THE MULTICHANNEL SPECTRUM ANALYZER

A. M. Peterson, K. S. Chen, and I. R. Linscott
Space, Telecommunications and Radioscience Laboratory
Stanford University
Stanford, CA 94305

ABSTRACT. The MCSA is a special-purpose digital signal processor. Its main function is to filter a wide-band signal into many narrower bands, so that each of the output bands has a bandwidth that is a better match to the signal being searched for.

The basic MCSA provides simultaneous output bandwidths of approximately 1 HZ, 32 Hz, 1024 Hz, and 74 kHz over a spectrum that is about 8 MHz wide. The input to the MCSA consists of a complex signal sampled at 10 MHz, and the outputs consist of either complex samples or power (square-law-detected) samples. In addition, the MCSA provides an accumulator for taking the integral of the power of the output bands for periods up to 1000 sec.

The MCSA hardware is constructed using wire-wrap technology. The implementation of the hardware is done with the aid of a computer program developed specifically for the design of the MCSA. Care has been taken in the MCSA design to ensure that engineering tradeoffs do not adversely affect the performance of the system.

1. GENERAL DESCRIPTION

Instead of using a single large Fast Fourier Transform (FFT), the MCSA derives its narrow bands by cascading two stages of digital bandpass filters with moderate-sized Discrete Fourier Transforms (DFT). FFT operations do not yield convenient signals for deriving the intermediate bandwidths that the MCSA delivers. Furthermore, it is possible to provide better RFI rejection with the bandpass filter technique. An FFT has a worst-case sidelobe (adjacent bin) response that is only 13 dB below the response of the main lobe. A bandpass filter can be designed to give more than 70 dB of adjacent channel rejection.

The first bandpass filter splits the input signal into 112 bands, each approximately 74 kHz wide. Each of these 74-kHz-wide signals is then filtered by a second bandpass filter which further subdivides the signal into 72 bands. Each of the resultant bands is about 1024 Hz wide.

M. D. Papagiannis (ed.), The Search for Extraterrestrial Life: Recent Developments, 373–383.

The 1024-Hz signals are then fed either to a 36-point DFT or to a 1152-point DFT to form the final 32-Hz or 1-Hz outputs, respectively. Each of these bandwidths (1 Hz, 32 Hz, 1024 Hz, and 74 kHz) is available as an output of the MCSA. The magnitude squared value of each output sample is computed and is available as the square-law-detected power output. Except for the 74-kHz bandwidth. The complex signals from the other bands are also available as outputs.

The memory required to do real-time transforms is equal to twice the size of the transform -- one block of memory being required to buffer a block of data while the second block is having the transform operations done to it. By using relatively small transform sizes in the final DFTs, the need for a main memory that is twice the size of the MCSA bandwidth can be avoided. Thus, a substantial saving in memory cost is made by applying the bandpass filter/DFT technique compared to a single large FFT, offsetting the extra cost needed to implement the bandpass filters.

Internal to the MCSA, the sampling frequency of a given signal is kept a fraction larger than its analysis bandwidth to prevent additional noise from aliasing into the band of interest. If this were not done, it would be impossible to avoid a substantial degradation of the S/N near the extremes of the band, however sharp the cutoff of the filter is. A sharper filter reduces the size of the region that gets severe aliasing, but could not prevent a 3-dB loss of S/N at the edge of the band due to aliasing. The sample frequencies used were chosen such that for each filter, the additional aliased noise is kept below 0.1 dB.

The use of oversampling causes the DFT transform sizes to not be powers of two. This is not a disadvantage, however. Through the use of newly developed DFT algorithms (called the Generalized Winograd DFT Algorithms), the computational efficiencies achieved are comparable with those of power-of-two FFT algorithms. Hardware complexity has increased (more hardware is needed in some cases, faster hardware is needed in others) because of over-sampling. This, however, is unavoidable unless a substantial loss in S/N can be tolerated.

2. BANDPASS FILTER 1

This first bandpass filter (see section, Digital Bandpass Filter) operates on a complex input that is sampled at 10 MHz. Each of the real-imaginary pair of samples is made up of a pair of 8-bit binary numbers taken by an analog-to-digital (A/D) convertor subsystem that precedes the MCSA.

Before being given to the bandpass filter, each 8-bit sample is further quantized into a 4-bit representation to reduce the arithmetical complexity of this very-high-speed stage. It was found that, with a signal that has Gaussian statistics, the loss in S/N due to the quantization of the signal to 4 bits is no more than 0.05 dB. Coarse quantization of a signal has been a necessity in many systems that are required to operate at high speeds. However, careful analyses have shown that very low S/N losses can indeed be achieved without having

to quantize the signal into very fine levels, providing the quantization parameters are held to some given values (see section, Digital Bandpass Filter).

The first bandpass filter is implemented at two identical banks of filters. Alternate 74-kHz passbands appear at the outputs of either of the two filters -- the even 74-kHz bands appearing at the outputs of the first filter bank and the odd 74-kHz bands appearing at the outputs of the second filter bank. The basic filter bank separates its input into 72 bands, of which 16 are discarded, retaining 56 of the bands as outputs. The 16 discarded bands represent the part of the input signal that has substantial amounts of noise aliased into it by the sampling process.

The input signal, following the 4-bit quantization, is applied directly to the input of the first filter bank. Before being applied to the second filter bank, the input signal is shifted in frequency by 74 kHz. The frequency shifting is achieved by the multiplication of the input signal by a complex sinusoid at the appropriate frequency, done digitally. The output of the frequency shifter, quantized to 4 bits, is then applied to the input of the second filter bank.

Each of the two filter banks consists of a 288th order Finite Impulse Response (FIR) filter and a 72-point DFT. The 72-point DFT is implemented with an 8-point FFT followed by a 9-point DFT.

The 8-point FFT is realized with hard-wired pipelined logic circuits. A table-lookup Read-Only Memory (ROM) is used to implement the only multiplication, the scaling of a value by $\sqrt{2}$, required to perform an 8-point FFT. The number representation at this stage (and successive stages) is kept at 16 bits.

The 9-point DFT is implemented with eight special-purpose programmable processors (see section on DFT Processor) working in parallel. A single microprogrammed controller controls all eight DFT Processors (also the 8 DFT Processor in the other bank of the filter).

The multiplication by the filter weights of the FIR filter is done with table lookups using ROMs. The FIR filter is done with table lookups using ROMs. The FIR filter taps are partitioned into four sets of taps. One table is used for each of these sets of coefficients. For each table, a datum from tap (4 bits) and the index of one of 72 filter coefficients (7 bits) form the address to the table. The output of the table-lookup is a 16-bit number which represents the product of the datum and the filter coefficient pointed to by the 7-bit index. This implementation resulted in substantial savings in hardware which would otherwise be needed if actual multiplications were performed.

As a consequence of using the table lookup technique to perform multiplications, the input data need not be the result of a uniform quantization of the input signal. Each 4-bit datum can be merely some known representation of the signal value, the actual conversion of the representation to its actual value taking place within the table-lookup process itself. Nonuniform quantization provides both a better S/N and less susceptibility of the S/N to changes in the system gain.

3. BANDPASS FILTER 2

Each of 112 74-kHz-wide signals at the output of the first bandpass filter is further filtered into bands that are 1024 Hz wide. Conceptually, this is done with 112 bandpass filters. In actual implementation, this is done with only 28 separate filters, each filter capable of performing the task of bandpass filtering four different 72-kHz-wide signals.

Each of the second-stage bandpass filters consists of a 2016th order FIR filter together with a 144-point DFT.

The FIR filters in the second bandpass filter are implemented with actual multiplications rather than with lookup tables. The representation of the signal at this point is a 16-bit quantity, making any lookup table prohibitively large. The summation of the 14 active taps (eq. (2) of the section, Digital Bandpass Filter) is performed within an integrated circuit multiplier accumulator (MAC). Since the MAC maintains a 35-bit internal accumulator, no precision is lost through the arithmetic operations even though the number of active taps is high. This enables us to implement filters which have stopband rejections that are better than 70 dB.

The 144-point DFT is performed by microprogramming the same special-purpose processor that is used in the first bandpass filter. The transform algorithm consists of a 16-point DFT followed by a 9-point DFT.

The signal at the input of this stage is oversampled by a factor of 2. Thus, of the 144 bands available at the output of each of the second bandpass filters, half are discarded, retaining only 72 bands, each 1024 Hz wide.

4. MAIN MEMORY

The main memory is logically partitioned into blocks of 384 bytes (96k complex words, 16-bit real and 16-bit imaginary components) each. Each block provides enough memory to process a 74-kHz slice of the spectrum in real time. The entire 8-MHz bandwidth (8,257,536 channels, at the 1-Hz resolution) is covered with 112 memory banks, making up approximately 43 Mbytes (10.75 million complex words) of memory. The memory uses dynamic random-access memory circuits. Data are written into the memory from the output of the second bandpass filter (bandwidth of 1024 Hz). The stored data are then read out of the memory and forwarded to the final DFT Processors, in the order that they are expected. Thus, the function of the main memory can be envisioned as a permutation operation on the data.

5. 36- AND 1152-POINT DISCRETE FOURIER TRANSFORMS

Two sets of DFT Processors are used to filter the 1024-Hz signals into the final 32-Hz and 1-Hz bins. Fourteen processors are used in parallel to perform the transforms required to produce the 32-Hz out-

puts, and 28 processors are required for the tranforms which results in the 1-Hz output bins.

The first set of DFT Processors is programmed to implement 36-point Fourier Transforms. The 1024-Hz signals are oversampled by a factor of 1.125. After the 36-point transform, four of the output bins are discarded, leaving 32 bins, each covering 32 Hz of bandwidth. The second set of processors is programmed to perform 1152-point transforms. Only 1024 of the output bins, representing 1024 Hz worth of bandwidth at the resolution of 1 Hz, are retained.

6. DIGITAL BANDPASS FILTERS

The digital bandpass filters in the MCSA are implemented by combining the operations of a finite impulse response filter with an inverse DFT.

An nth order FIR filter (also known as a transversal filter or a tapped delay-line filter) consists of a delay-line of length n. Each of the delay stages has a tap brought out and multiplied by a gain constant. These weighted taps are then summed together at a single common node, forming the output of the FIR filter.

Consider an input sequence to the FIR filter that is in the form of a single impulse at time t_o. As time increases, the impulse appears at successive taps of the delay-line. Since the output is simply the sum of the weighted taps, the response of the FIR filter to a single impulse is just the time-ordering of the weights of the FIR filter, which, by definition, is the impulse response of the filter. Notice that the impulse response is identically zero before time t_o and after time t_{n-1} (n - 1 time units later) -- which leads to the term finite impulse response filter.

For an arbitrary input sequence x_ν, the output sequence of an nth order FIR filter with weights h_μ, $\mu = 0, 1, \ldots, n-1$ is given therefore by

$$y_\nu = \sum_{\mu=0}^{n-1} x_{\nu-\mu}\, h_\mu \qquad (1)$$

A linear time-invariant filter is uniquely determined by its frequency response, which is just the Fourier transform of its impulse response. Thus, by an appropriate choice of the weights for the FIR filter, we can approximate various filter responses -- in particular, lowpass filter responses, which are of primary interest here.

Various techniques exist for the determination of the FIR filter weights (impulse response) for realizing lowpass filters. The filters in the MCSA were designed with the Remez algorithm and the application of Dolph-Cebysev windows.

Now, consider a FIR filter of order nm, for integers n and m, defined by its filter weights h_λ, $\lambda = 0, 1, \ldots, nm-1$.

Given an input sequence x_ν, we define y_ν, $\kappa = 0, 1, \ldots, n-1$ by

https://doi.org/10.1017/S0074180900146716 Published online by Cambridge University Press

$$y_\nu^\kappa = \sum_{\mu=0}^{m-1} x_{\nu-\mu n-\kappa}\, h_{\mu n+\kappa} \tag{2}$$

This is a slight variant from the simple FIR filter in that there are n outputs, each computed as a sum of a subset of the taps of the delay-line. Each of the m components of the impulse response h_μ in equation (2) is termed an active tap of the filter. From this, we take the nth order inverse DFT of y^κ, to obtain the sequence z^κ, i.e.,

$$z_\nu^\kappa = \sum_{\mu=0}^{n-1} y_\nu^\mu\, e^{i(2\pi\mu\kappa)/n} \tag{3}$$

Defining

$$H_\mu^\kappa = h_\mu\, e^{i(2\pi\mu\kappa)/n} \tag{4}$$

equations (2) and (3) can be combined and rewritten as

$$z_\nu^\kappa = \sum_{\mu=0}^{nm-1} x_{\nu-\mu}\, H_\mu^\kappa \tag{5}$$

Notice that equation (5) has the same form as equation (1), exept for a different impulse response for each index κ. From equation (4), we see that each of the impulse responses H^κ is simply the impulse response of the prototype filter, h, that has been translated in the frequency domain by the amount κ/n.

Thus, if the prototype filter were a lowpass filter, each output signal z^k would be a bandpass-filtered output of the input signal x, the location of the passband being a function of κ. If the prototype lowpass filter has a bandwidth of $1/n$, then the κ outputs would represent passbands which are non-overlapping and span the bandwidth of the input signal x.

Furthermore, given that each band is precisely $1/n$ in width by undersampling each of the output signals by a factor of $1/n$, each signal is folded precisely into a lowpass (baseband) signal that is $1/n$ wide, but representing its original bandpass. This implies that the entire arithmetical operation need be done only once each n time units. The computational complexity of an n band bandpass filter has thus been reduced approximately to that of a single FIR filter together with an n-point inverse DFT, at a computational rate of one iteration per n time units.

It should be pointed out that the inverse DFT is used here in an operational sense and not in the manner the inverse DFT is usually interpreted to be. Notice from equation (5) that even though the input signal has gone through a DFT, the signal z^κ, for each κ, remains a time-domain signal; i.e., for a given κ, z_n^k, z_{n+1}^κ, z_{n+2}^κ, ... for a time sequence.

https://doi.org/10.1017/S0074180900146716 Published online by Cambridge University Press

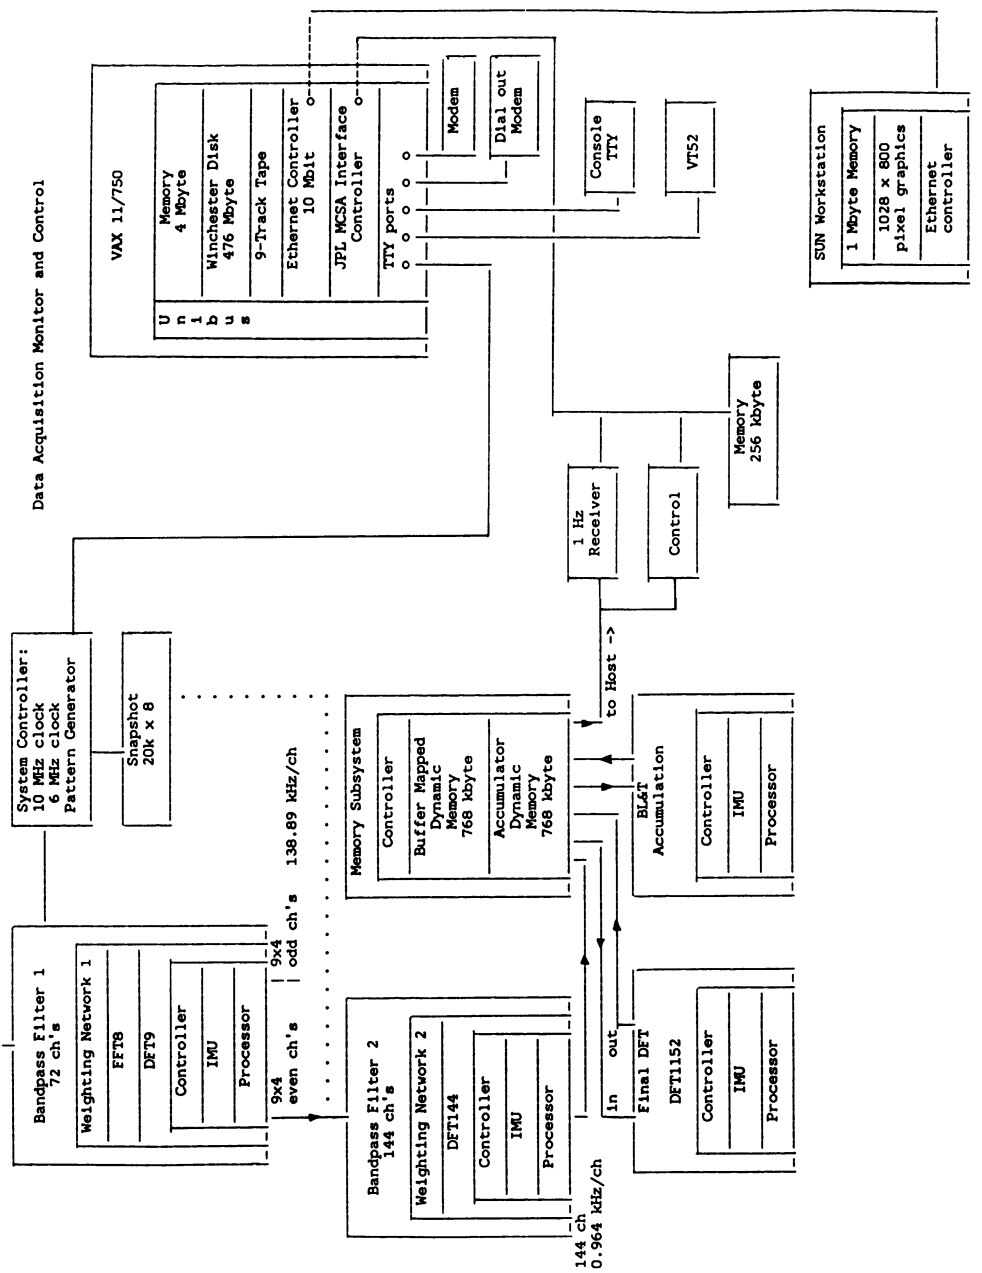

Figure 1. Block Diagram of SETI Prototype MCSA with Monitor and Control Computer

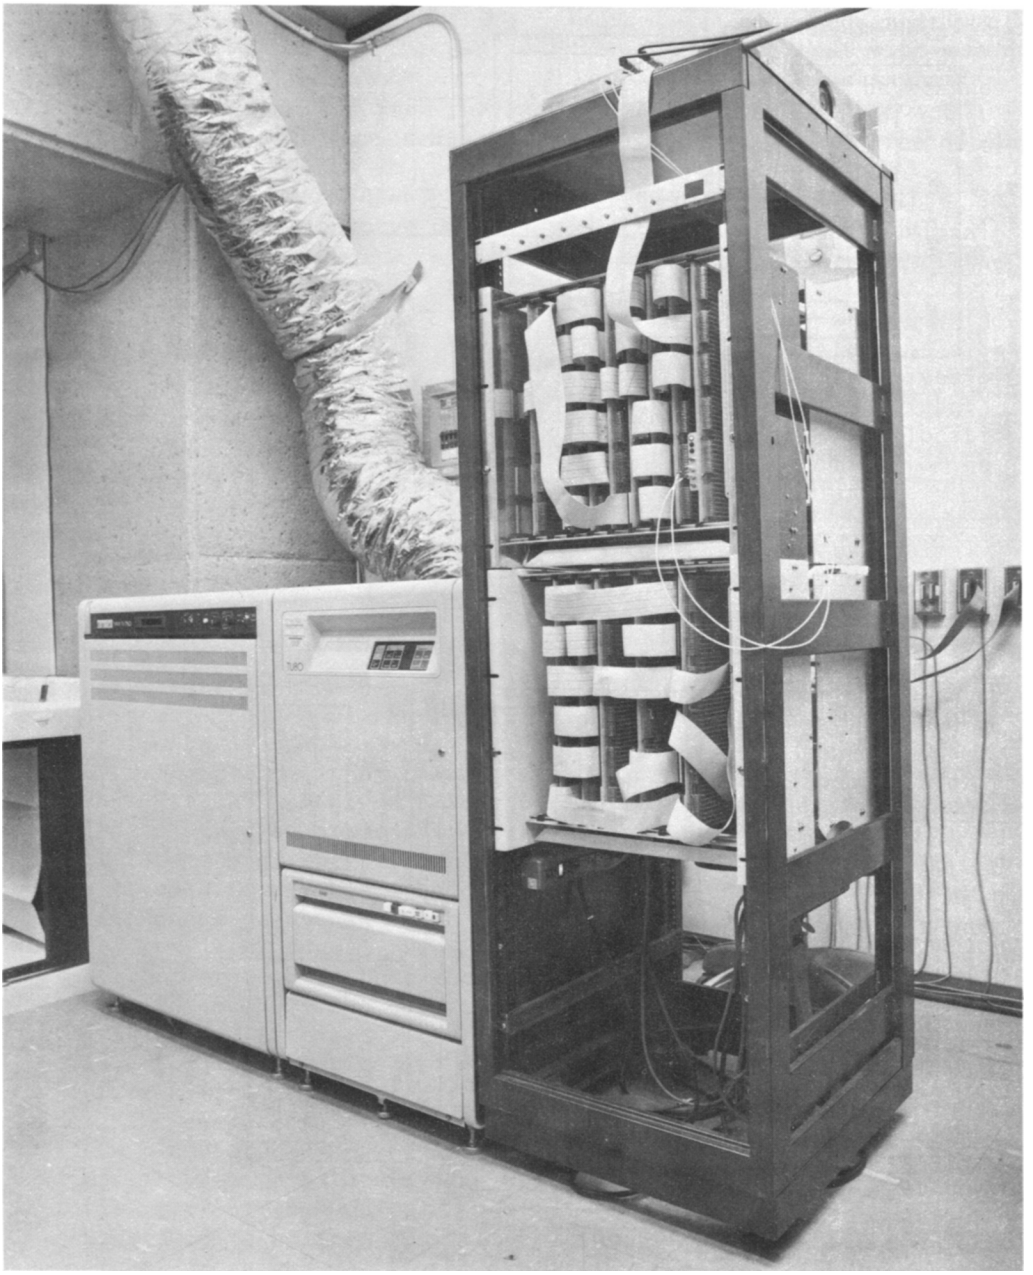

Figure 2. Photo of Prototype MCSA and VAX 11/750 Monitor and Control Computer.

https://doi.org/10.1017/S0074180900146716 Published online by Cambridge University Press

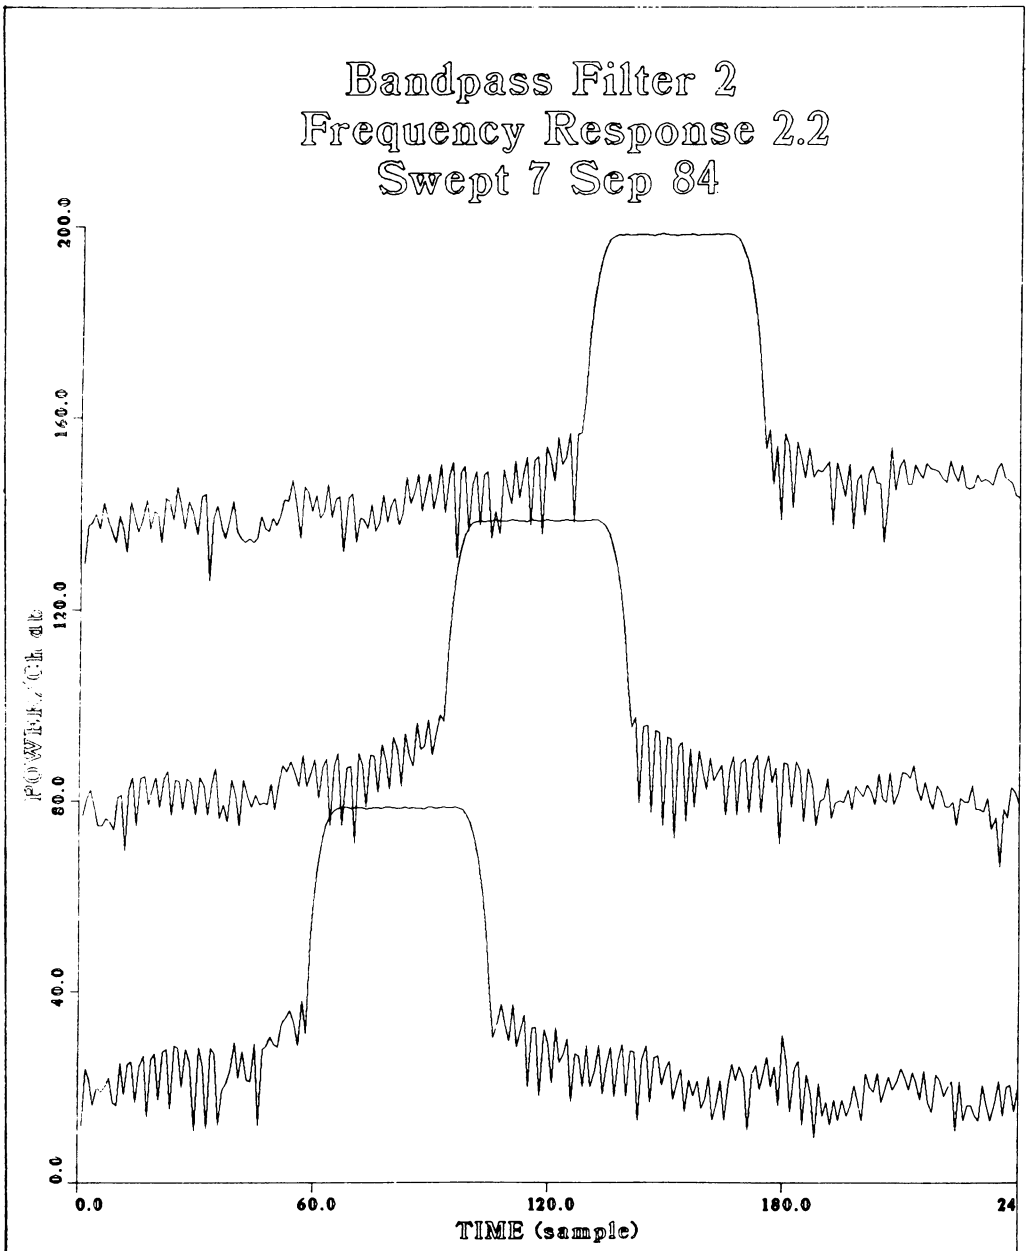

Figure 3. Frequency Response of Three Adjancent Channels of Bandpass Filter Bank Number 2

https://doi.org/10.1017/S0074180900146716 Published online by Cambridge University Press

7. MCSA SYSTEM

Figure 1 is a block diagram of the MCSA system. The overall system is monitored and controlled by a VAX 11/750 computer and selected output channels are displayed on a high resolution (1028 x 800 pixels) graphics unit. Figure 2 is a photograph of the MCSA prototype unit together with the VAX control computer. The discrete Fourier transforms required within the MCSA are computed by microprogrammed pipelined digital signal processors. Each of these processors has two arithmetic logic units (ALU) and a multiplier accumulator (MAC). Each processor can perform 18 million arithmetic operations per second. In the SETI prototype system (74,000 channels), 5 of these processors are required. A complete 8 million channel MCSA will require the parallel operation of many processors each processing different channels of data. Since the different channels of data require identical processing, a single microprogrammed controller can control many processors operating in parallel. The required microprograms are provided by the VAX 750 computer when the MCSA is placed in operation. Different microprograms are required for individual units within the MCSA: i.e. Bandpass Filter Bank 1, Filter Bank 2, Final DFT, etc.

In addition to providing the desired filtering and spectral analysis the MCSA also provides a noise background baseline and identifies signals which exceed predetermined threshold levels above the baseline level.

Figure 3 shows the measured frequency response of three adjacent channels of bandpass filter bank number 2. Notice that adjacent channels overlap at the channel edges where the response is changing rapidly and that the response is very flat within the channel passband region (± .05 dB). This insures that the sensitivity will be very nearly constant at all frequencies. As a result the channel edge regions can be eliminated since signals which occur in the edge of one channel will be observed in the uniform response region of the adjacent channel.

8. REFERENCES

Kolba, D.P., and T.W. Parks, 'A prime factor FFT algorithm using high speed convolution,' IEEE Trans. Acoustics, Speech, Signal Processing, **ASSP-25**, pp. 281-294, August 1977.

Narasimha, M.J., and A.M. Peterson, 'Design of a 24-Channel Transmultiplexer,' IEEE Trans. Acoustics, Speech, Signal Processing, **ASSP-27**, No. 6, pp. 752-762, December 1979.

Narasimha, M.J., and A.M. Peterson, 'Design and Applications of Uniform Digital Bandpass Filter Banks,' International Conf. on ASSP, Tulsa, OK, April 10-12, 1978, IEEE ICASSP Record, pp. 499-503, 1978.

Narasimha, M.J., A.M. Peterson, and S.S. Narayan, 'Implementation of Real-Time Digital Signal Processing Systems,' SPIE 22nd Annual Technical Symposium, Sand Diego, CA, August 28-31, SPIE Proceedings, **154**, pp. 81-90, 1978.

Narayan, S.S., 'Topics in Digital Signal Processing,' Ph.D. Dissertation, pub. Stanford Electronics Laboratories, Stanford University, Stanford, CA, June, 1981.

Narayan, S.S., M.J. Narasimha and A.M. Peterson, 'DFT Algorithms -- Analysis and Implementation,' Tech. Rpt. 3606-12, SU-SEL-78-021, Stanford Electronics Laboratories, Stanford University, Stanford, CA, May, 1978.