

Incremental Refinement of DFT and STFT Approximations

Joseph M. Winograd and S. Hamid Nawab

Abstract— We present a class of multistage algorithms for carrying out *incremental refinement* of DFT and STFT approximations. Each stage is designed to improve the previous stage's approximation in terms of frequency coverage, frequency resolution, and SNR. These algorithms rely almost exclusively on vector summation operations, and they can be designed to exhibit a variety of tradeoffs between improvement in approximation quality and computational cost per stage. The performance of incremental STFT refinement on real data serves to illustrate the relevance of such algorithms to application systems with dynamic real-time constraints.

I. INTRODUCTION

ALGORITHMS that can quickly produce approximate results and then refine those results in an incremental manner are useful for achieving graceful degradation of performance as available resources diminish in systems with dynamic real-time constraints [1], [2]. Results have previously been reported [3]–[5] on the design and evaluation of algorithms for the approximation of the DFT and the short-time Fourier transform (STFT) in order to meet predetermined bounds on computational cost. In this letter, we extend that work to include algorithms that perform incremental refinement of solutions.

II. SUCCESSIVE APPROXIMATION OF THE DFT

Assume the N -point signal frame $x(n)$ under analysis to be real valued and represented in B -bit two's complement binary fraction format. Denoting bit b of the n th sample as $x_b(n)$, this signal may be expressed as a linear combination of the bit vectors $x_0(n), x_1(n), x_2(n), \dots, x_{B-1}(n)$ by

$$x(n) = \sum_{b=0}^{B-1} \beta(b)x_b(n) \quad (1)$$

with $\beta(0) = -1$, and $\beta(b) = 2^{-b}$ otherwise.

Using a backward differencing approach [3], [5] for DFT evaluation of quantized frames, we define the result of the i th

Manuscript received September 17, 1994; approved December 13, 1994. This work was supported 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. The associate editor coordinating the review of this letter and approving it for publication was Prof. A. E. Yagle.

The authors are with the Electrical, Computer, and Systems Engineering Department, Boston University, Boston, MA 02215 USA.

IEEE Log Number 9409097.

successive approximation to the DFT of $x(n)$ as

$$\hat{X}_i(k) = \sum_{b=0}^{v_i-1} \sum_{n=0}^{r_i-1} g_b(n)G_{n,b}(k), \quad 1 \leq k \leq c_i \quad (2)$$

where $g_b(n)$ is obtained by applying backward differencing to $x_b(n)$:

$$g_b(n) = \begin{cases} x_b(0) - x_b(N-1), & n = 0 \\ x_b(n) - x_b(n-1), & 1 \leq n \leq N-1 \end{cases} \quad (3)$$

and $G_{n,b}(k) = \beta(b)e^{-j2\pi/Nkn}/(1 - e^{-j2\pi/Nk})$. The indexing bounds c_i, r_i and v_i may be considered control variables for the refinement process. To ensure that solution quality monotonically increases, we require that c_i, r_i , and v_i be nondecreasing with i and that $c_i + r_i + v_i < c_{i+1} + r_{i+1} + v_{i+1}$. The control variable values for a particular sequence of approximations may be summarized with a *control matrix* \mathbf{A} in which the i th row contains the values of the control variables c_i, r_i , and v_i (in that order) for the i th successive approximation.

The quality of $\hat{X}_i(k)$ with respect to frequency coverage is $2\pi c_i/N$ rad/sample. Its frequency resolution is approximately $2\pi/r_i$ rad/sample. The additive noise in $\hat{X}_i(k)$ due to quantization is approximately $6v_i$ dB SNR.

III. INCREMENTAL DFT REFINEMENT ALGORITHMS

The i th stage of an incremental refinement algorithm for evaluating (2) has to perform the following operations:

$$\hat{X}_i(k) = \begin{cases} \hat{X}_{i-1}(k) + C_i(k), & c_{i-1} < k \leq c_i \\ \hat{X}_{i-1}(k) + R_i(k) + V_i(k), & 1 \leq k \leq c_{i-1} \end{cases} \quad (4)$$

where $C_i(k)$ is the *coverage update*, which is defined as

$$C_i(k) = \sum_{b=0}^{v_i} \sum_{n=0}^{r_i} g_b(n)G_{n,b}(k) \quad (5)$$

where $R_i(k)$ is the *resolution update*, which is defined as

$$R_i(k) = \sum_{b=0}^{v_i} \sum_{n=r_{i-1}+1}^{r_i} g_b(n)G_{n,b}(k), \quad (6)$$

and $V_i(k)$ is the *SNR update*, which is defined as

$$V_i(k) = \sum_{b=v_{i-1}+1}^{v_i} \sum_{n=0}^{r_{i-1}} g_b(n)G_{n,b}(k). \quad (7)$$

The initial conditions for these updates are taken to be $c_0 = r_0 = v_0 = 0$ and $\hat{X}_0(k) = 0$ for all k . Each of the stage i

updates is nontrivial only when an improvement is made in its corresponding dimension of quality.¹

The updates may be implemented without multiplies through a vector summation approach reported previously [4], [5], which requires that we precompute and store in memory the vectors $G_{n,b}(k)$. Each summation term in the update equations may be implemented as the addition of a portion of one of these vectors to the corresponding elements of $\hat{X}_{i-1}(k)$. The vectors associated with zero-valued points in $g_b(n)$ may be omitted from the summation process.

In approximations for which frequency coverage is reduced, it is desirable for the range of frequency samples included in $\hat{X}_i(k)$ to contain significant energy. The center frequency of the energy distribution of $x(n)$ may be estimated from the number of nonzero values in $g_0(n)$ using an efficient method described in [4] and [5]. The frequency coverage is then constrained to be symmetric about this center frequency.

IV. COST CHARACTERIZATION

The cost of performing the vector summations of the update equations is $2c_i r_i v_i \gamma_i$ real additions through stage i , where

$$\gamma_i = \frac{1}{r_i v_i} \sum_{b=0}^{v_i-1} \sum_{n=0}^{r_i-1} |g_b(n)| \quad (8)$$

and γ_i takes on values in the range [0, 1].

If we assume a uniform and independent distribution of quantization levels in $x(n)$, the expected value of γ_i can be shown to be 0.5. A dependence of γ_i on the frequency content of the signal under analysis has been observed, and an efficient frequency reversal technique [4], [5] has been developed that reduces the effective value of γ_i for many signals.

V. INCREMENTAL STFT REFINEMENT

The algorithms for incremental DFT refinement may also be used for STFT frames. The STFT context allows results (or intermediate results) produced for one signal frame to be stored and reused to reduce the computational requirements of neighboring signal frames. For example, overall computation may be approximately halved using the methods described in [3] and [5] when the analysis window is rectangular, there is half-window overlap between consecutive signal frames, and uniform quality is maintained for all STFT frames at each stage.

We have applied an incremental STFT refinement algorithm to a digitized recording of a violin playing a sequence of two notes. Two seconds of this signal, sampled at a rate of 8 kHz and quantized to 256 levels, were analyzed using a four-stage incremental STFT refinement algorithm with $N = 256$, a 128-point rectangular analysis window, a temporal decimation factor of 64, and the control matrix

$$\mathbf{A} = \begin{bmatrix} 65 & 45 & 1 \\ 75 & 85 & 1 \\ 85 & 85 & 2 \\ 95 & 120 & 2 \end{bmatrix} \quad (9)$$

¹A summation for which the lower bound is greater than the upper bound is considered to evaluate to 0.

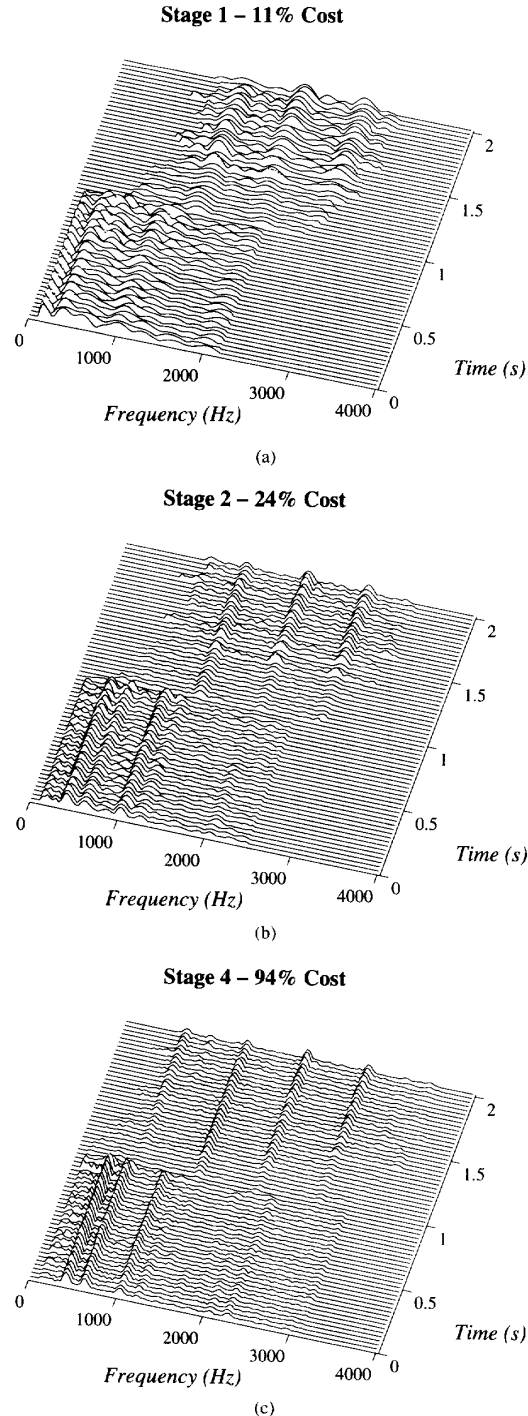


Fig. 1. Incremental refinement of STFT approximations. STFT magnitude is shown with net cost given as a percentage of the number of arithmetic operations required for FFT-based exact STFT analysis: (a) Result of stage 1: initial approximation; (b) result of stage 2: refinement in coverage and resolution; (c) result of stage 4: refinement in coverage, resolution, and SNR.

for all STFT frames. The algorithm incorporates frequency reversal, frequency centering, and interframe result reuse tech-

niques. The successive approximations generated at the completion of stages 1, 2, and 4 are shown in Fig. 1(a)–(c), respectively. The result in Fig. 1(c) has a 12-dB SNR and almost full frequency coverage and resolution. The figures also indicate the net cost up to the i th stage as a percentage of the number of arithmetic operations in an FFT-based evaluation of the exact STFT. We note that although 40% of the arithmetic operations performed in the FFT-based evaluation of the STFT are multiplications, the approximate STFT uses only additions.

VI. CONCLUSION

We have illustrated a framework for generating DFT and STFT algorithms that can quickly produce approximate results and subsequently improve their quality in an incremental

manner. Further details of this framework and its extensions will be addressed in a future paper.

REFERENCES

- [1] V. R. Lesser, J. Pavlin, and E. Durfee, "Approximate processing in real-time problem solving," *AI Mag.*, pp. 49–61, Spring 1988.
- [2] J. W. S. Liu, W. K. Shih, K. J. Lin, R. Bettati, and J. Y. Chung, "Imprecise computations," *Proc. IEEE*, vol. 82, no. 1, pp. 83–94, Jan. 1994.
- [3] S. H. Nawab and E. Dorken, "Efficient STFT approximation using a quantization and differencing method," in *Proc. IEEE ICASSP '93*, Minneapolis, Apr. 1993, pp. 587–590, vol. 3.
- [4] E. Dorken and S. H. Nawab, "Frame-adaptive techniques for quality versus efficiency tradeoffs in STFT analysis," in *Proc. IEEE ICASSP '94*, Adelaide, Apr. 1994, pp. 449–452, vol. 3.
- [5] S. H. Nawab and E. Dorken, "A framework for quality versus efficiency tradeoffs in STFT analysis," to appear in *IEEE Trans. Signal Processing*, Apr. 1995.