ON THE FFT OF 1-BIT DITHER-QUANTIZED DETERMINISTIC SIGNALS

L. Cheded, S. Akhtar
Systems Engineering Department
King Fahd University of Petroleum and Minerals
KFUPMBoox: 116, Dhaharan: 31261, Saudi Arabia
Email: cheded@kfupm.edu.sa

ABSTRACT

Reducing the quantization resolution of the FFT input is bound to adversely affect its computational accuracy. This loss in accuracy will in turn prohibit exploiting the practical advantages that an FFT algorithm, with a coarsely-quantized input, will otherwise enjoy. This paper proposes a new theory that allows the use of coarse quantization with only a negligible effect on the FFT computational accuracy. The proposed theory is tested on deterministic signals that undergo the crudest quantization possible (i.e. 1 bit) and is very well supported by the simulation results. Finally, these results provide a strong encouragement to extend this theory to noisy signals as well as to numerous other important transforms.

1. INTRODUCTION

The Fourier transform is undoubtedly a vitally important tool in a variety of engineering fields such as signal and image processing, filtering, spectral analysis, communication , radar and sonar signal processing, optics, acoustics, seismics, etc. As a computational tool, its digital implementation became hugely successful and widely applied when the direct algorithm, known as the discrete Fourier transform (DFT), underwent a radical structural change with the introduction of the now-famous radix-2 fast Fourier transform (FFT) [1]. According to [2], the FFT algorithm was in fact discovered by others before the authors of [1] did. Nevertheless, several important improvements to the original radix-2 FFT algorithm were made and successfully implemented (see for e.g. [3]). An interesting comparison and tutorial review of the state of that art of FFT algorithms can be found in [4,5].

It is interesting to note here that both the original FFT algorithm and all its variants aimed to improve the computational speed of the DFT through structural changes by exploiting symmetry properties and using the “divide to conquer” principle. All these fast algorithms operated on signals that enjoy sufficiently high quantization resolution ($\geq$ 8 bits). It therefore seems logical to expect the speed, cost and structural complexity of a digital implementation of any of these fast algorithms to be further improved upon if the signal quantization resolution were to be reduced since this would then lead to shorter wordlengths and all their associated benefits. However, as the quantization coarsemess increases, so does the inevitable loss in accuracy to the point where such loss may simply be prohibitive. Clearly, any scheme that aims to improve the FFT computational speed by allowing the use of coarse quantization schemes (i.e. simpler signal coding schemes), without impairing the FFT accuracy, would be a welcome addition to the vast FFT computational arsenal.

The main aim of this paper is to introduce such a scheme and show that the FFT computational accuracy is only negligibly impaired even when the
coarsest (i.e. 1-bit), and hence simplest, fastest and least costly quantization scheme is used. This simpler and attractive computational scheme owes its success to the use of our theory of exact moments recovery (EMR) based on a non-subtractively dithered quantization (NSDQ) scheme [6]. This theory was also successfully applied to other areas such as higher-order statistics [7] and frequency response estimation [8]. The 1-bit NSDQ-quantized single channel estimation scheme considered in this paper is block-diagrammatically depicted in Figure 1 where the block NSDQ quantizer is equivalent to a classical quantizer with a dithered deterministic input. Note here that quantization is applied to the input signal and not to the Fourier kernel (i.e. the complex exponential factor e^{-jnw}). As such, this scheme can also be called a modified-relay FFT (MR-FFT) estimator.

2. A KEY NEW THEOREM ON THE NSDQ QUANTIZATION-BASED SINGLE CHANNEL ESTIMATION OF THE DFT/FFT

This section presents some basic definitions and a characterization related to the NSDQ quantization [6], gives the statement of the new frequency-domain first-order theorem (Theorem 1) and then shows how it can be applied to exactly recover the DFT/FFT of a deterministic signal from the DFT/FFT of only the 1-bit (i.e. binary) version of this signal.

2.1 Definition of the NSDQ quantization scheme

Given an input x and a (user-defined) dither signal D that is statistically independent of x, then a non-subtractively dithered quantization (NSDQ) of x is equivalent to the classical quantization (Q) of the dithered signal $y = x + D$, i.e.

$$x \rightarrow x_{NSDQ} = NSDQ(x) = Q_a(y) = y_Q \qquad (1)$$

Here, Q_a represents the entire class of uniform classical quantizers parametrized by the step (q) and the shift factor given by $a \in \left[-\frac{1}{2}, \frac{1}{2}\right)$, i.e.:

$$y_Q = \left(a + n + \frac{1}{2}\right)q \qquad if \quad y \in [(a+n)q, (a+n+1)q\,] \qquad (2)$$

2.2 Definition of the p^th order class of linearizing dither signals $\mathcal{D}_p$

Given an ergodic and stationary dither signal D and its characteristic function $W_D(u)$, then:

$$D \in \mathrm{D_P} \Leftrightarrow W_D^{(r)}\left(\frac{2np}{q}\right) = 0, \forall\, r \in [0, p-1]\, and\ n \neq 0 \qquad (3)$$

According to the closure property of $\mathcal{D}_p$ [4], we can say that if $D \in \mathcal{D}_p$ and for any signal x that is statistically independent of D, then the dithered signal $y = (x + D) \in \mathcal{D}_p$.

2.3 A key forst-order Theorem and its application to the exact recovery of the FFT based on the MR-FFT estimator

First let us recall fom [6] that the NSDQ quantizer is characterized by its p^th-order ($p \geq 1$) moment-sense input/output function (MSIOF), $h_p(x)$, given by:

$$h_p(x(n)) = \sum_{k=0}^{p} c_k\, x^k(n) \quad where$$

$$c_k = \sum_{t=0}^{p-k} \left\{ \begin{array}{l} \frac{p!}{(p-k-t+1)!\ k!\ t!}\left(\frac{q}{2}\right)^{p-k-t} . \\ E\left[R^t\right][p \oplus k \oplus t \oplus 1] \end{array} \right\} \qquad (4)$$

with $\oplus$ denoting modulo-2 operation.

In this paper, we only consider the first-order case, when $p = 1$. A higher-order case will be reported later.

2.3.1 <u>Theorem 1 (First-order case):</u>

*Given a 1-D NSDQ quantizer whose MSIOF, input, output and dither signals are given respectively by $h_p(x), x, x_{NSDQ}$, and D where D is both zero-mean and statistically independent of x, and given that $X(w_i) = DFT[x(n)]$ and $X_{NSDQ}(w_i) = DFT[x_{NSDQ}(n)]$ $\forall i, n \in [1, N]$, then: **(a)** $X_{NSDQ_p}(w_i)$ and the first-order frequency-domain polynomial mapping, defined by $H_1(w_i) = DFT[h_1(x(n))]$, are moment-sense equivalent and hence **(b)** $X(w_i)$ can be exactly recovered from the expected value of $X_{NSDQ}(w_i)$, i.e.:*

$$E\left[X_{NSDQ_p}(w_i)\right] = E[H_1(w_i)] = X(w_i) \qquad (5)$$

Note here that since $h_1(x(n))$ gives an average I/O function of the NSDQ quantizer, its DFT, $H_1(w_i)$, gives the equivalent average frequency-domain characterization of the NSDQ quantizer.

2.3.2 Application of Theorem 1 to the MR-FFT estimation

Part **(a)** of Theorem 1 can be proved using the EMR theory of [6]. For part **(b),** it can be easily shown that for p=1, (4) reduces to the perfectly linear 1^st-order (i.e. average) MSIOF of the NSDQ quantizer, i.e. $h_1(x(n)) = x(n)$. Hence, $H_1(w_i) = DFT[(x(n)]$, which leads directly to . the desired equality of (5), i.e. the FFT of the unquantized signal $x(n)$ can be exactly recovered from the FFT of its binary version.

3. SIMULATION RESULTS AND CONCLUSION

The simulation setup is shown in Figure 1 where a uniformly-distributed dither signal was used. Three cases were simulated as explained below. All 3 cases are based on a 1-bit NSDQ quantizer.

First, a sinusoidal signal of amplitude *A=10* and frequency *f=500 Hz* was used as the input signal. Fig.2 shows the amplitude spectra for the original (i.e. unquantized) signal and a non-dithered 1-bit quantized version of it. It is very clear from the spectrum of the 1-bit quantized signal that it is very different from the spectrum of the unquantized signal in that there are very noticeable signal harmonics and also, the spectrum amplitude (i.e. 3) at the test frequency is much lower than the desired one (i.e. 5), hence leading to a large (40%) estimation error. Next, the same input signal was 1-bit dither-quantized, i.e. a zero-mean uniformly-distributed dither signal was first added to it and the sum signal 1-bit quantized. The results are shown in Figure 3 which compares the amplitude spectrum of the unquantized signal with that of the 1-bit dither-quantized version of the same signal. It is very clear from Figure 3 that the 1-bit NSDQ quantization scheme has not only recovered, with a very small error (about 3%), the desired amplitude spectrum value of ($A/2$) , i.e. 5, but also replaced the distinctly structured error, due to the contribution of the signal harmonics (see Fig. 2), with a totally unstructured (i.e. random-like) error that is spread across the entire frequency range and whose peak magnitude can be decreased by increasing the estimation time.

In the second simulation, we selected a summation of two sinusoidal signals of amplitude $A_1=10$ and $A_2=8$ and frequencies $f_1=400\ Hz$ and $f_2=600\ Hz$ respectively. Although not shown here, the plot comparing both the spectra of the unquantized signal with that of the undithered 1-bit quantized signal display the same kind of discrepancies and error magnitude as those described above in the single sine case. However, The plot of Figure 4 which compares the spectra of the unquantized signal and that of its 1-bit dither-quantized version shows an excellent spectrum recovery.

In our third and final simulation, we selected the output signal of a second-order dynamical system representing the rotational motion of a satellite controlled by thrusters. This system is described by the following difference equation:

$$y(n) - y(n-1) + 0.5\,y(n-2) = 0.5\,x(n) + 0.5\,x(n-1) \qquad (6)$$

A sinusoidal input (cos(2ðn/10) was applied to this system. The amplitude spectra of the original and 1-bit dither-quantized system outputs are shown in Fig. 5, which again shows an excellent spectrum recovery performance. It is to be pointed out that in all our simulation, the estimation time for the dither-based estimators was taken longer than that for the undithered ones. But that is only a small price to pay for the attractive practical benefits gained in return, such as a simplicity, low-cost, high computational speed and excellent estimation accuracy of the resulting estimator.
In conclusion, it can be said that the excellent results obtained here are very encouraging and provide the driving force to apply this technique to noisy signals and extend it to the estimation of a variety of other transforms.

4. REFERENCES

[1] J.W. Cooley and J.W. Tuckey, " An algorithm for the machine computation of complex Fourier series", Math. Comp. Vol. 19, pp 297-301, April 1965

[2] G.Heath**,** " FFT history" at http://www.lns.conell.edu/spr/2002-08/ msg0043452. html

[3] C.S. Burrus and T.W. Parks, " DFT/FFT and convolution algorithms", Wiley 1984.
See also C.S. Burrus: " Notes on the FFT " at : **http://www-dsp.rice.edu/res/fft/fftnote.asc**

[4] A. Ganapathiraju et al , " A comparative analysis of FFT algorithms " at : **http://www.isip.msstate.edu/publications**

[5] P. Duhamel and M. Vetterli, "Fast Fourier transforms: A tutorial review and a state of the art ", Signal Processing, Vol. 19, pp 259-299, 1990

[6] L. Cheded, **"**"Exact recovery of higher-order moments", IEEE Trans. Info. Theory, Bol. 44, No.2, March 1998, pp 851- 858.

[7] L. Cheded**, "**Exact recovery of Cumulants", Int. Conf. On Signal Proc., ICSP'98, Beijing, China.

[8] L. Cheded and S. Akhtar, " A new and fast frequency response estimation technique for noisy systems", 35th Asilomar conference on computers, circuits and Signals, Asilomar, Orange Grove, CA, USA, 2001

Acknowledgement:

The authors acknowledge KFUPM for its support.

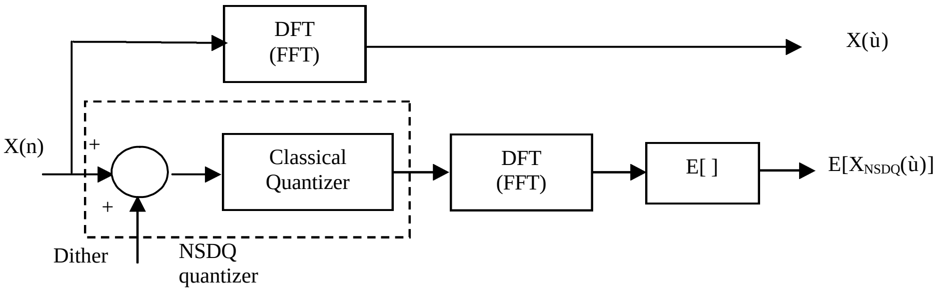

Fig.1: NSDQ quantization-based FFT Estimation FFT (MR_FFT)

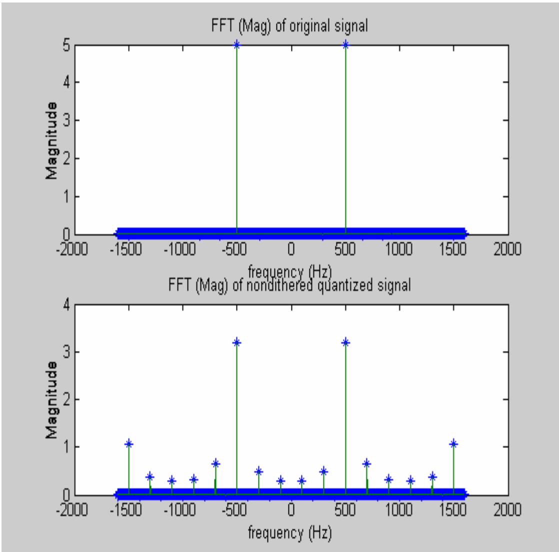

Fig.2: FFT (Mag) spectra of an unquantized and a 1-bit quantized single sine signal

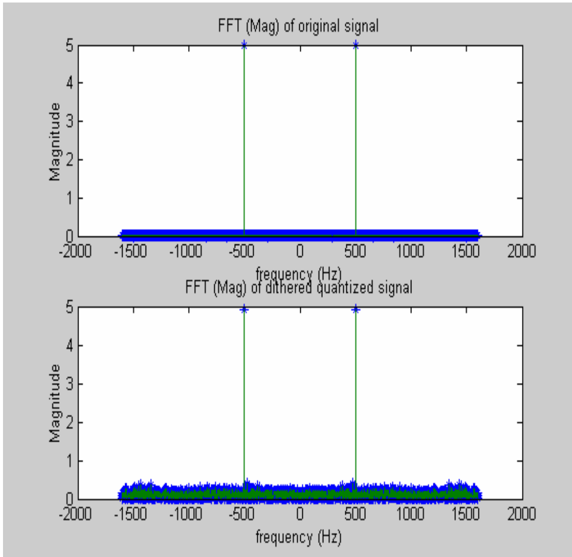

Fig.3: FFT (Mag) spectra of an unquantized and 1-bit dither-quantized single sine signal

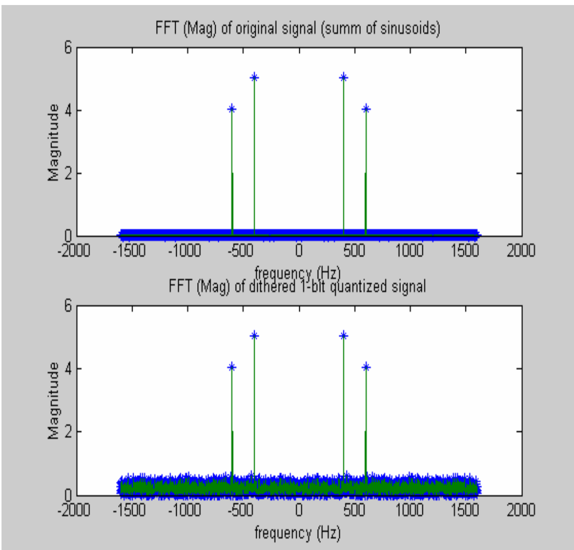

Fig.4: FFT (Mag) spectra of an unquantized and a 1-bit dither-quantized two-sine signal

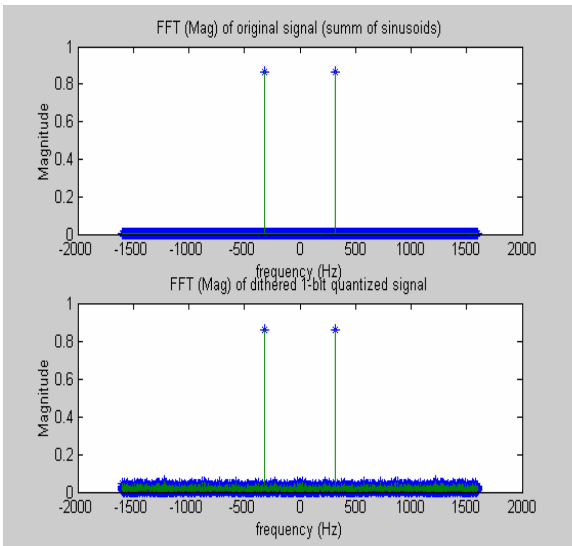

Fig.5: FFT (Mag) spectra of an unquantized and 1-bit dither-quantized of the satellite system output signal