FRAME-ADAPTIVE TECHNIQUES FOR QUALITY VERSUS EFFICIENCY TRADEOFFS IN STFT ANALYSIS

Erkan Dorken *S. Hamid Nawab* *

ECS Department, Boston University, Boston, MA

ABSTRACT

In this paper, we present techniques for computing frame-adaptive approximations to the STFT which can satisfy arbitrarily specified bounds on the number of arithmetic operations per frame. The central idea is to represent the samples in each signal frame with a very small number Q of quantization levels, yielding an SNR of $10log(Q^2)$ dB. This leads to the formulation of the frame's DFT computation as a process of summation among pre-stored vectors. Experimental studies are used to propose and verify mathematical models for how the number of vector additions needed in this process depends on the frequency content of a frame. The models are utilized to design frame-adaptive techniques for excluding various subsets of vector elements from the summation process in order to keep the number of additions per frame from exceeding any specified bound, B.

1. INTRODUCTION

Last year, we reported [1] a frame quantization and differencing method for calculating an approximate STFT through a vector summation process and *no multiplications.* With this approach, the use of Q-level quantization theoretically leads to an SNR of $10log(Q^2)$. Recommended values of Q are 3 and 5, which lead to SNRs of 9 dB and 14 dB respectively. In this paper, we present techniques which can be used to ensure that the number of additions in the vector-summation process does not exceed a specified bound, B, for each frame. An example is used to illustrate how important time-frequency features of the signals are preserved even when the bound B is *significantly lower* than the number of additions required by FFT-based approaches. This type of computational efficiency is achieved by sacrificing frequency resolution and/or frequency coverage in a manner which adapts to the frequency content of each frame. The overhead computation in these techniques is also shown to be comparable to the overhead computation in FFT-based approaches.

*This work was sponsored in part by the Rome Laboratories of the Air Force Systems Command under contract number F30602-91-C-0038 and in part by the Office of Naval Research under contract number N00014-93-1-0686.

Let us consider the discrete STFT of a signal $x(n)$:

$$X_{mL}(k) = \sum_{n=mL-N_w+1}^{mL-N_w+N} f_{mL}(n) e^{-j\frac{2\pi}{N}kn} \quad (1)$$

where $f_{mL}(n) = x(n)w(mL-n)$ and $k = 0, 1, ..., N-1$. Suppose $f_{mL}(n)$ undergoes a Q-level quantization. The quantized frames, $f^Q_{mL}(n)$, can be used to obtain an "approximate" STFT:

$$X^Q_{mL}(k) = \sum_{n=mL-N_w+1}^{mL-N_w+N} g^Q_{mL}(n) W_n(k) \; ; \; k = 1, ..., N/2 \quad (2)$$

where

$$W_n(k) = \left(\frac{e^{-j\frac{2\pi}{N}kn}}{1 - e^{-j\frac{2\pi}{N}k}} \right) \quad (3)$$

and

$$g^Q_{mL}(n) = \begin{cases} f^Q_{mL}(n) - f^Q_{mL}(n+N-1) \\ \text{for } n = mL - N_w + 1 \\ f^Q_{mL}(n) - f^Q_{mL}(n-1) \\ \text{for } mL - Nw + 1 < n \le mL - N_w + N \end{cases} \quad (4)$$

The expression in (2) for $X^Q_{mL}(k)$ can be viewed in terms of a *vector-summation* operation performed among column vectors each of which consists of the multiplicative coefficients $W_n(k)$ for a particular n scaled by one of the $(2Q-1)$ quantization levels used to represent $g^Q_{mL}(n)$.

Four basic variations on our vector-summation approach for the evaluation of the approximate STFT are presented in this paper. When the number of additions required by unrestricted vector-summation exceeds B for a particular frame, these techniques sacrifice frequency resolution, frequency coverage or some combination of the two in proportion to the number of additions which have to be eliminated. Furthermore, the reduction of frequency coverage is designed to be sensitive to the frequency content of the corresponding frame.

The loss in frequency resolution or frequency coverage results from the *exclusion* of a subset of vector elements from the vector-summation process. Let us assume that for the mth frame there are N_v vectors remaining after the exclusion of zero-valued vectors. Since each vector is $N/2$ elements long and in general each element is complex, the unrestricted vector-summation process for the mth frame requires $N \times N_v$ real additions. If $B \geq N \times N_v$, there is no need for further computational efficiency. The major differences among the four techniques arise when $B < N \times N_v$.

2. TIME-NARROWING TECHNIQUE

The simplest of the four approaches, frame-adaptive time narrowing involves the use of a counter to ensure that the number of vectors summed for the mth frame does not exceed $\lfloor B/N \rfloor$. For each $g^Q_{mL}(n)$, as defined in (4), a counter n_v is initialized to zero and the value of the argument n is incremented starting from a value of $mL - N_w + 1$. Whenever the value of $g^Q_{mL}(n)$ is non-zero, the corresponding vector addition is performed and the counter n_v is incremented and compared to $\lfloor B/N \rfloor$. The process for that frame stops when n_v matches $\lfloor B/N \rfloor$ or the full length of $g^Q_{mL}(n)$ has been covered by the incrementing variable n. Since the number of non-zero elements in $g^Q_{mL}(n)$ is sensitive to the frequency content of the corresponding signal frame, the amount of time narrowing introduced by this approach is similarly frame dependent. Generally speaking, frames with spectral energy concentrated in lower frequency regions experience a lesser degree of loss in frequency resolution.

3. FREQUENCY-NARROWING TECHNIQUE

The idea behind this approach is to compute $X^Q_{mL}(k)$ for a restricted set of values for k in order to keep the number of additions below B. It is desirable for the restricted range to include frequencies at which the corresponding signal frame has its most significant spectral energy. Toward this end, experiments were conducted to find the dependence of the number of non-zero samples in each frame to the frequency content of the frame [2] [3]. The results are presented in Figure 1.

Consider the data in Figure 1 for the case of 3-level quantization. For each frame length, N_w, this data may be modeled by a monotonically increasing function:

$$N_v = N_w(1 - e^{-6f}); \quad 0 \leq f \leq 0.5 \tag{5}$$

where f denotes normalized discrete-time frequency. The data of Figure 1 along with the corresponding models are shown in Figure 2. Obviously, the actual data is not monotonically increasing in the higher frequency regions. However, we model it to be a monotonically increasing function because such a function may be inverted to obtain an expression for frequency in terms of the number of non-zero samples detected in $g^Q_{mL}(n)$:

$$f = -(1/6)\ log\left(1 - \frac{N_v}{N_w}\right) \tag{6}$$

Our approach to frequency narrowing for the mth signal frame starts by counting the total number N_v of non-zero samples in $g^Q_{mL}(n)$. Using (6), the value f_o for f is obtained and converted to the integer k_o, calculated as $\lfloor f_o N \rfloor$. Letting R be $\lfloor B/2N_v \rfloor$, the values of k_1 and k_2 are then selected such that the frequencies at which the N-point DFT is calculated are centered as much as possible around the frequency f_o:

$$\begin{aligned} k_1 &= k_o - \lfloor \tfrac{R}{2} \rfloor \\ k_2 &= k_o + \lfloor \tfrac{R}{2} \rfloor - 1 \end{aligned} \tag{7}$$

When $k_o - \lfloor R/2 \rfloor < 1$, we set k_1 to 1, and k_2 to R. Similarly, when $k_o + \lfloor R/2 \rfloor > N/2$, we set k_2 to $N/2$ and k_1 to $N/2 - R + 1$. The flowgraph for an algorithm based on this adaptive frequency-narrowing approach is shown in Figure 3. The centering of the restricted frequency region about f_o is based on the assumption that the underlying frame has a significant amount of spectral energy in the vicinity of the frequency which corresponds to N_v through the relationship in (6).

As alluded to earlier, the actual relationship between f and N_v (as represented in Figure 2) is not monotonic at higher frequencies. In particular, this means that each value of N_v beyond a particular threshold corresponds to more than a single value of f. The modeling of the data in that region by a monotonic curve can therefore lead to substantial inaccuracies in determining the dominant frequency in a frame on the basis of its N_v measurement. For values of B which force the frequency coverage to be severely limited, such inaccuracies in the estimation of the center frequency can cause frequency regions with significant spectral energy to be missed. To circumvent this problem, we have formulated one approach which combines time and frequency narrowing and another one which carries out a type of frequency reversal when high-frequency energy is determined to dominate a frame.

4. HYBRID NARROWING TECHNIQUE

This approach is designed to guarantee minimum frequency coverage in order to mitigate the effects of unreliable "center frequency" estimates obtained through (6). This is achieved by first using frequency narrowing with the constraint that the width of the narrowed frequency range, $R = \lfloor B/2N_v \rfloor$, is not smaller than a pre-specified number R_{min}. When $R < R_{min}$, frequency coverage is kept at R_{min}, and time narrowing is used for reducing the number of vectors in the summation process to $\lfloor B/2R_{min} \rfloor$. By setting the value of R_{min} in the hybrid narrowing approach to specific values, it is possible to obtain STFT approximations with a "balanced" sacrifice of time-resolution and frequency-coverage.

5. FREQUENCY REVERSAL TECHNIQUE

In this technique we apply backward differencing on both the quantized frame, $f^Q_{mL}(n)$, and its "frequency reversed" version, $(-1)^n f^Q_{mL}(n)$. The vector summation is applied to whichever result has the smaller number of nonzero elements. This procedure has the advantage that if $f^Q_{mL}(n)$ is a high-frequency frame, its frequency reversed version is a low-frequency frame. Consequently, applying backward differencing to $(-1)^n f^Q_{mL}(n)$ yields a smaller number of nonzero samples and therefore fewer additions are needed in the vector summation process.

6. EXAMPLE

In Figure 4, we illustrate quality versus efficiency tradeoffs which can be obtained with our techniques in the case of a musical signal corresponding to two consecutive notes played on a violin. Figure 4(a) corresponds to the exact FFT-based STFT computed with 8960 real arithmetic operations (40% are multiplications) per frame. Figures 4(b), 4(c), and 4(d) correspond to the Approx. STFT performance with 4480, 2240, and 1120 real addition operations respectively. The number of real multiplications is less than 3% of the number of real additions in the case of the Approx. STFT.

REFERENCES

[1] S.H. Nawab, E. Dorken, "Efficient STFT Approximation Using a Quantization and Differencing Method", *Proceed. IEEE 1993 International Conference on Acoustic, Speech and Signal Processing*, vol. 3, pp. 587-90.

[2] Erkan Dorken, "Approximate Processing and Knowledge Based Reprocessing of Non Stationary Signals", *PhD. Thesis Dissertation*, Boston University, Sept. 93.

[3] E. Dorken, S.H. Nawab, "A framework for Quality Versus Efficiency Tradeoffs in STFT Analysis", *Submitted to IEEE Trans. Signal Processing.*

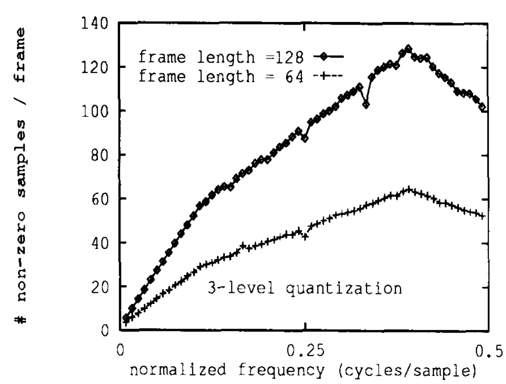

Figure 1. The number of non-zero samples N_v obtained after 3-level quantization and backward differencing of sinusoidal frames with different frequencies. The curve joining the points marked with ◇'s represents the results for frames of length 128 samples, whereas the curve joining the points marked with +'s depicts the results for 64-sample frames.

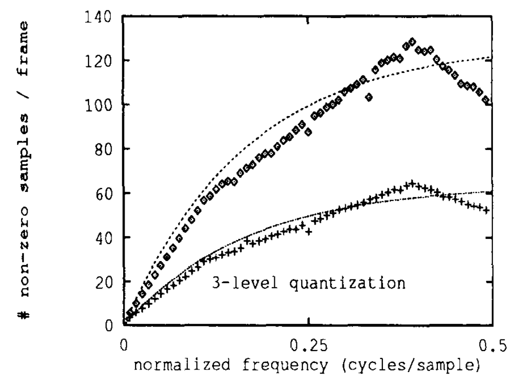

Figure 2. The data shown in Figure 1 plotted along with the corresponding models expressed in (5). The curves with dashed lines depict the models.

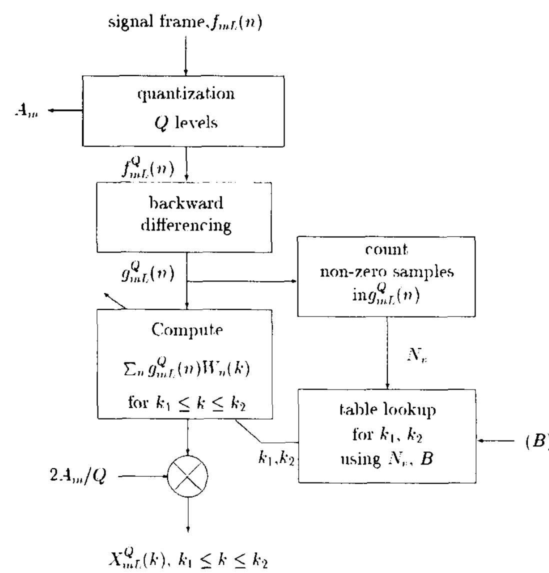

Figure 3. The flowgraph of the algorithm for calculating the approximate STFT based on the frequency narrowing approach.

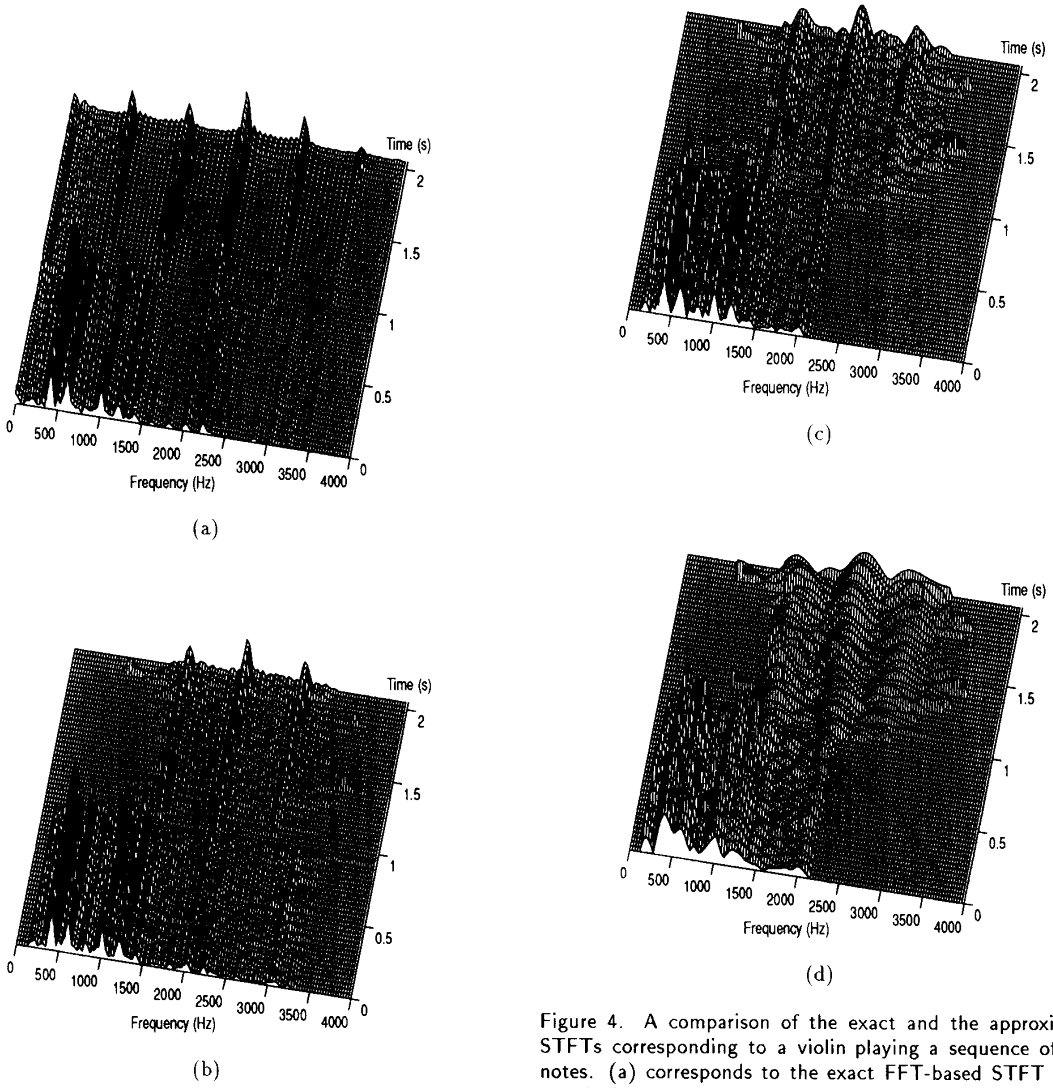

Figure 4. A comparison of the exact and the approximate STFTs corresponding to a violin playing a sequence of two notes. (a) corresponds to the exact FFT-based STFT computed with 8960 real arithmetic operations (40% are multiplications) per frame. (b),(c) and (d) correspond to the Approx. STFT performance with 4480, 2240, and 1120 real addition operations respectively.