number. This requires a reordering of the memory addresses on each pass through the data. Third, the rotational values required have the same periodicity as the data displacement of $N/2^p$. The values required always change by a number corresponding to a stepped binary counter, with its output bits reversed end for end. Fourth, the final set of answers are not generated in order, but its frequency bin can be obtained from the bit reversed binary memory address. Another factor illustrated by the graph is that independent but interlaced data sets can be processed. For in-

stance, two eight-point transforms based on odd and even data points in each set, can be obtained by stopping the process after the third pass.

The FFT Unit

Keys to the high-speed implementation of the algorithm as used in the digital Fourier analyzer (DFA) are: (1) a shift register organization that eliminates need for memory addressing in the re-ordering required, (2) a pipeline arithmetic unit capable of obtaining vector products at the shift register clocking rate, and (3) a sine-cosine generator also capable of operating at the shift register rates. These items will be discussed in sequence.

The 16 point transform will again be used as an example to establish flow patterns and control of the reiterative FFT unit organized as shown in Fig. 2. Data are initially stored in shift registers A and B in order such that shift register A contains the first 8 (or $N/2$) words and shift register B contains the last 8 words. The first pass is obtained by shifting both SRA (shift register A) and SRB (shift register B). With reference to the flow diagram, the SRA and SRB data streams correspond to the dashed and solid lines on the diagram

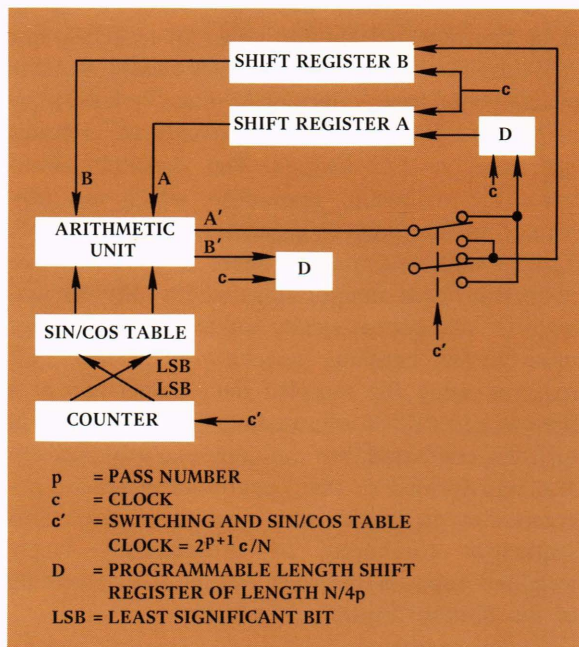


Fig. 2—Shift register FFT organization based on programmable length shift registers.

respectively. As output pairs of words are obtained as per Eq. (4), the data stream at the B' output is delayed by 4, or $N/2^{p+1}$ for the general case. This delay makes it possible to switch between the A' and delayed B' outputs so that data words leaving the switch are staggered in time but in proper order for the next pass. The switch position is changed every 4 (or $N/2^{p+1}$) clock pulses. The data entering SRA are delayed by 4 clock pulses to eliminate the staggering without stopping the SRA clock. This allows the use of dynamic type metal oxide semiconductor (MOS) shift registers and simplifies system timing. Successive passes through the intermediate data points follow the same format, with delays reduced on each pass until the single delay, last pass has been completed. If the output data can be used at the processing rate, new data may enter the system during the last pass. Thus, the minimum time between data sets becomes $1/2 N t_c (1 + \log_2 N)$ where t_c is the clock period.

The addressing of the sine-cosine table is accomplished by stepping a memory address counter every $N/2^{p+1}$ during a pass. The counter bit weights are used in reverse order, i.e., the least significant bit of the counter is used to address the most significant bit of the memory, etc.

The above two paragraphs establish the control mechanics for any value of N where $N = 2^\gamma$ and γ is an integer. Since all values of one pass are available prior to the start of the following pass, array scaling or "normalization" of the entire data block is used to reduce the number of bits required in the shift registers and the arithmetic unit.

For a small number of points per data set, digitally controlled variable delay units offer the most economical implementation. For systems requiring large values of N , the B register, with additional switching, can be substituted for the variable delay as shown in Fig. 3. Again, switching occurs every $N/2^{p+1}$, and shifting is continuous. The processing time reduces to $1/2 N t_c \log_2 N$.

The above two formats are combined in the FFT unit used in the analyzer to reduce hardware. For shifting operations requiring delays of 1 to 64 bits, a variable shift register is used to save switching circuitry. For larger delays, switching circuits are used. This results in a savings of shift register stages and increases the processing

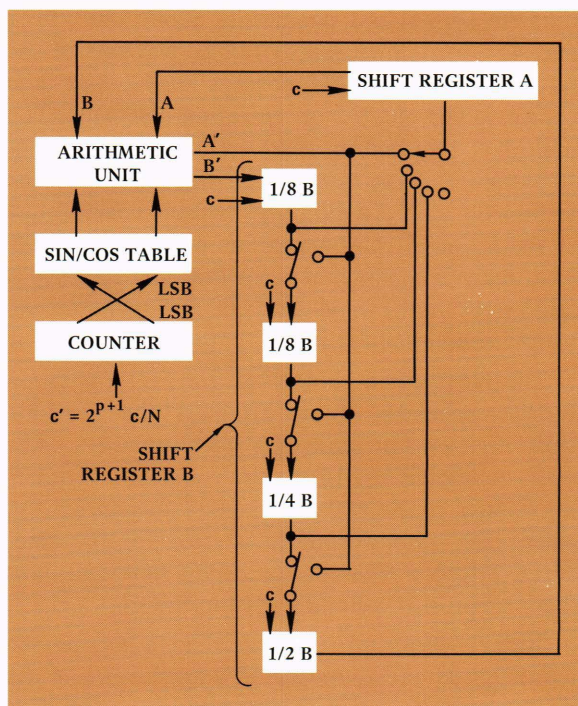


Fig. 3—Shift register FFT organization based on switching of shift register segments.

rate at the expense of additional data switching. The net processing time for $N > 64$ equals $(128 + N/2 \log_2 N) t_c$. Maximum clock rate for the 4160 word by 26 bits per word shift registers used in the DFA is 2.5 MHz.

The arithmetic unit used in the FFT unit operates in rectangular coordinates. It is of pipeline configuration, i.e., it has internal delays for a particular data increment to "flow" through the unit, but can process the data at the clock rate. In order to use standard transistor-transistor logic, some parallelism is also required. Four 12×12 bit multipliers with signs carried externally, are used to perform the vector rotations required. They are built from an array of four bit full adders and quad two-input nor gates. Propagation delays in the multipliers are reduced by subdividing the partial products so that run out delays occur simultaneously in each subarray.

The trigonometric function generator supplies 13 bit sine and cosine values to the arithmetic unit at the 2.5 MHz rate. The values required in the transform are represented by terms which are integer powers of $\exp(j2\pi/N)$. The generator, therefore, is addressed in the binary angle meas-

urement system and symmetries about 180° , 90° , and 45° axes are used to reduce the size of the table. The basic table consists of 32 sine and cosine values plus 32 interpolation constants. The multiplier required in using the interpolation constants are of array type similar to that used in the arithmetic unit. To produce the sine-cosine pairs at the required rate, a pipeline approach is used. The delays encountered from the time the unit is addressed to the time answers are available are simply accounted for in the timing control.

The hard wired control of the FFT operation consists primarily of two counters plus a 12 bit shift register for indicating pass numbers. The main counter is used for shift count and reordering control of the variable length shift register and gates. The second counter is used for the sine-cosine address. The control circuit represents less than 1.5% of the circuits used in the FFT unit.

System errors are introduced in digitizing the input signals and rounding off word lengths in the processing. The errors obtained are also a function of the input signal statistics, input signal level, and amount of coherent processing.

The analog-to-digital noise can be broken up into two parts. For high level signals, energy is spread over many frequency bins and quantization noise is usually negligible. However, as the input signal is lowered, the quantization noise relative to signal energy will increase, degrading system performance. The other effect is due to digital saturation. Since the digitizing comes after the band limiter, this can introduce harmonics well above the Nyquist frequency. For noise-like inputs, this has the effect of adding a noise floor to the output, thereby limiting dynamic range.

The discrete quantization of the signal affects a less than maximum amplitude sinusoid more with no noise background than with noise. The maximum undistorted sinewave output power to the output power for minimum input is $6.02b - 3.01$ db where b is the number of bits in the input word. With additive white noise, the difference between the quantized wave and input wave is a random variable between $-1/2$ and $+1/2$. This relates to a power of $1/12$ in the quantization noise, or $6.02 b + 1.76$ db. Within the 69 db two-signal accuracy limit imposed by the word size used in the FFT, the maximum dynamic range of the transform output is $(6.02 b + 3.01 \log_2 N + 1.76)$ db.

Postprocessing Unit

When analyzing signals containing wideband noise, the short term spectra are themselves statistically random and, when detected as in the FFT output power spectra, exhibit a chi-squared distribution. These short term statistical fluctuations of the spectra tend to mask the long term, relatively stationary lines desired. Averaging (or summing) of a number of ensembles improves the quality of the estimate in direct proportion to the square root of the number of statistically independent spectra averaged. Since the processor is programmed to take in essentially contiguous time data sets, the noise spectra samples for the data blocks processed are independent.

In off-line processing, a straight summation of a number of time-adjacent data sets is all that is required for averaging. However, for a real-time processor, it is desirable to display a running or sliding window. This can take the form of

$$P(iT) = \sum_{h=0}^{H-1} p(iT - hT), \quad (5)$$

where T = time between ensemble sets, $p(i)$ represents a point from the frequency ensemble at time i , and H corresponds to the number of data sets to be summed. To implement this directly would require a storage of the latest H values for each filter bin. To minimize the storage memory to one word per frequency bin, an inverse time exponential weighting is used. This can be expressed in nonrecursive form as

$$P(iT) = \sum_{h=0}^{\infty} K^h p(iT - hT), \quad (6)$$

where K is always less than one. Expressed as a first order linear difference equation, the equation, as implemented in the DFA becomes

$$P(i) = p(i) + (1 - 1/2^k)P(i - 1), \quad (7)$$

where K , equal to $(1 - 1/2^k)$, is limited to values formed by $k = 0, 1, 2 \dots 7$.

Except for the discrete nature of the above difference equation, it is similar to a resistor-condenser low pass filter with a gain of 2^k . For values of k equal 3 or more, the effective time constant of the circuit is equal to $2^k T$. T in this case can be controlled independently of the data gathering time. However, in normal usage T is the time required to gather one data set ($T = N/f_s$).

The three card postprocessing unit represents a

blend of speed, performance, and circuit minimization. Its circuit can be used for the power summing of ensemble sets as well as the filtering of Eq. (7). A large dynamic range is preserved at the expense of some accuracy.

The basic steps performed in the unit are: (1) convert the power spectra words into floating point format; (2) interrogate core memory for previous filter value; (3) update old filter value by multiplying by a K factor, and adding the product to the new value obtained in step (1); and (4) store updated value into the original core memory position of step (2). These steps are repeated for all frequency bins (8192 max.) and are confined to one clock period per frequency bin update to be compatible with the memory read-modify-write cycle.

The inputs to the postprocessing unit are two 24-bit words (the squares of the real and imaginary parts of the Fourier coefficients) plus 5 bits of block floating point. The block floating point number represents an exponent common to the entire array of coefficients. This represents a dynamic range equivalent to 27 bits (~ 81 db) for a 4096 point transform with an 8-bit input word as implemented in the DFA. An additional 7 bits of range is required for implementation of the $k = 7$ mode of the recursive filter for a maximum possible dynamic range of approximately 100 db.

For useful outputs, the above dynamic range should be preserved. From an engineering viewpoint, however, accuracy of representation seldom requires magnitudes to be expressed to greater accuracy than ± 0.15 db. Limiting the accuracy re-

duces the floating point conversion hardware. Thus, the raw power spectra can be represented by a 6-bit exponent (characteristic) plus 5 accuracy bits (mantissa). However, a minimum of 7 more bits are required in the mantissa of the filtered output, $P(i)$, since there is an integration gain of 2^7 for $k = 7$.

The Overall System

A block diagram of the system is shown in Fig. 4. Band-limited analog data, either real or complex, can enter the system via two, 8-bit analog-to-digital converters, each with a maximum conversion rate of 625 kHz. The sampled data are stored in the core memory for multiplexing of data sets. Word blocks of 4096 samples are sent to the fast Fourier transform unit. If the Hanning weighting option is selected, samples are weighted as they are loaded into the shift register store. The Fourier coefficients are formed in the shift register after $\log_2 N$ circulations or passes through the arithmetic unit. One additional pass is required to convert the coefficients into power spectra. As the power spectra points are generated, they are used to update postprocessing filter bins located in the core memory. Spectral data, either with or without postprocessing, is displayed on an oscilloscope. The processor is interfaced with a computer for additional processing and displays.

The FFT shift register memory has storage capabilities for 4096 complex data points which can be divided into a single set of 4096 points or M sets from M different input channels. The number of points per set (N) must satisfy the require-

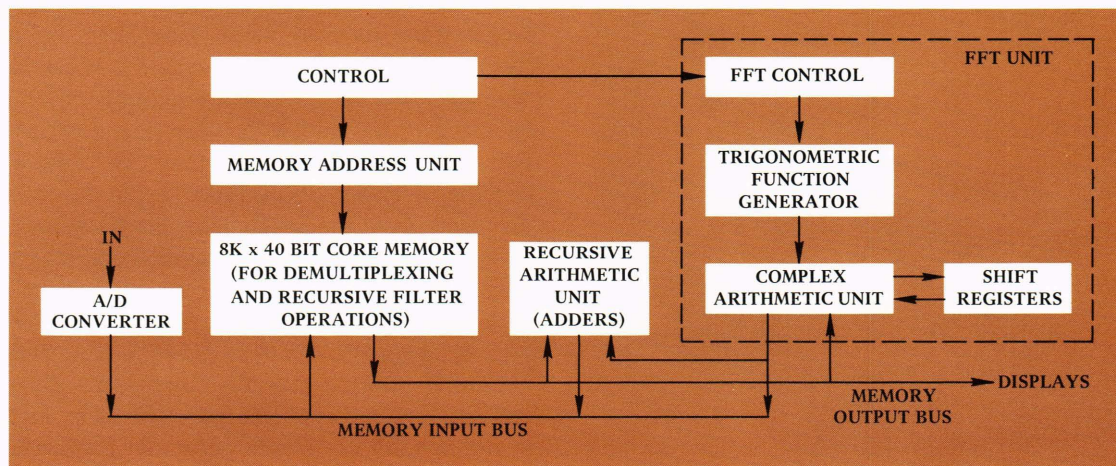


Fig. 4—Digital Fourier analyzer block diagram.

ment that $NM = 4096$. The input memory extends the number of 4096 blocks the system is capable of handling. For reasons of economy, both the input multiplexing memory and the postprocessing unit share the single core memory. This imposes a restriction on overall speed and display update rate. The core memory, as now programmed, can store up to 8192 complex input samples and 8192 postprocessing filter positions. Normal procedure is to load raw data into core, process in the FFT unit, and update the postprocessing filters as the processed data are returned to core.

The processor is controlled by a 40-bit control register that can be loaded from the front panel or from a computer. Programs are hard wired with built-in expansion capabilities. All timing is based on an internal 20 MHz crystal clock. To aid checkout, internal registers can be examined visually via display light at any cycle in the program. Checkout control is made via the front panel.

The system, pictured in Fig. 5, is packaged in two drawers plus the core memory. The construction used in the DFA was designed to minimize time between paper design completion and final hardware by the elimination of etched wiring artwork. A high density universal card was designed initially for the DFA to mount the dual-in-line packaged circuits. Interconnections are made by point-to-point wiring. Connections were made by a welding procedure perfected by APL for satellite



Fig. 5—Digital Fourier analyzer.

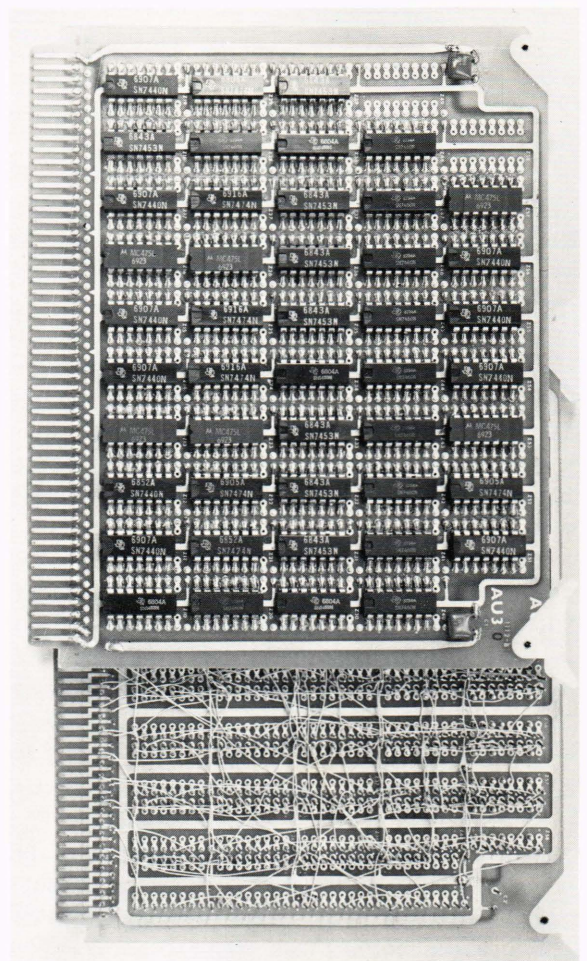


Fig. 6—DFA card.

construction. Welded wiring boards were used instead of wire wrap boards to decrease card wiring costs and to decrease board to board spacing. Costs were reduced since welds can be made through the insulation, removing the necessity for wire stripping. Also, only one weld operation need be made in the middle of wire runs compared to two operations for wire wrap run. Interboard spacing is also reduced since a pin for welding protrudes from the wiring side approximately 1/4 inch compared to 0.5 inch (min.) for wire wrap pins. A total of 42 cards per drawer can be used with 3/4 inch card spacing. The card, shown in Fig. 6, measures 4-3/8 x 5-3/4 inches and has mounting positions for 48 circuits. Its design is consistent with recent trends that show an increase in the number of circuits per card. This trend results from the need to minimize card connectors, backplane wiring, and lead lengths.

Discussion

The processor has completed its first 1½ years of usage without component failure. This indicates the high reliability of the relatively new MOS circuits. The characteristics of digital filtering, i.e., large output dynamic range, stability, high Q filters regardless of frequencies, uniformity of filter bin spacing and response, are all inherent in the processor. In addition, the high-speed operation of the DFA makes it especially useful for wideband digital processing of coherent radar signals and multiple channel sonar processing.

Its most used mode is to process analog inputs and send either raw or filtered data to a computer. The data then can undergo some additional modification prior to either storage, power spectra plotting, or direct lofargram generation. This arrangement places the digitizing and high speed

processing burden on the DFA with the control and lower order processing on the computer.

The analyzer reflects use of the medium scale integration circuits available at the time of construction. Using the latest technology, reclocked array arithmetic units can now operate at 18 MHz. MOS shift registers are now available at 20 MHz clocking rates. Therefore, a new implementation of the basic organization presented could result in a speed increase of seven and an improvement in the cost to computation rate ratio. Also, more parallelism could be incorporated into the design to increase speed even further. However, any new analyzer should be considered for the total system in which it is to be used. Preprocessing, postprocessing, and data handling factors already can readily overshadow the basic transform unit in cost and complexity.

PUBLICATIONS

Compilation of principal recently published books and technical articles written by APL staff members.

- K. Green, B. Hastings (The Johns Hopkins School of Medicine), and M. H. Friedman (APL), "Sodium Ion Binding in Isolated Corneal Stroma," *Amer. J. Physiology* **220**, No. 2, Feb. 1971, 520-525.
- A. G. Schulz, L. G. Knowles, L. C. Kohlenstein, and L. J. Maroglio, "A Digital Simulation of the Rectilinear Scanning Process," *Proc. Symp. on Sharing Computer Programs and Tech. in Radiation Therapy and Nuclear Medicine*, Oak Ridge, Tenn., Apr. 2-3, 1971.
- F. T. Fraunfelder (The Johns Hopkins School of Medicine), L. J. Viernstein (APL), "Intraocular Pressure Variation during Xenon and Ruby Laser Photocoagulation," *Am. J. Ophthalmology* **71**, No. 6, June 1971, 1261-1266.
- L. W. Ehrlich, "Solving the Biharmonic Equation as Coupled Finite Difference Equations," *SIAM, J. Numer. Anal.* **8**, No. 2, June 1971, 278-287.
- R. M. Fristrom, "The Chemistry of Polymer Flames—Propagation and Suppression," *Proc. Polymer Conference Series, Flammability Characteristics of Polymeric Materials*, Section 8; June 1971.
- R. C. Orth and H. B. Land, "A Production Type GC Analysis System for Light Gases," *J. Chromatographic Sci.* **9**, June 1971, 359-363.
- F. J. Adrian, "Contribution of $S_0 \leftrightarrow T_{\pm 1}$ Intersystem Crossing in Radical Pairs to Chemically Induced Nuclear and Electron Spin Polarizations," *Chem. Phys. Letters* **10**, No. 1, July 1, 1971, 70-74.
- T. P. Armstrong (Univ. of Kansas), S. M. Krimigis (APL), "Statistical Study of Solar Protons, Alpha Particles, and $Z \geq 3$ Nuclei in 1967-1968," *J. Geophys. Res.* **76**, No. 19, July 1, 1971, 4230-4244.
- J. R. Apel and T. O. Poehler (APL) and C. R. Westgate and R. I. Joseph (The Johns Hopkins University), "Study of the Shape of Cyclotron-Resonance Lines in Indium Antimonide Using a Far-Infrared Laser," *Phys. Rev. B* **4**, No. 2, July 15, 1971, 436-451.
- D. M. Silver, "Metric Evaluation in the Space of Exponential and Gaussian Functions," *Chem. Phys. Letters* **10**, No. 2, July 15, 1971, 227-229.
- M. H. Friedman (APL) and K. Green (The Johns Hopkins School of Medicine), "Ion Binding and Donnan Equilibria in Rabbit Corneal Stroma," *Amer. J. Physiology* **221**, No. 1, July 1971, 356-362.
- K. Green (The Johns Hopkins School of Medicine) and M. H. Friedman (APL), "Potassium and Calcium Binding in Corneal Stroma and the Effect on Sodium Binding," *Amer. J. Physiology* **221**, No. 1, July 1971, 363-367.
- B. E. Tossman, "Variable Parameter Nutation Damper for SAS-A," *J. Spacecraft and Rockets* **8**, No. 7, July 1971, 743-746.
- D. M. Silver, "Electron Pair Correlation: Products of $N(N-1)/2$ Geminals for N Electrons," *J. Chem. Phys.* **55**, No. 3, Aug. 1, 1971, 1461-1467.
- T. A. Potemra (APL) and L. J. Lanzerotti (Bell Tel. Labs.), "Equatorial and Precipitating Solar Protons in the Magnetosphere. 2. Riometer Observations," *J.*