CCD FILTER AND TRANSFORM TECHNIQUES FOR INTERFERENCE EXCISION

G. M. Borsuk
R. N. DeWitt

ITT Avionics Division, Electro-Physics Laboratories
Columbia, Maryland

ABSTRACT

The theoretical and some experimental results of a study aimed at applying CCD filter and transform techniques to the problem of interference excision within communications channels are presented. Adaptive noise (interference) suppression can be achieved by the modification of received signals such that they are orthogonal to the recently measured noise field. This technique does not require real-time signal processing but does require a stationary spectral distribution to achieve good performance. On the other hand, real-time excision or limiting of noise can yield the best performance because a stationary noise spectrum is not assumed. However, the implementation of real-time spectrum processing with conventional components can be impracticable. Consequently, CCD techniques have been examined to develop real-time noise excision processing. They are recursive filters, circulating filter banks, transversal filter banks, an optical implementation of the chirp Z transform, and a CCD analog FFT.

I. INTRODUCTION

The purpose of the work presented here was to determine the utility of charge-coupled devices for a signal-to-noise enhancing process here referred to as interference excision. In this introductory section, we will describe the interference excision process, discuss practical considerations that limit its applicability, and describe the role that CCDs can play in extending the range of practicable applicability of the process. In subsequent sections, several different implementations of interference excision processes will be discussed, both analytically and in terms of experimental results. Finally, a section is devoted to intercomparing the various implementations of interference excision, particularly with regard to the range of bandwidths and to the dimensionality of the signal space to which the various processing techniques can be applied.

A. ADAPTIVE NOISE SUPPRESSION

In the presence of stationary white Gaussian noise, the optimum filter that can be used to separate signal from noise is the matched filter which has an equivalent embodiment in the cross-correlation receiver. The specification of such a filter can be derived completely from the known signalling waveform. However, if the noise field in which the signal is immersed is non-white in a way which is in some sense predictable, then there are other forms of filtering which will better serve to improve signal-to-noise ratios than does the matched filter. Generally, these filters will be designed such as to reject some signal energy if by doing so they can be made to reject a disproportionately greater amount of noise energy. Usually, the noise field cannot be sufficiently predicted beforehand to permit optimization of a determined filter. Instead, some form of adaptation of the filter to the realized noise environment is attempted.

There are two general strategies that are employed to effect this adaptation. In the first, a set of m observations are in some sense statistically analyzed for a persistent characteristic that can be used to reject the noise. Most commonly, this characteristic will be the noise's spectral distribution (or in the case of adaptive arraying, the angular distribution). On the basis of these observed characteristics, the filter is optimized using some quantity such as SNR as the measure of performance.

Alternatively, a second strategy can be used. Sets of signal samples are transformed into a representation in which the signal energy is spread among many degrees of freedom of the representation, while the interference is spread among a relatively small subset of the degrees of freedom. Limiting or excising is performed on the data in this representation so that the total noise energy is greatly reduced, while the highly diluted signal is left little affected because none of the excised degrees of freedom contains a significant fraction of the signal energy.

There are some examples where this form of processing has successfully been implemented. The most familiar involves the rejection of impulsive atmospheric noise that occurs in the presence of desired narrow-band signals. In this case, the "excision" is often accomplished by limiting the time samples of the signal space (time domain representation). The bandwidth of the system at the clipping point must be as wide as it can possibly be without including other narrow-band interference that may be present at nearby frequencies. In such a representation, the narrow-band signal of interest is greatly diluted (No one time sample contains a large amount of signal energy.) while the impulsive atmospherics are confined to a relatively small number of the time samples, so excision or clipping in this domain has a disproportionally large effect on eliminating noise energy and

relatively little effect on the signal.

Some experimentation has also been carried out with converse processes in which wideband signals are to be extracted from an interference environment comprising many narrow-band interferers. Now the representation for the limiting or exciser process must be that in which the interferers affect a minimal number of degrees of freedom, i.e., a frequency domain representation in which the analysis bandwidth is comparable to the bandwidth of the interfering signals. However, in order that the desired signal not be rejected, it must be distributed more or less uniformly over much of the analyzed band; i.e., it must be a spread-spectrum signal. For example, the above procedure would be applicable were the system being designed to detect and measure the kind of atmospheric impulses that were being discriminated against in the first example.

This second form of real-time noise excision is somewhat less familiar for two reasons. First, signals of interest have usually been relatively narrow in bandwidth. This is because current allocation of the frequency spectrum largely restrict signals to such a form. Secondly, it is only recently that advances in signal processing techniques have made it possible to even consider processing of this kind except for specialized applications. It is the purpose of this paper to present such considerations.

B. PRACTICAL CONSIDERATIONS LIMITING REAL-TIME ADAPTATION

The present discussion will concern itself with the real-time excision of narrow-band interference from a spread-spectrum space. For an example, we will consider the noise and interference that typically appear in the HF bands. Observations and analyses carried out in the band indicate that the benefits of the excision process improves with increased frequency resolution down to bandwidths as little as a few hundred cycles (Ref. 1). Consequently, the representation desired for the excision process is one in which time is divided into resolvable units measuring 5 to 10 ms and for each resolvable time cell there exists a set of 100 to 200 Hz resolution cells for frequency, respectively. Suppose that a 100 kHz band is to be so resolved into 100 Hz subbands. Then the transformation from the observed time series to the excision representation involves 1000 points. The processor must be capable of performing the transformation in 10 ms.

A transform of this size requires about 20,000 complex multiplies; therefore, a complex multiply must be performed in about 500 nanoseconds. This is faster than can be accomplished with present-day microprocessors, although it is not beyond the reach of some form of processor designed specifically for the function. Furthermore, several less powerful processors can accomplish the task by dividing it up in such a way that they work in parallel. Such a division of labor could take the form of having the processors "leap frog" in time or might involve assigning different processors to operate on portions of the overall frequency passband so that each processor would have a smaller sized transform to perform.

It is the purpose of the present study to attempt to fit such a processing task to the capabilities of a particular kind of processor. Specifically, we will attempt to determine the capability of charge-coupled devices to perform the function, operating in an analog-sampled data mode. Such devices, although not unlimited in dynamic range, may nonetheless be quite suitable for the noise excision processing, and their use would permit considerable simplification of the processor by eliminating the need for analog-to-digital converters, digital multipliers, etc.

Much of the following discussion deals with the implementation of various forms of the Fourier transform. For the purpose at hand, noise excision, it is necessary that the transform process be capable of inversion. This implies that phase information be retained for the various frequency components and precludes the use of some methods, notably the sliding chirped Z algorithm, which for other applications has many attractive features for its implementation with CCDs.

In the discussions of the various spectral analysis techniques that follow, emphasis is placed on the sizes of the transforms that can be accomplished within the current state-of-the-art of CCDs. Generally, the size of the transform will be limited by the number of transfers which can be carried out without suffering significant degradation of the signal from charge transfer inefficiency. This limitation proves to be the one that most severely imposes a limit on transform size. To some extent, the effects of charge transfer inefficiency can be mitigated by techniques such as carrying the signal in two or more adjacent CCD memory locations and by providing vacant cells between the signal bearing cells. However, the use of these methods require the use of longer CCDs clocked at higher rates, and ultimately these magnitudes reach an upper limit which in turn imposes a limit on the transform size.

II. ANALYSES AND RESULTS

A. METHODS USING UNTAPPED CCDs

The simplest form of CCD and that which at this time is most readily available is the untapped analog shift register. Although the output of such a device is nothing more than a delayed version of what was earlier entered into it, these devices constitute a memory that can be

used to perform a spectral analysis.

In this section we treat three different forms of filters implemented with untapped CCDs. There are (1) the recursive filter, (2) the circulating filter bank and the similar coherent memory filter (CMF), and (3) the analog FFT.

Since transformation from the time to frequency domain involves a many-to-one and, conversely, a one-to-many relationship between the input and output of the transformation process, it may be anticipated that access to each data point must be repeated. To accomplish this with an untapped delay line requires that the data be circulated repeatedly through the device. As a consequence, the number of CCD transfers required to accomplish the transformation can be large, and so in the following discussion considerable attention is paid to overcoming the deleterious effects of charge transfer inefficiency.

1. Recursive Filters

Perhaps the simplest form of filtering that can be implemented with CCDs is the recursive filter. Such filters can be implemented with an untapped delay line, so the simplest devices are suitable components of such filters. However, the recursive filter does not provide a full transformation to the frequency domain of the kind required for real-time adaptation. Instead, it provides a single bandpass or bandstop filter. It is possible that such filtering can be adapted to suppress individual strong interferers; because of the potential benefit of this mode of operation under some circumstances, recursive filters are treated here in some detail.

Using a single stage of delay, either a feedback or a bilinear recursive filter can be implemented (Ref. 2). This delay can be implemented with a CCD. If the CCD has n stages, the delay T through the device will be $T = n/f_c$ where f_c is the clocking rate through the CCD, i.e, the clocking rate at which the signal transits from stage to stage. Fig. 1 shows the configuration of the bilinear recursive filter with the various feed-back gains indicated.

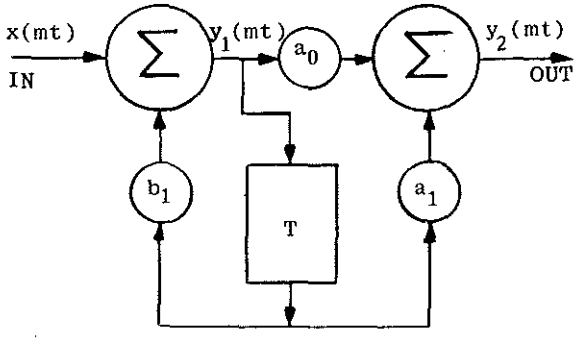

Fig. 1. Canonical form of a recursive filter

If either a_0 or a_1 is set equal to zero, the filter is equivalent to the standard recursive filter. Consequently, the more general bilinear configuration is treated and used to derive the response of the standard filter by setting the appropriate parameter equal to zero.

The response of the filter in Fig. 1 to an input x(mT) is given by a pair of difference equations. Thus,

$$y_1(mT) = x(mT) + b_1 y_1(mT-T) \qquad (1)$$

$$y_2(mT) = a_0 y_1(mT) + a_1 y_1(mT-T) \qquad (2)$$

A solution for the output can be obtained by taking the Z transform, solving for Y (Z) when the input function is a complex exponential, and inverting the transform by the method of residues. The result of this, for the steady state solution, is

$$Y(nT) = \frac{a_0 e^{j\omega T} + a_1}{e^{j\omega T} - b_1} e^{jn\omega T} = H(\omega) e^{j\omega nT} \qquad (3)$$

By suitable choice of the quantities a_0, a_1, and b_1, the transfer function $H(\omega)$ can be made to be that of a low-pass filter, high-pass filter, bilinear low- or high-pass filter, a low-pass canceller, or a high-pass canceller. In Table 1 these quantities are related to the frequency of the peak filter response and the 3 dB corner frequency for the 6 kinds of filters mentioned above.

Feed-back and bilinear recursive filters were designed and tested using the Fairchild CCD-311 and CCD-321 devices. The CCD-311 is a 260-stage two-phase buried channel serial in/serial out device. The CCD-321 consists of two independent 455 stage buried channel serial in/serial out shift registers (Ref. 3). Both types of CCDs were characterized in non-recursive circuits before their use as recursive filter elements. Second harmonic distortion of the CCD-311 was measured at -42 dB for a 945 kHz tone at a clock frequency of 10 MHz while a similar test using the CCD-321 showed second harmonic distortion to be -52 dB.

Fig. 2 shows the swept frequency response of a feed-back low-pass integrator and of the corresponding bilinear integrator using the CCD-311. The total time delay of 51.28 μs (obtained by clocking the CCD at 5.07 MHz) results in poles at frequency intervals of 19.5 kHz. The 3-dB bandwidth of the passband response was measured as 150 Hz, indicating a quality factor, Q, of 130. Fig. 3 shows similar swept frequency results obtained using one serial shift register of the CCD-321.

The harmonic content of single tones applied to both CCD-311 and 321 recursive circuits were studied. Although the plots

of Figs. 2 and 3 show peak to null ratios of ~40 dB for the feed-back circuit, we noted substantial relative increase in harmonic distortion when a single tone was placed exactly at a pole frequency. For example, second harmonic distortion increased 13 dB to a level of -25 dB of the fundamental for a signal of 39 kHz applied to the CCD-311 feed-back filter clocked at 5.07 MHz. Similar results were obtained using the CCD-321 circuit. This non-linear increase in harmonic distortion at the poles can be attributed to the finite linearity of the CCDs themselves. All measurements were made with input signals which were kept small compared to the absolute signal handling capability of the devices. Distortion due to filling the CCD potential wells to a point that interaction occurred with the interface were thus avoided, as were non-linear responses of the input and output stages.

The effects of transfer inefficiency are to shift the position of the pole frequencies from their ideal values (Ref. 4) and to cause the transfer frunction to be nonuniform in frequency. The transfer efficiency of the CCD-321 can be manipulated by changing the voltage of the static electrode set. We were therefore able to demonstrate a shift in the pole frequencies with a change in the transfer efficiency. This shift can be interpreted in terms of a non-linearity of the CCD's phase response with frequency. This increases with increased charge transfer inefficiency, and consequently, the frequency at which there occurs a particular phase shift through the device changes as the charge transfer efficiency changes.

A second effect of charge transfer inefficiency is a diminished amplitude of the device's transfer function as the Nyquist frequency is approached. Not surprisingly, this same diminuition of gain appears in the filter circuits implemented with CCDs. The amplitude of the transfer function of the CCD-311 feed-back low-pass integrator is shown for the frequency range zero to 5 MHz in Fig. 4. In Fig. 5 is shown the change of the transfer function of the CCD-311 bilinear low-pass integrator that results as the relative values of the feed forward coefficients a_0 and a_1 are changed.

2. The Circulating Filter Bank

Although the recursive filter is not suitable for transforming a data set into a representation suitable for narrow-band interference excision, there is another technique somewhat similar to recursive filtering which does provide a full transformation that may be suitable for narrow-band excision and which has an easily implemented inverse transformation. We here refer to the technique as the circulating filter bank. A block diagram of this method is shown in Fig. 6.

The circulating filter bank produces a spectral measure of the form

$$F_k(i) = \sum_{j=0}^{\infty} W^{k(i-j)} e^{-\alpha j} f_{i-j} \qquad (4)$$

where f_j is the input signal and

$$W^{jk} = e^{i2\pi jk/N} \qquad (5)$$

We wish to regard this as a transform and to determine an inverse transformation that permits restoration of the signal in the time domain. We consider the operation

$$g_i = \sum_{k=0}^{N} F_k(i)\, W^{-ki} \qquad (6)$$

$$= \sum_{j=0}^{\infty} f_{i-j}\, e^{-\alpha j} \sum_{k=0}^{N} W^{-ki}\, W^{k(i-j)}$$

$$= \sum_{j=0}^{\infty} f_{i-j}\, e^{-\alpha j}\, N\, \delta_{i,i-j} = N f_i$$

and find it is the desired inverse transform.

To implement this, the output of the filter is multiplied by the kernel function W^{ij} and summed. This can readily be done using an "integrate and dump" scheme shown below in Fig. 7

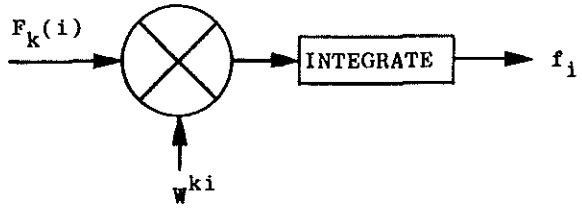

Fig. 7. Inverse transform of circulating filter bank

Note that the sum is only over a single set of N coefficients. Therefore, the dumping can be performed at the end of each circulation of the filter.

The circulating filter bank, shown in Fig. 6, is related to an older spectrum analysis technique which appears in the literature under the name coherent memory filter (CMF). To show their relationship, let the transform computed by the circulating filter be denoted $F_k(i)$ where k is the frequency index and i is the time index. Then we have

$$F_k(i) = \sum_{j=0}^{\infty} W^{k(i-j)} e^{-\alpha j} f_{i-j} \qquad (7)$$

The sum is over j, so the kernal $W^{k(i-j)}$ may be factored into $W^{ki}\, W^{-kj}$ and the first term taken out of the sum and transposed to the left hand side of the equation. Thus, we have

$$W^{-ki} F_k(i) = \sum_{j=0}^{\infty} (W^{-k} e^{-\alpha})^j f_{i-j} \qquad (8)$$

The right hand side of the above equation describes the processing done by the CMF analyzer. That the mixing and single sideband filtering is equivalent to multiplication by W^{-k} may be seen from

$$S(t) \cos 2\pi \frac{k}{N} = S(t) \left(\frac{e^{i2\pi k/N} + e^{-i2\pi k/N}}{2}\right) \qquad (9)$$

The single sideband filtering eliminates the left hand term in parenthesis,

$$\left[S(t) \cos \frac{k}{N}\right]_{\text{filtered}} = S(t) \frac{e^{i2\pi k/N}}{2} = S(t) W^{-k}/2 \qquad (10)$$

In both the CMF and circulating filters, the CCD is required to circulate the information once for each time sample. The number of transfers required then increases as N^2 where N is the size of the transform, if the device can be tailored to the function, or as nN if an available delay of n stages is used (where $n > N$). Because charge-coupled devices are not perfect in their ability to transfer charge, the cumulative effect of many such transfers can be a gradual dispersion of the charge packet representing a signal into following charge packets. This effect places an upper limit on the size of the transformation that can be accomplished. Techniques for increasing the number of filters implementable with a single device are considered next.

Typical charge-transfer devices operate with a charge transfer efficiency of approximately 0.99995. That is to say, each time a charge packet is transferred from one well to another, there is a probability $\varepsilon = 0.00005$ that any given charge quanta will fail to transfer and so will remain in a well where it will serve to corrupt the signal respresented by a succeeding charge packet.

The transfer of charge from one cell to another may be treated as a Bernoulli process with a probability η of success and a probability $\varepsilon = 1 - \eta$ that the charge element will not transfer. If it is attempted to transfer the charge N positions along the device, after the Nth clock cycle, the charge that was initially within a single packet will be distributed among the initial cell, the intended final cell, and the various intervening cells according to the binomial distribution, which may be written

$$p(i,N) = \binom{N}{i} \eta^{N-i} \varepsilon^i \qquad (11)$$

where $\binom{N}{i}$ is the binomial coefficient:

$$\binom{N}{i} = \frac{N!}{(N-i)!\,i!} \qquad (12)$$

The laggard charge is sensed only after it has completely traveled through the device to the output point. The probability that it will require N+k attempted transfers before the specified N successful transfers occur may be written in terms of the binomial distribution as

$$p(N,N+k) = \binom{N+k}{N} \eta^N \varepsilon^k \qquad (13)$$

Parametric computations showing how $p(N,N+k)$ varies with N and k have been made. They show that if it is desired to keep the effects of charge transfer inefficiency to within a few percent, the number of transfers cannot much exceed 1024, nor the number of implemented filters, 32, if this simple mode of operation is to be utilized.

Better performance can be achieved if a cluster of adjacent cells are used to represent each signal sample. The extra cells of each cluster are zeroed with each circulation so that laggard charge is disposed of before it can corrupt the succeeding signal bearing cell, except in the relatively unlikely event that the charge falls far enough behind in a single circulation to reach the following signal sample. This method can provide good isolation of the signal samples, but it suffers from a deterioration of the samples themselves because the laggard charge is lost. To overcome this effect, a second use can be found for the multiple cells of a cluster. In this technique, the signal bearing cell is preceded by one or more cells that are filled to the same level as the signal bearing cell. Then the charge that is lost from the signal bearing cell to the following insulating cells is partially compensated by the charge it gains from the preceding cells.

The effective charge transfer efficiency that can be achieved by these methods can be computed by regarding the lagging of charge during a single pass through the CCD as a transition in a Markov process. Where k-1 preceding cells are filled to the same level as the signal bearing cell, the probability of no transition (no lag) is enhanced by the next k-1 transition probabilities. However, the probability that a signal cell is corrupted by laggard charge from a preceding signal sample must also be increased by summing over k appropriately chosen transition probabilities.

Table 2 shows the transition probabilities that result when a CCD is used in the above described way. The number of signal bearing cells is k; and the number of cells in a cluster is m; so m-k is the number of insulating cells. The transition probabilities are for the specified number of circulations, the number being appropriate for use of the device as a circulating filter bank or coherent memory filter.

The sinusoidal portion of a thirty-two point circulating filter bank has been designed and tested using the CCD-311. The diagram of the analyzer is as shown in Fig. 6. The analyzer was implemented using the techniques addressed above for increasing effective transfer efficiency. Control of the CCD was accomplished by digital logic. Accumulated signals spanned four CCD clock cycles (i.e., four samples) followed by four zeroed reference signals (except for the thirty-second element which consisted of eight zeroes). The fourth sample of each signal set from the CCD output was used to set the signal level for the four new samples taken into the CCD. The algorithm described by equation (5) was implemented using digital shift registers. The computed coefficients were then applied to an analog multiplier after digital-to-analog conversion. Fig. 8 shows several of the initial kernel functions thus generated. Figure 9 shows a complete 32 x 32 data set. Some preliminary experimental results obtained from the analyzer for three different input tones are shown in Fig. 10. The CCD clock rate was 5.1 MHz corresponding to an analysis bandwidth of 5 kHz. A future implementation of this technique is planned using the CCD-321. In this case, one of the 455 bit registers will be used to perform in-phase processing while the other register will be processing the quadrature channel. A ninty-one point transform will be implemented.

3. Analog FFT

Although the recursive filter and the circulating filter discussed above are relatively simple and straightforward to implement, they suffer from a common shortcoming when implemented in so simple a form: there is little control of the window function used in computing the transforms, with the resulting lack of control of sidelobe levels. This shortcoming can be overcome by employing additional delay elements to implement a higher order recursive filter, or to approximate more desirable weighting functions in the circulating filter. This process carried to its logical extreme would require CCDs with as many taps as those required for the implementation of transversal filters of the kind to be discussed in a subsequent section.

There exists at least one other method that can be used to compute Fourier transforms of arbitrarily windowed functions using CCDs that are not tapped. This method derives from the Fast Fourier Transform that is a well known technique in digital signal processing. The following discussion of this method assumes some familiarity with the Fast Fourier Transform, and no effort will here be made to justify or derive the algorithm.

A typical form of the algorithm used in the digital computation of the FFT is called the "in place" transform. Its advantage for this purpose lies in the fact that the results of each butterfly operation can be stored in the same memory locations as were the input data to the butterfly operation. The replaced data are the only data required for that particular operation, and these data are needed nowhere else in the calculation. This property of the algorithm has two advantages: it minimizes the amount of memory required, and it obviates the need to compute more than a single pair of addresses for each butterfly operation. However, for our purpose, where the memory is to be a charge-coupled device operated as an analog shift register, these advantages are not realized. The CCD is not a random access memory but a circulating serial memory. Therefore, a considerable advantage is realized if the FFT algorithm can be implemented in a modified form which assures that the data to be read are immediately available from the circulating memory, and the results of a calculation can be immediately stored into the first available memory locations. The tree graph for such an algorithm is shown in Fig. 11 (Ref. 5). numbers in the vertical columns indicate the positions which the corresponding data points would have taken had the "in place" algorithm been used. The mathematical process (butterfly operation) would, of course, have to be the same as that in the "in place" algorithm. Note the important property of the form of the algorithm for our purpose. Data are taken from "addresses" at the output of the circulating memory and from the midpoint of the memory, and the results of the butterfly operation are stored in adjacent memory locations (of a second memory). The same addressing scheme applies to all levels of the computation, although the arrows indicating the flow of data are shown only for the first level.

Fig. 12 is a block diagram of an analog processor implemented according to this scheme. All arithmetic operations are complex. The weights necessary at each point in the process can be stored in digital form and applied to a digital-to-analog converter before undergoing multiplication. If this procedure were used, the weights could be computed by digital logic as shown in Figure 13.

As described earlier for the other techniques, transfer inefficiency is the limiting physical process here. The effect places an upper limit on the number of transfers which can be performed and, consequently, to the size of the transform. If N is the greatest number of transforms which can be performed, then the largest transform which can be performed by the direct computation of a Fourier transform is $\sqrt{N}$. In comparison, the fast transform algorithm allows an n point transform to be computed with only $n \log_2 n$ transfers, so the upper limit on the size of the transform is governed by the equation

$$n \log_2 n \leq N \qquad (14)$$

The difference can be significant. For example, if 10^4 transfers can be performed, the direct algorithm would be limited to a 100 point transform while a 1000 point transform could be computed by the fast transform method.

C. SPECTRAL ANALYSIS METHODS UTILIZING TAPPED DELAY LINES

A tapped CCD, in which the individual analog samples are accessible for processing, permits the implementation of a transversal filter. We here consider two techniques involving transversal filters: the direct filtering of the signal into bandpass frequency components and the chirped Z transform method. The former technique leads to the signal being distributed on many lines on the basis of frequency but still in a time domain representation. These separated spectral components can be separately limited or excised, and the residuals can be summed back together to reconstitute the desired time series representation of signal and residual noise. The chirped Z transform results in an actual change of representation to the frequency domain, and there is required an inverse transformation if a return to the time domain is necessary.

1. Transversal Filter Bank

By means of a tapped CCD delay line, it is possible to compute the convolution of a discretely sampled signal, f_i, with a stored reference function, g_i. For example, with a device having N taps, the convolution

$$F(m) = \sum_{j=1}^{N} f_{m-j}\, g_j \qquad (15)$$

can be computed. If the function g_i is made to be the discrete Fourier kernel function

$$W^{-jk} = (e^{-i2\pi/N})^{jk} \qquad (16)$$

the result of the convolution is the kth coefficient of the discrete Fourier transform of f_j.

If I and Q phase references are used to reduce the signal to baseband, the signal may be written

$$f_j = \alpha_j + i\,\beta_j \qquad (17)$$

and this expansion may be used to obtain the convolution function in terms of real quantities. Thus

$$F_k(m) = \sum_{j=1}^{N} (\alpha_{m-j} + i\beta_{m-j})\,(C_{jk} + i\,S_{jk}) \qquad (18)$$

$$= \sum_{j=1}^{N} \alpha_{m-j}\, C_{jk} - \sum_{j=1}^{N} {}_{m-j}\, S_{jk} + i \sum_{j=1}^{N} \alpha_{m-j}\, S_{jk} + i \sum_{j=1}^{N} \beta_{m-j}\, C_{jk}$$

Thus, four convolutions are required to compute each complex Fourier coefficient. The first and third terms can be derived from different weights applied to the taps of the same CCD, and similarly the second and fourth terms can be derived from a second CCD's taps.

A method for implementing this approach by means of multiplicative conductive elements only using two identical tapped CCDs is shown in Figure 14. Here, the signal inputs to the tapped CCDs are in phase quadrature. (This condition is assured by using quadrature phase local oscillator injection voltages in the last mixer circuits.) The in-phase and quadrature phase time delayed signals can be separately weighted and combined to synthesize the desired filters. It is to be observed that the two, n tap each, delay lines allow for the generation of n orthogonal filters, since the in-phase and quadrature partial resultants can, themselves, be combined by either addition or subtraction to result in response to either an upper or lower sideband. (For the upper and lower sideband synthesized filter set to be contiguous in frequency, the time delays must be operated at baseband; otherwise, the sideband filter sets will be separated in the manner of hetrodyne images.) Thus, if the latter practice is followed, no penalty in capacity has been paid for the engineering convenience of resistive-only weighting; the number of filters synthesized is equal to one-half the total number of taps.

Alternatively, a relatively high frequency IF can be employed so that image response can be suppressed by conventional filtering at the stage preceding the mixer. A CCD with n taps can then be used to generate n/2 bandpass filters.

A method for cascading filters of this kind has been suggested by White (Ref. 6). In effect, this permits relatively large transforms to be implemented with a common building block--a bank of N filters implemented in the manner discussed above. By this means, N+1 such filter banks can be used to implement a bank of N^2 filters, and a $N^2 + N + 1$ filter bank can implement N^3 filters, etc. It should be noted that each successive level of these CCDs are clocked at a rate of 1/N that of the preceding level, so while the first such CCD may be required to process signals at its maximum clock rate, the subsequent CCDs operate rather below their maximum rates.

The experimental properties of bandpass filters simultaneously synthesized using tapped CCDs have been studied. The device used to perform these experiments was a Westinghouse designed and fabricated twenty tap, four phase, surface channel CCD (Ref. 7). The device has been described in detail elsewhere (Ref. 8). In the following examples, only a single device was used resulting in baseband processing of up to ten independent filters,

(i.e., we assume all image frequencies have been removed) as shown in Fig. 15. As in the case of recursive filter implementations, the eventual need for analog pre- and post-filtering of sampled signals is recognized although not specifically addressed here. Each independent filter consisted of twenty variable resistors operating as the conductances of a current summing differential amplifier. Provision for negative weights was made by providing a negative summing bus which was subtracted from the positive bus at the output of each filter. A sample and hold was used at the output of these differential amplifiers. Each CCD tap output was buffered by a low output impedance, high current gain, line driver to minimize cross coupling between resistor weights of different filters. The uniform weighted impulse response of a filter whose bandpass occurs at 0.1 fc is shown in Fig. 16. The frequency response of the filter is shown centered at 2 kHz when clocked at 20 kHz in Fig. 17. The impulse of this filter when a Hamming window is impassed upon the tap weights is shown in Fig. 18. The impulse response of two adjacent filters, the first centered at 0.1 fc and the second centered at .15 fc are shown in Fig. 19. Fig. 20 shows the frequency response of these filters clock at 20 kHz independently in (a) and (b) and summed together in (c). Besides the Hamming weighting, each set of filter tap weights has imposed upon it a relative uniform weighting to account for the differences in amplitude response between filters that otherwise would occur.

The viability of using this externally weighted approach in a system application depends largely upon the capability of implementing the resistor array as a reproduceable thick or thin film hybrid with sufficient resistor value accuracy. To obtain peak to sidelobe ratios of -40 dB, which is within the capabiltiy of the particular device tested, requires setting tap weights to an accuracy of <1%. This proved to be a formidable task when more than one filter was being synthesized.

2. Optically Tapped Chirped Z Transform

The transversal filters discussed in the preceding section requires the use of a CCD having each stage provided with an output tap. This electrical output can be divided to provide multiple weighted sums. However, it is also possible to employ a delay line that is provided with input taps at each stage. Now, however, there is only a single weighting function that can be applied, i.e., the data set can be convolved only with a single stored function rather than with the multiple functions of the prior case. Nevertheless, there exists a technique whereby the required spectral analysis can be accomplished through convolution with a single stored function, i.e., can be performed with a single transversal filter. In the following discussion, we consider how this single transversal filter canbe implemented with a line imaging CCD of the kind now commercially available. Of course, output taps of the kind discussed above could be used as well.

The discrete Fourier transform may be defined as

$$F_k = \sum_{n=0}^{N-1} f_n \, e^{-2\pi i n k/N} \qquad (19)$$

If the substitution is made,

$$2nk = n^2 + k^2 - (n-k)^2 \qquad (20)$$

equation ( ) can be put into the form

$$F_k = e^{-\pi i k^2/N} \sum_{n=0}^{N-1} (f_n \, e^{-\pi i n^2/N}) e^{\pi i (n-k)^2/N} \qquad (21)$$

If the signals are to be represented as real analog voltages, currents, or quantities of charge, it must be broken into two parallel channels, one representing the reals and the other, the imaginaries. Thus, for example, the set of real data samples $\{f_n\}$ must be multiplied by sinusoidal and cosinusoidal chirps, yielding

$$c_n = f_n \text{ cosine } (\pi n^2/N) = f_n C_n \quad \text{and} \qquad (22)$$

$$s_n = f_n \text{ sine } (\pi n^2/N) = f_n S_n$$

These must be convolved with the function

$$e^{\pi i (n-k)^2/N} = C_{n-k} + i \, S_{n-k} \qquad (23)$$

We here encounter the problem of having to represent a complex function in terms of real physical quantities since we are performing the convolution function by a light sensing CCD utilizing the technique first described by Lagnado and Whitehouse (Ref. 9). A light signal, varying as S_n or C_n is propagated through a mask having a transmission that is proportional to the real or imaginary parts of $e^{\pi i (n-k)^2/N}$. Transmission can be controlled either by graded optical density or by varying the area of transmission in an otherwise opaque mask. However, now both the light intensity, representing S_n or C_n and the mask transmission, representing the Re $e^{\pi i (n-k)^2/N}$ or Im $e^{\pi i (n-k)^2/N}$ are unipolar signals while the functions to be represented by them are bipolar.

One approach is to carry the bipolar signal on a pedestal. The bias levels for the premultipled signal and the convolution kernel may be denoted A and B, respectively. Then, a convolution such as

$$R_k = (A + s_n) * (B + C_{n-k}) \qquad (24)$$

$$= \sum_{n=0}^{N-1} (A + s_n)(B + C_{n-k})$$

is computed. Four resulting terms may be identified: (25)

$$R_k = NAB + A \sum_{n=0}^{N-1} C_{n-k} + B \sum_{n=0}^{N-1} s_n + \sum_{n=0}^{N-1} {}_n C_{n-k}$$

Only the last of these is the desired result, the rest being the spurious result of placing the signals on a pedestal.

Of these extraneous terms, the first is constant, and the second will vary predictably with k. However, the third term

$$B \sum_{n=0}^{N-1} s_n = B \sum_{n=0}^{N-1} f_n S_n \qquad (26)$$

will be a random variable having zero mean and an rms fluctuation determined by the mean temperature of the passband under analysis. This result is undesirable inasmuch as it represents an inability of the spectrum analyzer completely to resolve the band into frequency components. In order to avoid this problem, it is necessary to at least provide parallel convolvers so that, in effect, the mask function can be made to be positive and negative (i.e., the light can be assigned to either the positive or negative channel). The two channels can then be differenced to provide the desired result. Fig. outlines the arrangement of the CCDs and their illuminators.

This technique is being implemented using 256 x 1 linear imagers. To reduce the number of components and the optical complexity, a mask has been developed which spatially multiplexes the chirp waveform corresponding to sine +, cosine +, sine -, and cosine - functions. A computer generated plot of this mask function is shown in Fig.

The practical result that light incident upon one photosite will contaminate its adjacent neighbors either because of diffraction or poor optics has been analyzed. Table lists the computed results of the expected sidelobe response of the analyzer for various operating conditions.

At this time we have not yet demonstrated experimentally this analyzer technique. However, we have demonstrated (as have others, Ref. 9) the convolution properties of a 256 x 1 CCD imager. We obtained the impulse response and autocorrelation of a mask function which consisted of a 127-bit PN sequence.

III. INTERCOMPARISON OF TECHNIQUES

In the preceding sections, we have discussed five methods for transforming in real time from a time series to a relatively narrow band (approximately 100 Hz) frequency representation where effective interference excision can be performed. We here make comparisons among the various methods on the basis of the experimentation and analysis of the study.

Three methods were considered for implementing filters with untapped CCDs. The first of those was the simple recursive filter and various variations of it (e.g., bilinear and canceller forms). This method requires a single CCD delay line for each subband filter to be implemented, and in each such filter there are two or more gains that must be adjusted if the filter is to conform with other such filters. Also, each CCD requires a different clock rate which at the least is an inconvenience (however, recursive filter techniques have been developed which allow a single delay line or two to synthesize more than a single filter (Ref. 10)). Consequently, it seems unlikely that such an approach will be practicable except in cases where the suppression of a relatively small number of interferers offers significant improvement in system performance. In such an application the ability to dither the passband of the filter by control of the CCD clock is an advantage. The effective window of such a filter can be altered in its duration by adjusting the feed-back gain, but any control of the passband beyond that can only be achieved by a multipole filter.

The circulating filter bank, and similar coherent memory filter (CMF), have properties akin to those of the recursive filter except that now it is attempted to implement many filters in time multiplex form. Because the delay is the same for all filters in the set, it is necessary to make a correction of phase from frequency to frequency--in the feed-back loop for the CMF and at the input for the circulating filter bank. However, these methods suffer the same disadvantages as the recursive filter in that the time window is not easily controlled and thus neither are the analyzers' sidelobe responses. Also, unlike the recursive filter, the circulating filter bank requires that the signal at independent degrees of freedom be represented as charge in adjacent cells of the CCD. This places a high premium on the transfer efficiency of the device in order that the signal represented by charge in one cell does not badly corrupt that in another. Techniques were developed to mitigate this problem by representing the signal as one of a cluster of m cells which permitted insulation between adjacent signal cells and some compensation for the charge lost in the transfer process. These methods may permit filter banks numbering into the hundreds but at the cost of aggravating a second limitation--that of the clocking frequency. CCDs usually operate at clock rates to $\approx$ 15 MHz with a few having been tested at >100 MHz. Because the CCD must circulate once between each pair of samples, a bank of N filters, implemented with m cell clusters, can only process samples at a rate $f_c/(NM)$. For N=100 and m=3, a CCD clocked at 6 MHz would permit sampling at only a 20-kHz rate, corresponding to Nyquist bandwidth of 10 Hz. Attempts to implement larger banks of filters with the same resolution would require clocking rates that increase with at least the square of their number.

The analog FFT provides a means for circumventing many of the problems associated

with recursive filters and the circulating filter bank. First, it would not be difficult to introduce arbitrary window functions as the data are initially stored in the CCD analog memory. Because the number of transfers necessary increases only as N log N with the size of the transform, less severe demands are placed on the device's charge transfer efficiency and on the clocking rate. Also, the method requires no more elaborate components in its implementation than do the circulating filter bank and the CMF. Both require a memory for storage of the kernel functions (sines and cosines) and an analog multiplier. The greatest difference is that a more sophisticated control logic is required in the case of the analog FFT. It is projected that 1024 point transforms could be implemented by the FFT approach with currently available serial in/serial out CCDs. Such a transform would require on the order of 10^4 transfers, which at a rate of one MHz, could be completed in .01 seconds--which is approximately the time aperture required for 200 Hz resolution. Such an analysis would handle a sampling rate of 10^5 Hz.

Two methods were considered that made use of tapped CCD delay lines. Because there is no need to circulate the data in these methods, the signal sampling is the same as the CCD's clock rate. However, in the case of the externally tapped delay line filter bank, a point is quickly reached in which the number of pinouts becomes a practical limitation. For example, the device tested had twenty taps. If the technique of cascading is employed as described by White, et al., (Ref. 6), up to one hundred filters could be synthesized by eleven devices. However, unless the weighting array corresponding to the twenty taps of one device were contained in a package approximatley the same size as the device itself (or possibly on a common substrate) the implementation of this many filters would be ungainly.

The chirped Z method is very powerful because a single convolution (transversal filter) serves to analyze an entire set of signal samples. These samples can be taken at the clocking rate, so passbands measured in MHz can be processed by this method. The limitation on the size of the transform is imposed only by the limits in the size of the CCD and associated taps and summing circuitry that can be fabricated. When the convolution is performed by an electrically accessed CCD, the weights for the filter can readily be implemented by the split electrode technique. One advantage that an optically implemented CZT processor could have over its electrically implemented analog is the capability of tailoring a mask to compensate for the effects of charge transfer inefficiency.

There is one respect in which a tradeoff between the chirped Z method and the analog FFT may exist. This is in respect to the precision with which the weights can be implemented. The analog FFT is implemented with untapped CCDs that treat all samples identically, and so while it is necessary to match the performance of pairs of CCDs, the problem of cell-to-cell inqualities is obviated. Such inequalities must be kept rather low in the methods employing tapped delay lines. For example, if the peak response of each filter is to exceed the mean sidelobe level by a factor R, it is necessary that the tap weights be within $\sqrt{N/R}$ units of the correct weights.

IV. CONCLUSIONS

We have analyzed and demonstrated several CCD techniques which may find application in electronic systems which require excision of unwanted interferers (natural or otherwise). Depending upon the relative number of interferers expected in a particular application, a hierarchy of the filter techniques studies here can easily be made in which the complexity of implementation is weighted against the density of excisions required. In the case of just a few interferers, a CCD implemented recursive filter is justified. When a greater number of excisions are required, transform techniques become viable. The circulating filter bank is probably the easiest to implement but also the most limited from the standpoint of transform length. On the other hand, the analog CCD FFT has some attractive features which make it a competitor of the electrical implemented chirp Z technique when the total analysis bandwidth is less than one MHz. The analog CCD FFT uses similar hardware to the circulating filter bank and should not be difficult for us to demonstrate.

V. ACKNOWLEDGEMENT

We would like to cite Mr. Jim Taylor for his expert help in implementing the circulating filter bank and Mr. Joseph Green for his analyses of the transversal filter bank technique. We wish to thank Dr. Dean Baker of NRL for furnishing the tapped CCD. This work has been funded by the Defense Advanced Research Projects Agency under ONR Contract No. N00014-76-C-0496.

VI. REFERENCES

1. Wheeler, J., "Analysis of Noise Measured in the Upper MF-Lower MF Regime in the North Atlantic," ITT-EPL Project Report No. 255, ONR Contract N00014-76-C-0441, June, 1974. (U)

2. Tao, T. F., et al., "Sampled Analog CCD Recursive Comb Filters," Proc. International Conference of Charge-Coupled Devices," pp. 257-266, October, 27, 1975.

3. Wen, D., "A CCD Video Delay Line," Proc. ISSCC, pp. 204-205, February 18-20, 1976.

4. Sequin, C., and Tompsett, M., Charge Transfer Devices, Academic Press, Inc., New York, 1975, p. 216.

5. An Algorithm with similar properties can be found in: Gold, B., and Rader, C., Digital Processing of Signals, McGraw-Hill, New York, 1969, p. 183.

6. White, M., and Webb, W., "Study of the Use of Charge-Coupled Devices in Analog Signal Processing Systems," Final Report, NRL Contract N00014-75-C-0069, May, 1974.

7. White, M., et al., "An Analog CCD Transversal Filter with Floating Clock Electrode Sensor and Variable Tap Gains," Proc. ISSCC, pp. 194-195, February 18 - 20, 1976.

8. White, et al., "Charge-Coupled Device Analog Signal Processing," Final Report, NRL Contract N00014-75-C-0283, February, 1976.

9. Lagnado, I., and Whitehouse, "Signal Processing Image Sensor Using Charge-Coupled Devices," Proc. Int., Conf. Tech. and Appl. of CCD, Edinburgh, September, 1974, p. 198.

10. Mattern, J., and Lampe, D., "A Reprogrammable Filter Bank Using CCD Discrete Analog Signal Processin," Proc. ISSCC, February 12-14, 1975, p. 148.

Table 1. Design Equations for Recursive Filters

	K	a_0	a_1	ω_{peak}	G_p	
STANDARD LOW PASS	$e^{-\omega_x T}$	1	0	0	$\left\|\frac{a_0}{1-K}\right\|^2$	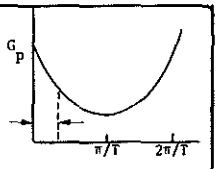
STANDARD HIGH PASS	$-e^{-\omega_x T}$	1	0	π/T	$\left\|\frac{a_0}{1+K}\right\|^2$	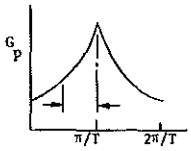
BILINEAR LOW PASS	$\left(\frac{1-\sin \omega_x T}{\cos \omega_x T}\right)$	1	1	0	$\left\|\frac{2a_0}{1-K}\right\|^2$	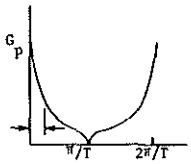
BILINEAR HIGH PASS	$-\left(\frac{1-\sin \omega_x T}{\cos \omega_x T}\right)$	1	-1	π/T	$\left\|\frac{2a_0}{1+K}\right\|^2$	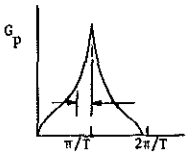
	K	a_0	a_1	ω_{peak}	G_p	
LOW PASS CANCELLER	$-\left(\frac{1-\sin \omega_\delta T}{\cos \omega_\delta T}\right)$	1	1	0	$\left\|\frac{2a_0}{1-K}\right\|^2$	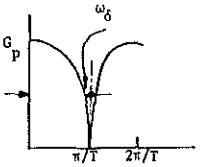
HIGH PASS CANCELLER	$\left(\frac{1-\sin \omega_\delta T}{\cos \omega_\delta T}\right)$	1	-1	π/T	$\left\|\frac{2a_0}{1+K}\right\|^2$	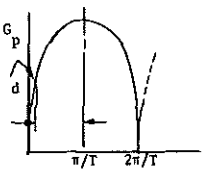

NOTATION:

- G_p peak gain
- ω_x frequency interval between frequency of peak response and 3 dB corner frequency
- ω_δ frequency interval between cancelled frequency and 3 dB corner frequency $\omega_\delta = \pi/T - \omega_x$
- T delay time. For CCD delay $T = f_c/N$
- f_c CCD clock frequency
- N number of CCD stages

Table 2. Transition Probability

NO. OF CIRCULATIONS	k	m	N	π_0	π_1	π_2	π_3	π_4
128	2	3	384	0.983	1.47×10^{-4}	1.12×10^{-8}	5.54×10^{-13}	2.05×10^{-17}
256	2	3	768	0.842	1.99×10^{-3}	2.35×10^{-6}	1.84×10^{-9}	1.08×10^{-12}
512	2	3	1536	0.243	8.93×10^{-3}	1.64×10^{-4}	2.00×10^{-6}	1.82×10^{-8}
512	3	3	1536	0.989	3.624×10^{-2}	6.627×10^{-4}	8.066×10^{-6}	7.350×10^{-8}
512	3	4	2048	0.992	2.047×10^{-3}	2.219×10^{-6}	1.601×10^{-9}	8.644×10^{-13}

Table 3. Effects of Optical Diffraction and Distortion Upon Optical CZT Performance

	QUANTIZING LEVELS				
	3	10	20	30	100
UNIFORM					
No Pedestal					
0 Spill				-13.291	-13.473
0.1 Spill					-13.291
0.2 Spill					
Pedestal					
0 Spill				-13.291	-13.382 (-13.324)
0.1 Spill					
0.2 Spill					
HAMMING					
No Pedestal					
0 Spill	-12.83	-27.943	-33.04	-37.688	-42
0.1 Spill	-12.56	-27.310	-29.131	-32.317	-33.592
0.2 Spill		-24.579			-27.674
Pedestal					
0 Spill				-25.945	
0.1 Spill		-25.581		-24.761	-24.761
0.2 Spill		-24.306			

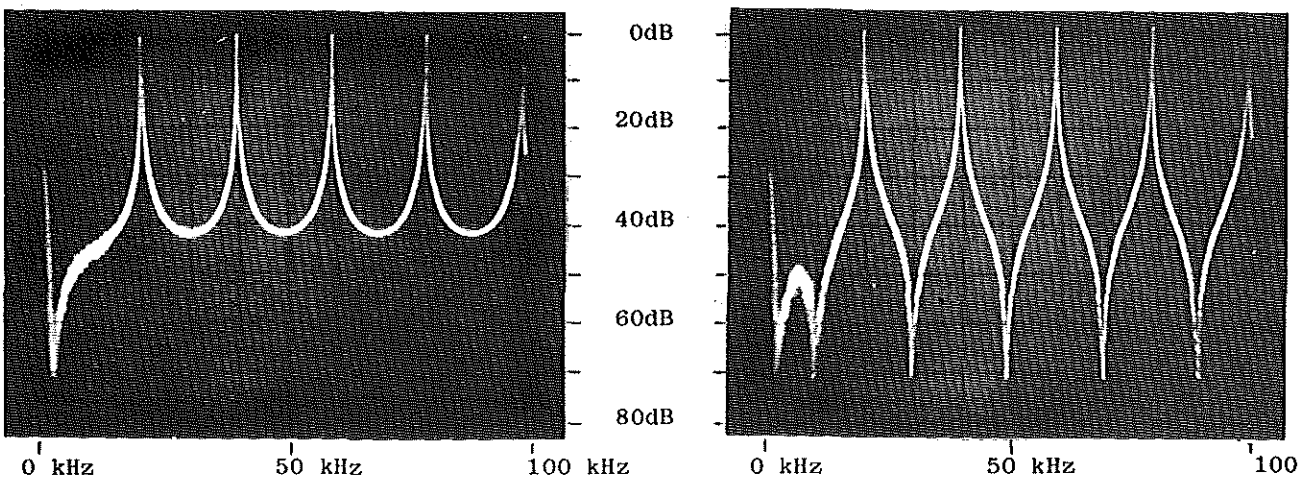

Fig. 2. Swept Frequency Response of Recursive Filters Using the CCD-311

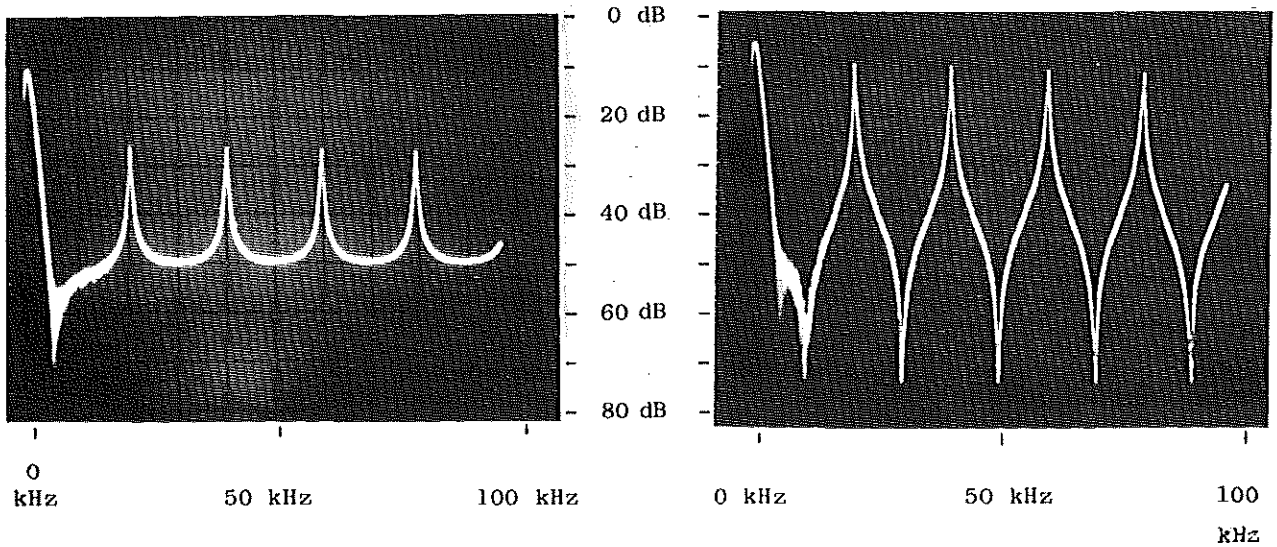

Fig. 3. Swept Frequency Response of Recursive Filters Using the CCD-321

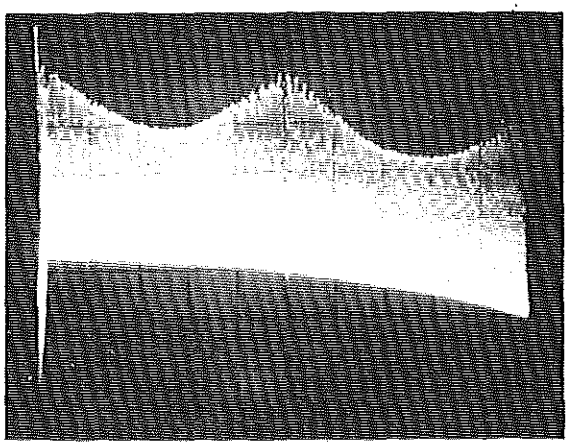

Fig. 4. Zero to 5 MHz Frequency Response of CCD-311 Feed-back Filter

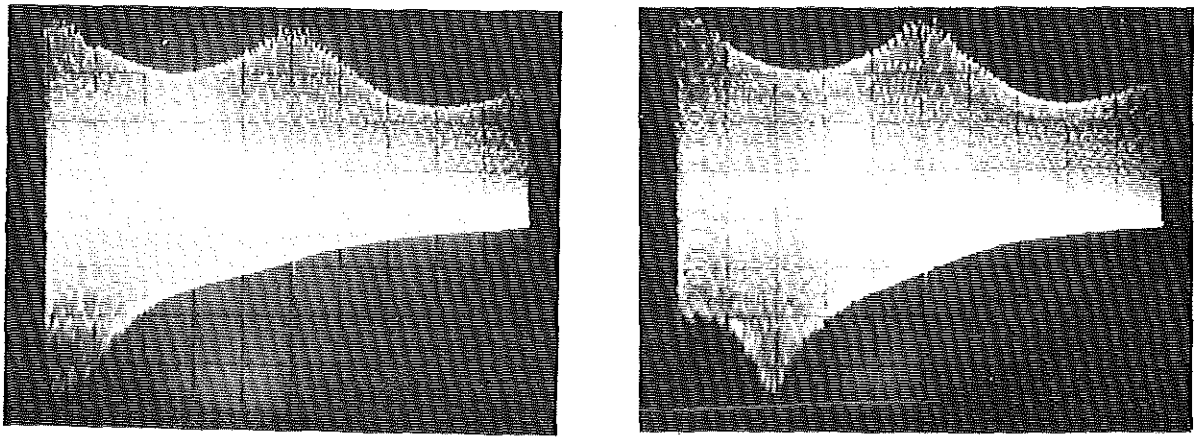

5(a) 5(b)

Fig. 5. Zero to 5 MHz Frequency Responses of CCD-311 Bilinear Filter for Several Relative Values of a_0 and a_1

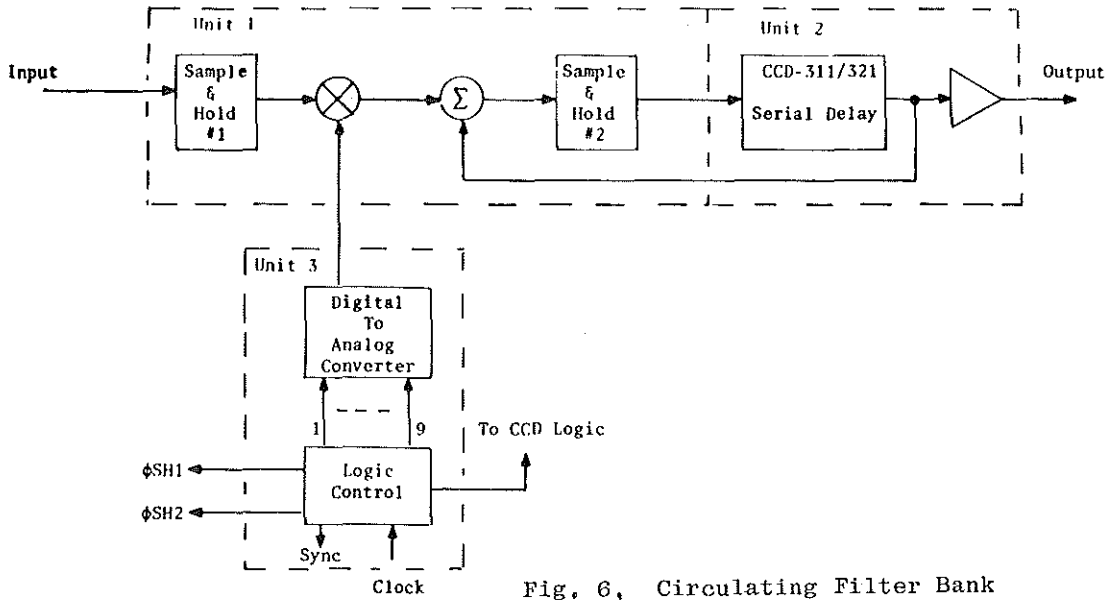

Fig. 6. Circulating Filter Bank

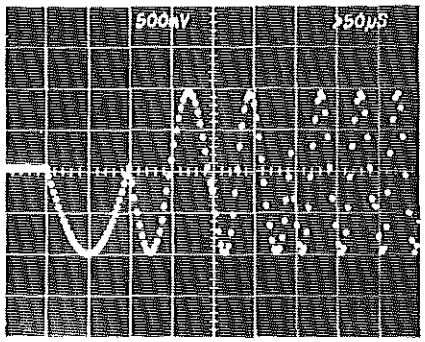

Fig. 8. DAC Output Showing Several Kernels

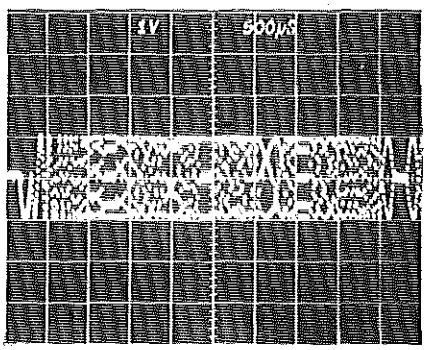

Fig. 9. The Kernel Function

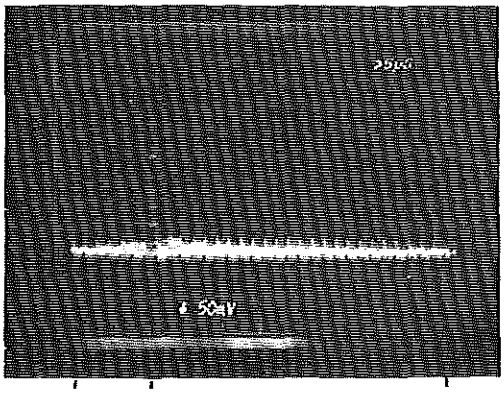

0.0 kHz 1 kHz 5 kHz

10(a)

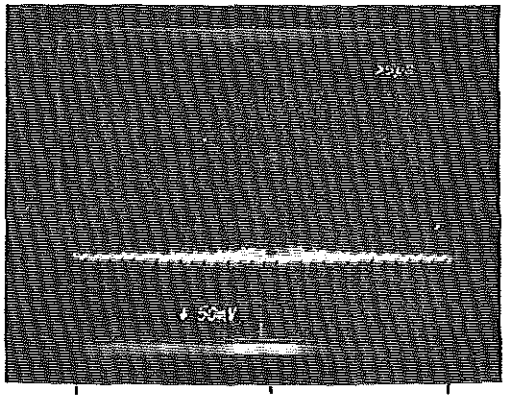

0.0 kHz 2.6 kHz 5 kHz

10(b)

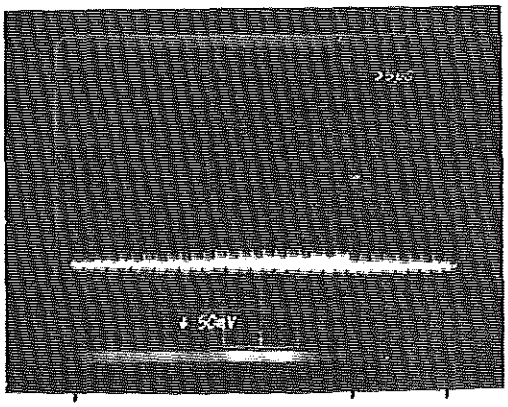

0 kHz 3.75 kHz 5 kHz

10(c)

Fig. 10. Circulating Filter Bank Using the CCD-311

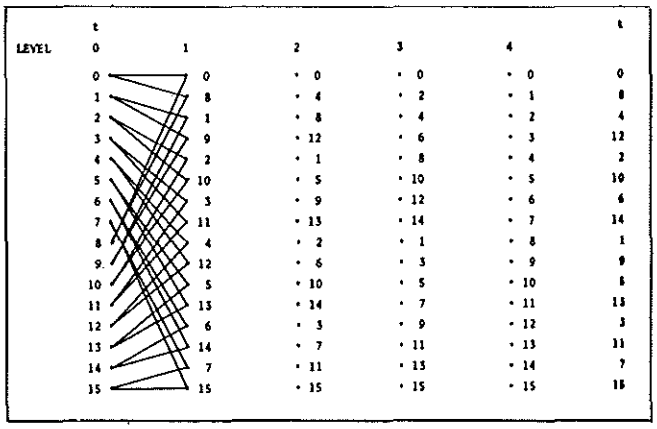

Fig. 11. Modified Tree Graph

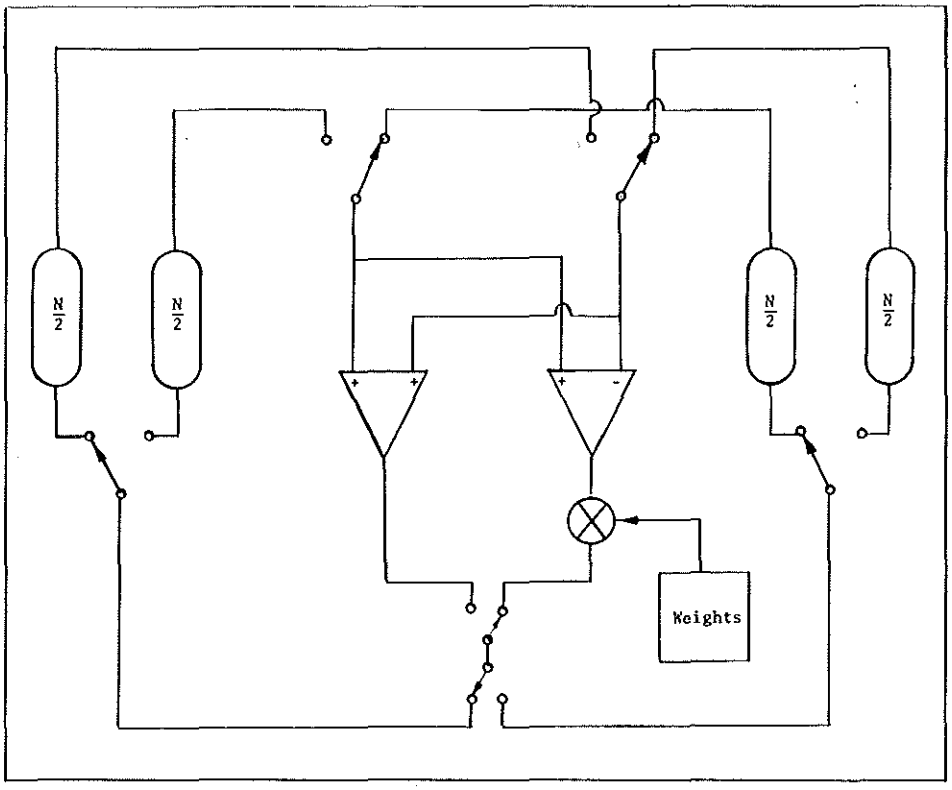

Fig. 12. Analog FFT Implemented with CCDs

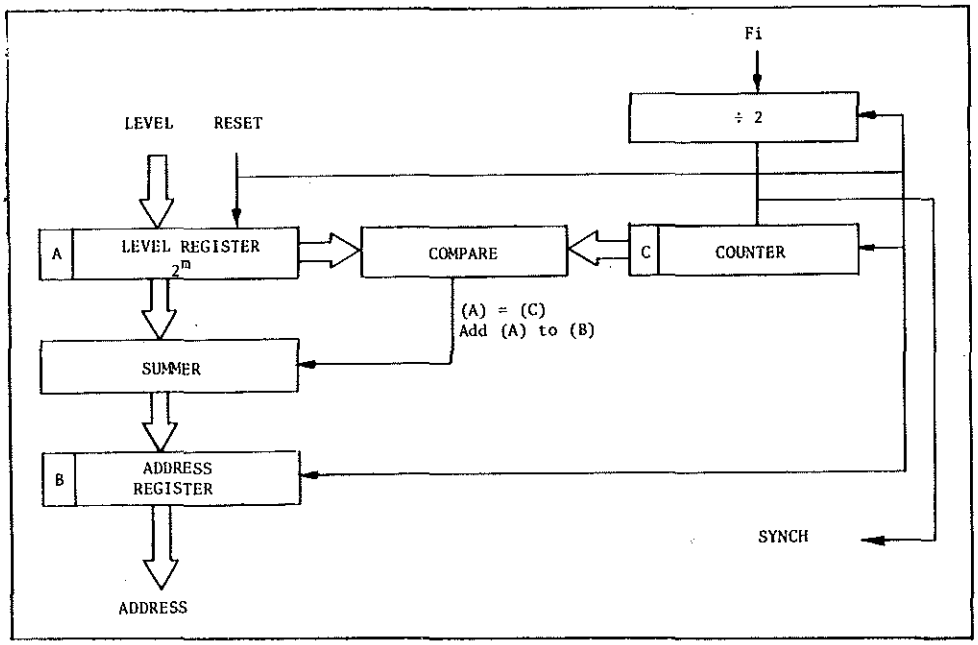

Fig. 13. Digital Logic for Generating Weight Addresses for Analog FFT

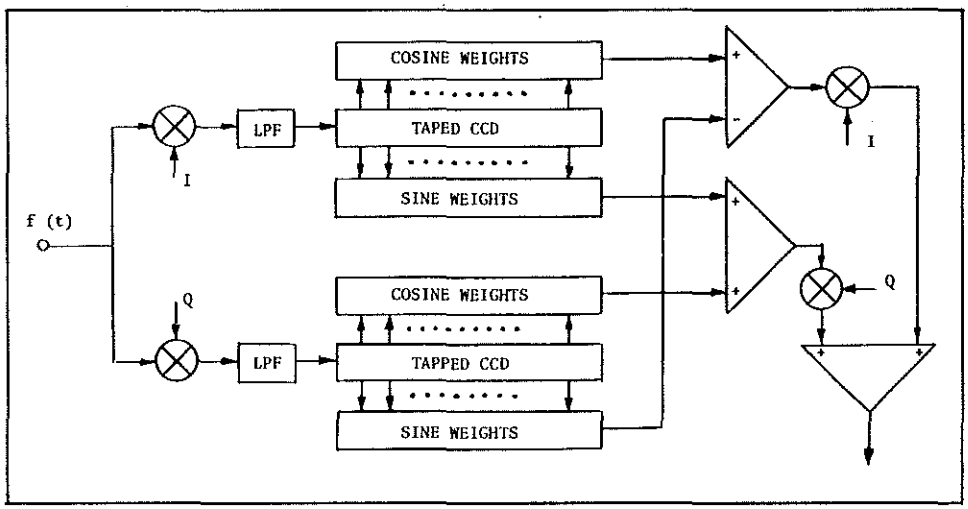

Fig. 14. Implementation of Transversal Filters Using Tapped CCDs

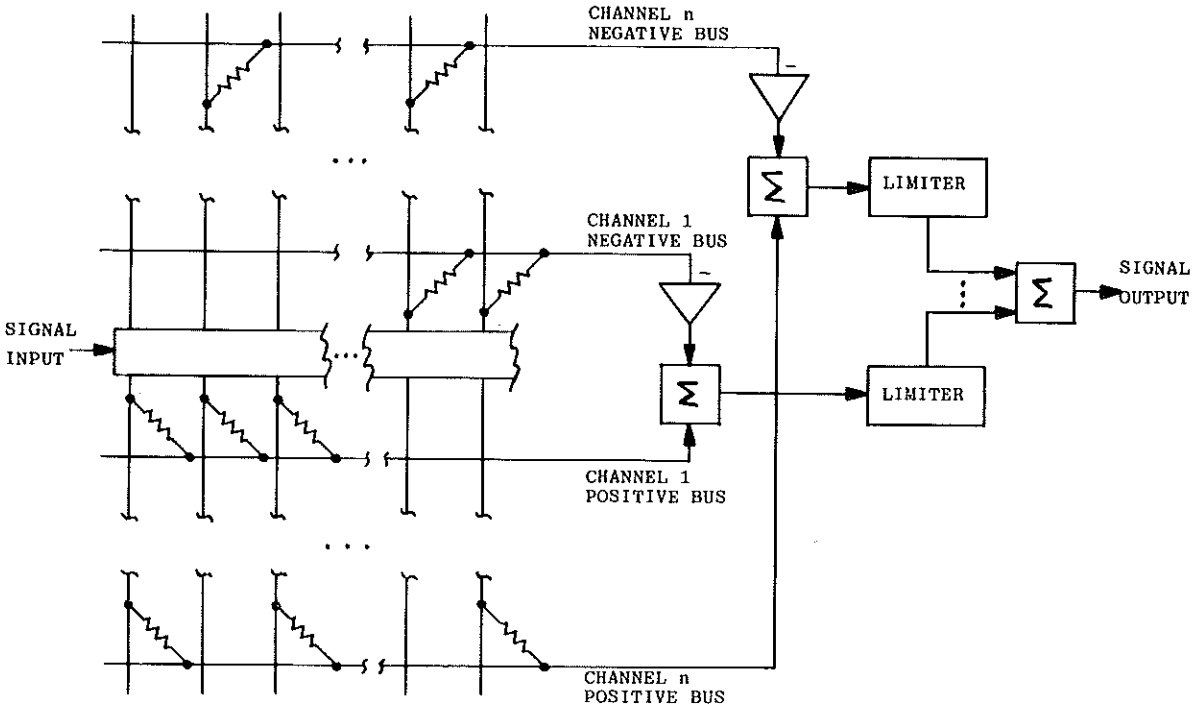

Fig. 15. Implementation of Transversal Filter Bank With Resistive Components

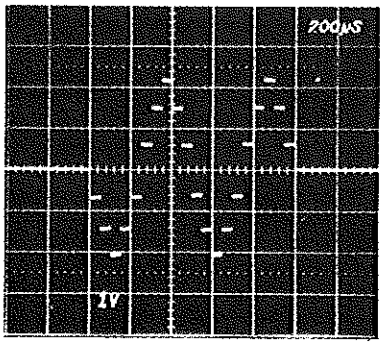

Fig. 16. Impulse Response of a Uniformly Weighted Transversal Filter for a Passband at 0.1 f_c

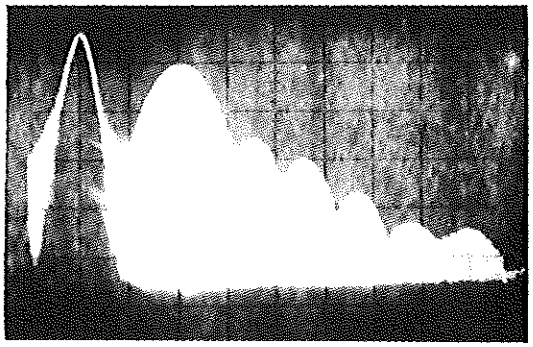

Fig. 17. Frequency Transfer Function for Uniformly Weighted Transversal Filter at 0.1 f_c

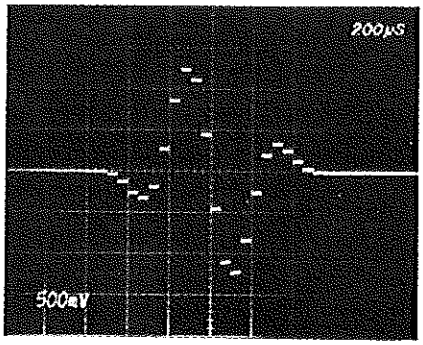

Fig. 18. Impulse Response of Transversal Filter with Hamming Window

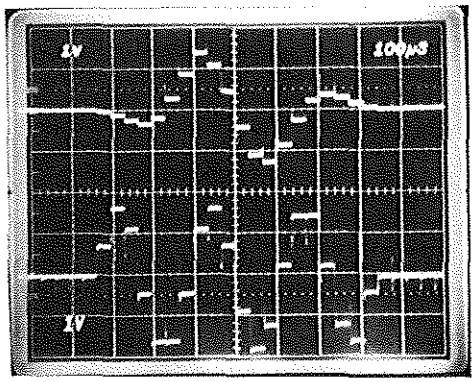

Fig. 19. Impulse Response of Two Simultaneously Synthesized Filters

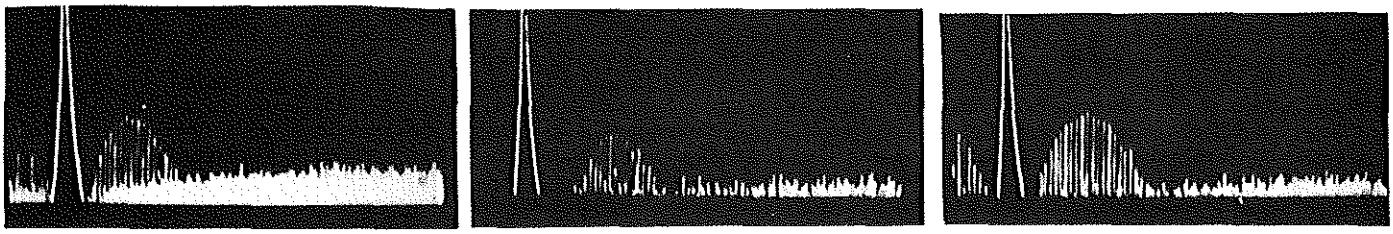

20(a) 0.1 f_c 20(b) 0.15 f_c 20(c) combined

Fig. 20. Transfer Functions of Adjacent Transversal Filters f_c = 20 kHz

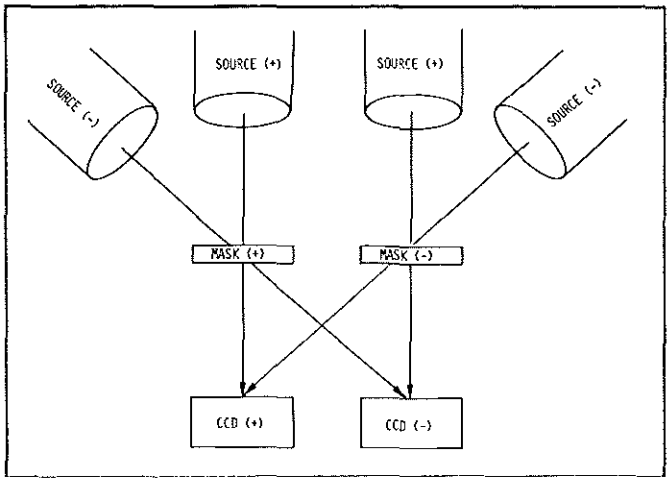

Fig. 21. Arrangement of Illuminators for Optically Implemented CZT-Diagonal Illuminators Can Be Eliminated if Pedestal on Mask Function is Tolerable

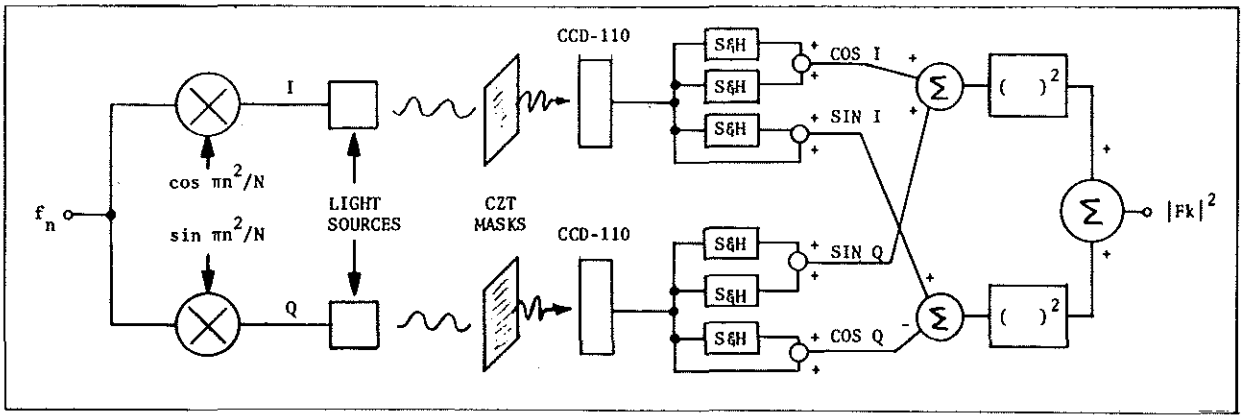

Fig. 22. Block Diagram of Circuit for Optically Implemented CZT

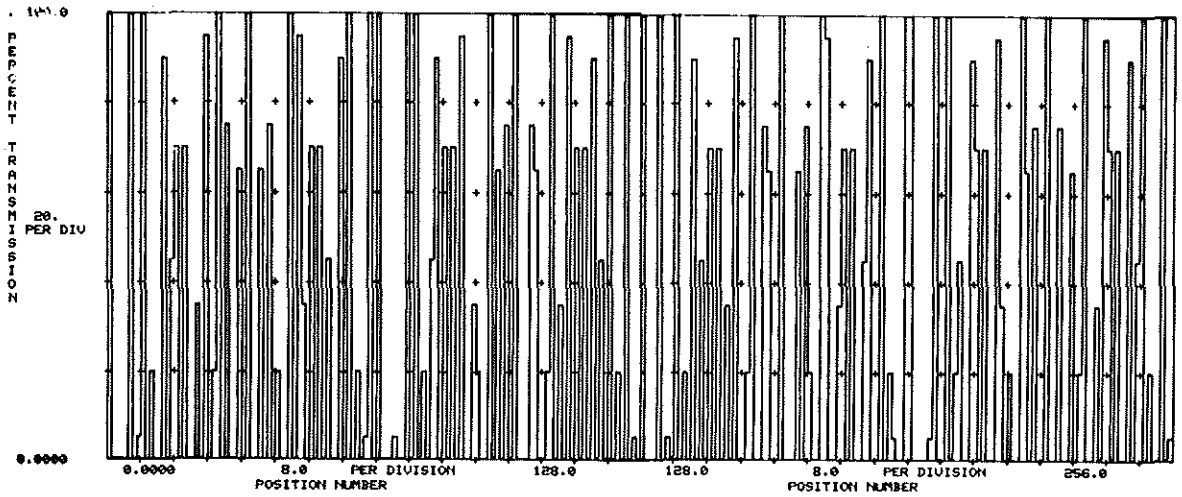

Fig. 23. Mask Function for Optical CZT