FFT-BASED INCREMENTAL REFINEMENT OF SUBOPTIMAL DETECTION

Joseph M. Winograd[†] *S. Hamid Nawab*[†] *Alan V. Oppenheim*[‡]

[†]ECS Department, Boston University, Boston, MA 02215

[‡]Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139

ABSTRACT

In the context of FFT-based maximum-likelihood (ML) detection of a complex sinusoid in noise, we consider the result of terminating the FFT at an intermediate stage of computation and applying the ML detection strategy to its unfinished results. We show that detection performance increases monotonically with the number of FFT stages completed, converging ultimately to that of the exact ML detector. The receiver operating characteristic associated with the completion of each FFT stage is derived. This enables the calculation of the minimum number of FFT stages that must be completed in order for desired detection and false alarm probabilities to be obtained.

1. INTRODUCTION

There is currently a growing interest in approximate digital signal processing techniques that are characterized by the fact that their intermediate results represent successive approximations to some desired solution. The improvement in solution quality from one intermediate result to the next may be viewed as a process of *incremental refinement* [1] [2]. In this paper we demonstrate that, in the context of detecting sinusoids in noise, the FFT possesses the incremental refinement property. By considering the performance of the maximum-likelihood (ML) detection strategy applied after successive FFT stages, we show that the performance of the resulting suboptimal detector improves incrementally, converging ultimately to that of the exact ML detector. This leads to important consequences such as the fact that for a wide range of SNR values at the input of the FFT, high probabilities of detection are obtained without the necessity of going to the last stage of the FFT.

We begin by describing the traditional FFT-based approach to ML detection and the analysis of its performance. Next, we consider the data obtained at intermediate stages of the FFT in the context of ML detection and show that at each successive stage of computation the effective SNR is doubled while the number of channels which could contain signal energy is halved. Compact expressions for the probability of detection, probability of false alarm, and the receiver operating characteristic (ROC) are then derived. We conclude with an illustration of how these performance results can be applied to the problem of selecting the minimum number of FFT stages that must be completed in order to achieve a desired level of detector performance.

This work was sponsored in part by the Department of the Navy, Office of the Chief of Naval Research, contract number N00014-93-1-0686 as part of the Advanced Research Projects Agency's RASSP program.

2. FFT-BASED ML DETECTION

The detection of a complex sinusoid of unknown frequency and phase in additive white Gaussian noise (WGN) can be formulated as a decision D between the two alternative hypotheses:

$$\begin{aligned} H_w &: x(n) = w(n), \\ H_s &: x(n) = s(n) + w(n), \end{aligned} \tag{1}$$

where $x(n)$ is the received data sequence, observed for $n = 0, 1, \ldots, N-1$, $w(n)$ is the white noise process, and $s(n)$ is the sinusoid whose detection is desired. The hypothesis H_w represents the case when only white noise is present, and H_s is the signal-present hypothesis.

We consider the detection of complex sinusoids of the form

$$s(n) = \sqrt{E} e^{j\frac{2\pi}{N}ln + j\phi}, \quad n = 0, 1, \ldots, N-1, \tag{2}$$

where E is the signal power (which is known), l is an unknown integral frequency index in the range $0 \le l \le N-1$, and ϕ is the unknown phase with possible values $0 \le \phi < 2\pi$. The complex-valued noise process $w(n)$ with power spectral density $N_0/2$ can be defined as

$$w(n) = q(n) + jr(n), \quad n = 0, 1, \ldots, N-1, \tag{3}$$

where $q(n)$ and $r(n)$ are both real-valued WGN processes with variance $N_0/4$.

The ML detector [3] for $s(n)$ consists of a bank of correlators followed by a comparator of their outputs and a threshold detector. Each correlator can be thought of as producing at its output the magnitude-squared of the output obtained at time $n = N-1$ from a filter matched to one of the possible sinusoidal frequencies. Denoting by $C(k)$ the output from the correlator associated with frequency index k, we obtain

$$C(k) = \left| \sum_{n=0}^{N-1} x(n) e^{-j\frac{2\pi}{N}kn} \right|^2, \quad k = 0, 1, \ldots, N-1. \tag{4}$$

This output is equivalent to the magnitude-squared of the N-point DFT of $x(n)$ and is typically implemented using the FFT algorithm at a significantly reduced computational cost in comparison with a filter-based implementation.

The ML detection strategy dictates that the output of these correlators be compared by selecting the maximal value over all $C(k)$. If this value is greater than a threshold η, then the sinusoid is determined to be present (i.e. $D = H_s$), otherwise it is declared absent (and $D = H_w$). Using the Neyman-Pearson detection criterion [3], the threshold value is selected so that a fixed false alarm probability ($P_{FA} = \text{Prob}[D = H_s | H_w]$) is obtained.

0-7803-3192-3/96 $5.00©1996 IEEE

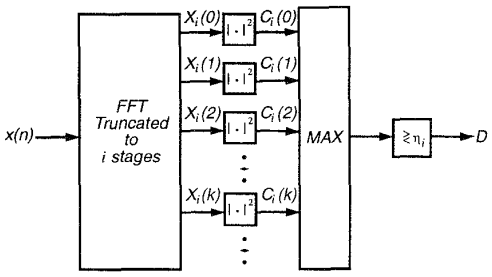

Figure 1. Incremental refinement detector of sinusoids in noise.

The probability of detection and ROC for this detector are determined by forming the distribution of the maximum energy value found across all elements of $C(k)$ under each input hypothesis [4]. Under hypothesis H_w, the FFT output consists of N complex-valued random variables each with real and imaginary parts that are independently Gaussian-distributed with zero mean and variance $N \cdot N_0/4$. This results in values of $C(k)$ that are independently χ^2-distributed with two degrees of freedom. Under hypothesis H_s, the FFT output for $k = l$ is equal to $N\sqrt{E}e^{j\phi}$ perturbed by a complex-valued noise component with independent Gaussian-distributed real and imaginary parts of zero mean and variance $N \cdot N_0/4$. Hence, $C(l)$ is noncentral χ^2-distributed with two degrees of freedom and noncentrality parameter N^2E. The values of $C(k)$ for $k \neq l$ have the same distribution as for the noise-only case.

In the following section, we derive the distributions obtained for $C(k)$ when only a subset of FFT stages are completed. These distributions are then used to obtain the corresponding detection performance.

3. DATA ANALYSIS FOR FFT STAGES

For applications in which further reduction in computation is desired, one may consider the result of terminating the FFT algorithm after an intermediate stage of computation and using its incomplete results as the basis for detection. The structure of a detector employing this approach is illustrated in Fig. 1, where we denote by $X_i(k)$ the output of the ith FFT stage and $C_i(k)$ is the associated correlator output.

In this paper, we consider the result of applying this detection strategy when $N = 2^v$ (for some positive integer v) and the computation of $X_i(k)$ is performed using either the decimation-in-time (DIT) or the decimation-in-frequency (DIF) radix-2 FFT [5]. Both of these FFT algorithms are comprised of v successive stages of computation, and each stage performs $N/2$ parallel "butterfly" operations on the results of the previous stage (or the input data, in the case of the first stage).

In order to ascertain the performance of this detector, we first consider the contents of $X_i(k)$ under each of the two input hypotheses. When only noise is present in the input sequence, we have $X_i(k) = W_i(k)$, where $W_i(k)$ denotes output of the ith stage of an FFT which has the sequence $w(n)$ as its input. Due to the linearity of all FFT operations, under hypothesis H_s we have $X_i(k) = S_i(k) + W_i(k)$, where $S_i(k)$ is the output of the ith stage of an FFT for which the input sequence is $s(n)$. We proceed in our analysis by independently considering the contents of $S_i(k)$ and $W_i(k)$ at each stage. These results are then used to obtain the distribution of $C_i(k)$ under each input hypothesis.

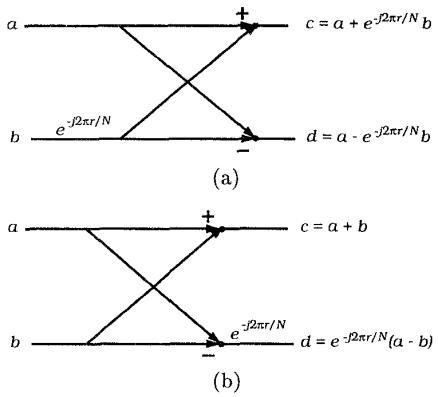

Figure 2. The butterfly computation structure used in the radix-2 (a) decimation-in-time and (b) decimation-in-frequency FFT.

3.1. Signal Data Analysis

At the input to the FFT, the complex sinusoid $s(n)$ consists of N elements of magnitude $\sqrt{E}$ as given in Eq. 2. At the output of the last stage, $S_v(k)$ contains only a single non-zero element at $k = l$, having the value $N\sqrt{E}e^{j\phi}$. If we consider the intermediate results obtained at each intervening stage, moving one stage at a time from the output of the last stage towards the input, we find that the number of non-zero elements doubles at each stage and the magnitude of each of these non-zero elements is successively halved. This results in $S_i(k) = 2^i\sqrt{E}e^{j\theta_k}$ for $N/2^i$ different values of k, and $S_i(k) = 0$ for all other values of k between 0 and $N-1$.[1]

Proof of this assertion for both the DIT and DIF radix-2 FFT can be obtained through careful examination of the FFT structure. That this relationship holds between the last two FFT stages can be shown using a basic property of the butterfly computation. For both the DIT and DIF butterfly operations (shown for reference in Fig. 2), when $c = Ae^{j\theta}$ and $d = 0$ (or when $c = 0$ and $d = Ae^{j\theta}$) for any A and θ, it must be the case that $a = Ae^{j\psi}/2$ and $b = Ae^{j\lambda}/2$ for some ψ and λ. The values of $S_{v-1}(k)$ must therefore consist of two elements with the values $N\sqrt{E}e^{j\psi}/2$ and $N\sqrt{E}e^{j\lambda}/2$ for some ψ and λ. The other elements of $S_{v-1}(k)$ are zero-valued, since a butterfly which is zero at both of its outputs must also have zero-valued elements at both of its inputs.

It is a further property of both the DIT and DIF FFT structures that each of the two non-zero elements of $S_{v-1}(k)$ lie at the outputs of distinct butterflies. Thus, this same property holds for each of them, and $S_{v-2}(k)$ must consist of exactly 4 non-zero elements, each with magnitude $N\sqrt{E}/4$. This behavior, in fact, occurs at *every* stage[2] of the radix-2

[1]The quantities θ_k are random variables which we need not characterize for the purposes of this analysis. We also do not require knowledge of the locations (values of k) of the non-zero elements of $S_i(k)$ at each stage.

[2]This can be seen informally by tracing all paths in a flowgraph representation of the radix-2 FFT which lead to a single output point. A formal proof is given in the appendix.

FFT. This results in twice as many non-zero elements with half the magnitude in each preceding FFT stage.

3.2. Noise Data Analysis

As discussed in Sec. 2, when the input sequence to the FFT is $w(n)$, the output of the last FFT stage, $W_v(k)$, consists of N independent complex-valued random variables with real and imaginary parts that are each independently Gaussian distributed with zero-mean and variance $N \cdot N_0/4$. By solving the DIT and DIF butterfly equations for their inputs in terms of there outputs, the elements of $W_{v-1}(k)$ can be shown to have the same distribution as $W_v(k)$, but with the variance of their real and imaginary parts halved to $N \cdot N_0/8$. Following the computation across successive stages, it can be shown that when i stages have been performed, the values of $W_i(k)$ are independent complex-valued random variables with Gaussian-distributed real and imaginary parts each having variance $2^{i-2}N_0$.

3.3. FFT Stage Distributions

The results of our signal and noise data analysis for intermediate FFT stages can be applied to determine the distribution of $C_i(k)$ under each of the two input hypotheses.

In the noise-only case, we have $C_i(k) = |W_i(k)|^2$. Based on the results from Sec. 3.2, the values of $C_i(k)$ for $k = 0, 1, \ldots, N-1$ possess the χ^2 probability distribution with two degrees of freedom and have the cumulative distribution function

$$P_w(x;i) = \left(1 - \exp\left(-\frac{x}{2^{i-1}N_0}\right)\right) u(x), \qquad (5)$$

where $u(x)$ is the unit step function.

Under hypothesis H_s, considering the contributions to $C_i(k)$ from both $S_i(k)$ and $W_i(k)$, there exist $N/2^i$ elements with signal plus noise and $N - (N/2^i)$ elements with noise only. The elements containing only noise are distributed according to Eq. (5). The elements with signal and noise are derived from squaring the magnitude of a deterministic component with magnitude $2^i\sqrt{E}$ summed with a complex-valued noise component with Gaussian-distributed real and imaginary parts of zero-mean and variance $2^{i-2}N_0$. This results in independent non-central χ^2 densities with two degrees of freedom for these elements of $C_i(k)$ and each of them has the cumulative distribution function

$$P_s(x;i) = 1 - Q\left(\sqrt{2^{i+1}\mathrm{SNR}_{\mathrm{in}}}, \sqrt{\frac{x}{2^{i-2}N_0}}\right), \qquad (6)$$

where $I_0(\cdot)$ is the zeroth-order modified Bessel function of the first kind, $Q(\cdot,\cdot)$ is Marcum's Q function [3], and $\mathrm{SNR}_{\mathrm{in}} = 2E/N_0$ is the SNR at the input to the detector.

4. SIGNAL DETECTION FROM FFT STAGES

By forming the distribution of the maximum energy value found across all elements of $C_i(k)$ under each input hypothesis [4], we can determine the threshold values required to obtain a given probability of false alarm, the resulting probability of detection, and the receiver operating characteristic achieved by applying the ML detection strategy after any FFT stage.

Application of the ML detector according to the Neyman-Pearson criterion requires that we obtain the threshold value which gives the desired probability of false alarm (P_{FA}). Since the noise distribution depends on i, so must the threshold, which we denote by η_i. The probability of producing a false alarm, given that a threshold of η_i is applied, is

$$P_{FA} = 1 - \left(1 - \exp\left(-\frac{\eta_i}{2^{i-1}N_0}\right)\right)^N. \qquad (7)$$

It follows that a given value of P_{FA} is obtained when

$$\eta_i = -2^{i-1}N_0 \ln\left[1 - (1 - P_{FA})^{1/N}\right]. \qquad (8)$$

The probability of detection $P_D(i)$ obtained when FFT processing is terminated after i stages can be derived from the distribution of $C_i(k)$ under hypothesis H_s:

$$P_D(i) = 1 - \left(1 - \exp\left(-\frac{\eta_i}{2^{i-1}N_0}\right)\right)^{N-(N/2^i)} \times \left(1 - Q\left(\sqrt{2^{i+1}\mathrm{SNR}_{\mathrm{in}}}, \sqrt{\frac{\eta_i}{2^{i-2}N_0}}\right)\right)^{N/2^i} \qquad (9)$$

The receiver operating characteristic is found by substituting Eq. (8) into Eq. (9):

$$P_D(i) = 1 - (1 - P_{FA})^{1-2^{-i}} \times \left(1 - Q\left(\sqrt{2^{i+1}\mathrm{SNR}_{\mathrm{in}}}, \sqrt{-2\ln\left[1 - (1 - P_{FA})^{1/N}\right]}\right)\right)^{\frac{N}{2^i}} \qquad (10)$$

This performance analysis enables us to verify that the detector performance improves monotonically across stages. By considering the first derivative of the ROC, taken with respect to i, and making term-wise comparisons on the infinite series expansion [6] of Marcum's Q function, we obtain that for any input SNR and false alarm probability, the probability of detection increases monotonically with i. Since $i = v$ corresponds to performing all FFT stages, the performance obviously converges to that of the exact ML detector.

For any fixed input SNR and P_{FA}, the improvement in $P_D(i)$ between successive FFT stages is non-uniform across any given FFT algorithm. This is exemplified by Fig. 3 which shows a typical characteristic for $P_D(i)$ when $P_{FA} \lll 1$. In general, the change in $P_D(i)$ between successive stages depends upon two counteracting effects: the doubling of SNR at the output of the FFT, which increases the probability of detection, and the halving of the number of channels containing signal energy, which decreases it. Since $P_D(i)$ increases monotonically, it follows that the shape of the Q function, which increases monotonically with its first parameter, is the primary influence on the change in detection probability at each stage. In comparison, the reduction in the number of channels containing signal energy is of secondary importance.

5. DISCUSSION

Using the performance analysis derived in the previous section, we can determine the number of FFT stages that must be completed in order to obtain a desired detection performance. Such information provides a sound basis for a system designer, system operator, or control algorithm to select the earliest stage at which the FFT calculation can be terminated while meeting application specific constraints on the detection and false alarm probabilities.

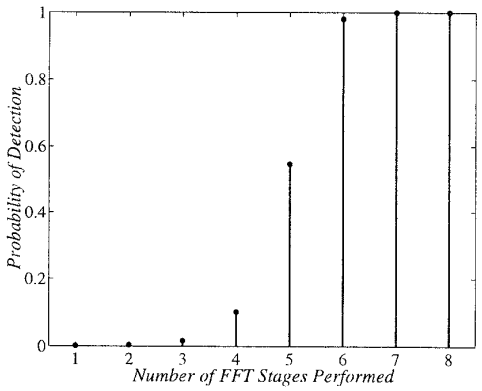

Figure 3. Detection probabilities at successive FFT stages when $\mathrm{SNR_{in}} = -6$ dB, $P_{FA} = 10^{-4}$, and $N = 256$.

Input SNR (dB)	*FFT Stages Needed to Guarantee $P_D \geq 0.7$*
6	0
4	1
2	2
0	3
-2	4
-4	5
-6	6
-8	7
-10	8
-12	-

Table 1. Number of FFT stages that must be performed in order to guarantee $P_D \geq 0.7$ when $N = 256$ and $P_{FA} = 10^{-4}$.

The number of stages needed to obtain a desired detection probability for any specific input SNR and P_{FA} can be obtained by evaluating the ROC in Eq. (10) for increasing values of i. This is illustrated in Table 1, which lists for different values of input SNR the minimum number of FFT stages that must be performed in order to guarantee a detection probability of 0.7.

6. ACKNOWLEDGMENTS

The authors wish to thank Prof. Carl W. Helstrom for kindly providing his computer software for the numerical evaluation of Marcum's Q function.

A APPENDIX

In this appendix, we prove that when a complex sinusoid is analyzed by a radix-2 FFT, each constituent butterfly has either one non-zero output and two non-zero inputs, or else all of its inputs and outputs are zero. We first show that the number of non-zero elements in the output of a butterfly is bounded below by one-half the number in its inputs. This bound applies across all butterflies in an FFT stage. Next, we demonstrate that this bound is met exactly between successive stages of processing for $S_i(k)$. Finally, we show that if there were a single butterfly with two non-zero inputs and two non-zero outputs, this bound would be contradicted by the remaining butterflies of the stage. Therefore, such a butterfly could not exist, which proves our claim.

Bound: The total energy in the butterfly outputs c and d (see Fig. 2) can be easily shown to be related to that in the inputs a and b by

$$|c|^2 + |d|^2 = 2(|a|^2 + |b|^2). \tag{11}$$

Thus, for any butterfly, the total energy at its outputs is twice that at its inputs. This implies that for each butterfly, if one (or both) of its input values is non-zero, then at least one of its output points must be non-zero. Applying this observation across the $N/2$ independent butterflies of a given FFT stage, we obtain that the number of non-zero elements in the input to stage i can be no fewer than one-half the number of non-zero elements in its output. Denoting by ξ_i the number of non-zero elements in $S_i(k)$, this can be written as

$$\xi_i \geq \xi_{i-1}/2, \quad i = 1, 2, \ldots, v. \tag{12}$$

Applying this inequality across successive stages, we obtain

$$\xi_i \geq \xi_j / 2^{i-j}, \quad 0 \leq j < i \leq v. \tag{13}$$

We use $i = 0$ to represent the FFT for which *no* stages have been performed; i.e. $S_0(k)$ represents the input signal.

Equality: For the complex sinusoidal input given in Eq. (2), we have $|S_0(k)|^2 = E$ for all $k = 0, 1, \ldots, N-1$, so $\xi_0 = N$. Considering the output of the final FFT stage (the DFT of $s(n)$), we have $|S_v(l)|^2 = N^2 E$ and $|S_v(k)|^2 = 0$ for $k \neq l$, so $\xi_v = 1$. This implies that the inequality given in Eq. (13) must hold as an equality and it follows that Eq. (12) also holds as an equality at each intermediate stage. Thus, for each $i = 0, 1, \ldots, v$, $S_i(k)$ has exactly $\xi_i = N/2^i$ non-zero elements.

Contradiction: If there were a butterfly in stage i with non-zero energy at both of its input and both of its output elements, then the other $(N/2) - 1$ butterflies must be non-zero at $(N/2^{i-1}) - 2$ of their input elements and also at $(N/2^i) - 2$ of their output elements. Because each of these other butterflies are governed by Eq. (11), it follows that Eq. (12) should hold for these $(N/2) - 1$ butterflies taken together. However, there are $(N/2^{i-1}) - 2$ non-zero values in the inputs to these butterflies and $(N/2^i) - 2$ in their outputs, contradicting Eq. (12).

REFERENCES

[1] J. M. Winograd and S. H. Nawab, "Incremental refinement of DFT and STFT approximations," *IEEE Signal Processing Letters*, vol. 2, pp. 25–28, Feb. 1995.

[2] S. H. Nawab and J. M. Winograd, "Approximate signal processing using incremental refinement and deadline-based algorithms," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing*, vol. 5, (Detroit), pp. 2857–2860, May 1995.

[3] A. D. Whalen, *Detection of Signals in Noise*. New York: Academic Press, 1971.

[4] R. J. Kenefic, "Generalized likelihood ratio detector performance for a tone with unknown parameters in Gaussian white noise," *IEEE Trans. Signal Processing*, vol. 39, pp. 978–980, Apr. 1991.

[5] A. V. Oppenheim and R. W. Schafer, *Discrete–Time Signal Processing*. Englewood Cliffs, NJ: Prentice Hall, 1989.

[6] S. O. Rice, "Mathematical analysis of random noise," *Bell Sys. Tech. J.*, vol. 23, pp. 282–332, July 1944. Also vol. 24, pp. 46–156, Jan. 1945.