Quantization Noise Conditioning Techniques for Digital Delta-Sigma Modulators

Sudhakar Pamarti

***Abstract*— This paper presents an overview of outstanding theoretical problems in delta-sigma modulator based electronic digital-to-analog circuits and outlines quantization noise conditioning techniques that are being employed to address these problems. Both the problems and the conditioning techniques are described in the context of a special class of electronic circuits called frequency synthesizers.**

I. INTRODUCTION

Digital-to-analog conversion circuits (DACs) are central components in integrated hardware for several wireless communication and consumer electronics applications. Their primary goal is to generate a continuous time signal, $y(t)$, according to a given sequence of numbers, $y[n]$:

$$y(t) = \sum_{n=0}^{\infty} y[n]p(t - nTs), \tag{1}$$

where T_s is a positive real number and $p(t)$ is an arbitrary signal with finite support e.g., a unit pulse of duration, T_s.

Since most DAC implementations employ a binary representation of the sequence, $y[n]$, DAC performance is often characterized in terms of number of effective bits used to represent $y[n]$ e.g., 2-bit DAC, 6-bit DAC, etc. In general, a p-bit DAC has higher performance than a q-bit DAC if $p > q$ all else being comparable. Note that DAC performance is characterized by several other metrics such as bandwidth, speed, and power consumed. However, this paper focuses on effective number of bits for the sake of simplicity. Transistor non-linearity, noise, and inevitable integrated circuit (IC) fabrication errors limit the effective number of bits of integrated DACs. However, the progressive scaling of IC fabrication technologies and Moore's law have made high speed DACs i.e. those with small values of T_s practical.

Delta-sigma modulator based DACs enable high performance by trading off DAC speed for higher effective number of bits. The block diagram of a generic delta-sigma modulator based DAC is shown in Fig. 1. A high resolution sequence, $x[n]$, band-limited to $B \subset [-\pi, \pi]$ and representing the desired analog signal that needs to be generated, is coarsely quantized by a delta-sigma modulator. Samples of the resultant sequence, $y[n]$, have fewer bits than those of $x[n]$ and require only a DAC with fewer number of effective bits. Of course, quantization induces error but the quantization error is suppressed within the signal band, B. Outside the signal band, B, the quantization error is suppressed by filtering subsequent to the DAC. Consequently, a delta-sigma based DAC can achieve very large effective number of bits even while the constituent DAC has only a few effective number of bits.

Fig. 1. Block diagram of a generic delta-sigma modulator based digital-to-analog converter

The aforementioned benefit has led to wide-spread use of delta-sigma modulator based DACs in several applications. However, their performance is plagued by several problems. First, in many delta-sigma modulators, the quantization error suppression within the signal band, B, depends strongly on the input sequence, $x[n]$: on its type e.g., a constant input or a sinusoidal input or a white input, on its variance e.g., large or small etc. This dependence makes quantitative predictions about the delta-sigma modulator based DAC difficult and in many cases unreliable. This is particularly problematic in the design of practical integrated DACs. Second, the amount of quantization noise suppression depends strongly on the number of bits needed to represent the delta-sigma modulator output bits of $y[n]$. In fact, in delta-sigma modulators that quantize to only a 1-bit sequence e.g. $y[n] = 0$ or 1, also known as 1-bit delta-sigma modulators, the quantization error suppression can be very poor for a reasonably large class of sequences, $x[n]$. Third, delta-sigma modulators are known to exhibit limit cycle behavior. Limit cycles result in strong spikes in the power spectral density of the delta-sigma modulator based DAC's output, $y(t)$: such spikes are extremely undesirable in most applications. Fourth, transistor non-linearity and inevitable IC fabrication errors also cause spikes in the power spectral density of DAC output and undo much of the quantization noise suppression achieved by the delta-sigma modulator. Both effects severely limit the performance of the overall delta-sigma modulator based DAC.

Most of the aforementioned problems can be addressed by appropriately modifying the statistical properties of the delta-sigma modulator's quantization noise. Such "quantization noise conditioning" can eliminate limit cycles, improve quantization noise suppression within the signal bandwidth, B, and more importantly, make the digital-to-analog conversion process insensitive to both transistor non-linearity and IC fabrication errorss.

This paper describes the aforementioned problems in

This work was not supported by any organization

S. Pamarti is with the Electrical Engineering Department, University of California, Los Angeles, CA 90095-1594, U.S.A. spamarti@ee.ucla.edu

ISBN 978-963-311-370-7

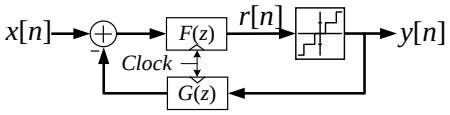

Fig. 2. Block diagram of a generic digital delta-sigma modulator

delta-sigma modulator based DAC design and outlines recent quantization noise conditioning techniques in detail. The problems and the conditioning techniques are described in the context of a special class of DACs called a frequency synthesizer which generates a continuous time periodic signal with instantaneous frequency changes according to the given sequence, $x[n]$. Sections II and III present an overview of the delta-sigma modulator and the frequency synthesizer respectively; Section IV describes two main problems that limit the performance of the frequency synthesizer and describes recent quantization noise conditioning techniques to address them; finally, Section VI concludes the paper.

II. OVERVIEW OF DIGITAL DELTA-SIGMA MODULATION

Fig. 2 shows the block diagram of a generic digital delta-sigma modulator (DDSM). It consists of a quantizer, "forward transmission" and "feedback" filters, represented by their Z-transforms, $F(z)$ and $G(z)$ respectively. The DDSM input is a sequence of rational numbers, $x[n]$, $0 \leq x[n] < 1$, where n is a sample index. It is band-limited to $B = [-\pi/R, \pi/R]$ where R is usually a positive integer called the "oversampling ratio" (OSR) i.e.

$$S_{xx}(e^{j\omega}) = 0, \forall \omega \notin B,$$

where only $\omega \in [-\pi, \pi]$ is considered. In practical implementations, it is common to represent the various sequences in the DDSM in binary form. For example, it is assumed that $x[n]$ is represented by a p bit binary number so, $x[n] \in \{0, 1, ..., 2^p - 1\}/2^p$. The DDSM coarsely quantizes the input sequence, $x[n]$, to a sequence of integer multiples of $1/M$, namely $v[n]$, where $M = 2^Q$, Q is a positive integer, and $Q < p$. The quantizer is assumed to be of the uniform mid-tread type whose operation is defined as:

$$v[n] = \frac{1}{M}\left\lfloor Mr[n] + \frac{1}{2} \right\rfloor,$$

where $\lfloor x \rfloor$ is the largest integer less than or equal to x, and $r[n]$ is the input to the quantizer as shown in the figure. The quantization error, $q[n]$ is defined as,

$$q[n] = v[n] - r[n], \tag{2}$$

and it can be shown that

$$q[n] \in -\frac{2^Q}{2} + \frac{1}{2^p}\{0, 1, ..., K-1\}$$

where $K = 2^p/2^Q$, i.e. $-M/2 \leq q[n] < M/2$. Often times, the range of $v[n]$ is limited intentionally to $\{v_{min}, v_{max}\}$, where $0 < v_{min}$ and $v_{max} < 1$ for practical reasons described later. When this happens, $q[n]$ may not be bounded so, and the situation is referred to as "quantizer overload". It can be shown that

$$v[n] = x[n] * stf[n] + q[n] * ntf[n], \tag{3}$$

where $*$ is the convolution operator, and $stf[n]$ and $ntf[n]$ are sequences whose Z-transforms are the so-called "signal" and "noise transfer functions,"

$$STF(z) = \frac{F(z)}{1 + F(z)G(z)}, \text{ and } NTF(z) = \frac{1}{1 + F(z)G(z)}.$$

The filtered quantization error, $e_{\Delta\Sigma}[n] = q[n] * ntf[n]$ is called "shaped quantization noise" to distinguish it from the quantization error, $q[n]$. The filters, $F(z)$ and $G(z)$, are chosen such that $NTF(z)$ suppresses the quantization error, $q[n]$, within the input sequence's bandwidth, B. The most popular DDSM, called the L^{th} order DDSM, sets

$$F(z) = z^{-L}(1 - z^{-1})^{-L}, G(z) = (1 - z)^L - z^L,$$

resulting in signal and noise transfer functions,

$$STF(z) = z^{-L}, NTF(z) = (1 - z^{-1})^L.$$

Consequently, while the input sequence, $x[n]$, is simply delayed by L samples, the shaped quantization noise, $e_{\Delta\Sigma}[n]$, is attenuated within the signal band, B. The power spectral density of the shaped quantization noise can be shown to be:

$$S_{ee}(e^{j\omega}) = \left|2sin\left(\frac{\omega}{2}\right)\right|^{2L} S_{qq}(e^{j\omega}),$$

if the power spectral density of quantization error, $S_{qq}(e^{j\omega})$, exists. Higher order DDSMs result in superior quantization error suppression within B at the expense of higher amplification outside i.e. in $[-\pi, \pi] - B$.

Note that in many cases, both $q[n]$ and $e_{\Delta\Sigma}[n]$ may not be wide sense stationary and hence, their power spectral densities may not even be defined. This issue is discussed in detail later in Section IV. In practice, it is assumed that $q[n]$ is independent of $q[m]$ for all integers $m \neq n$, independent of $x[m]$ for all integers m, and uniformly and identically distributed for all integers, n, such that,

$$S_{qq}(e^{j\omega}) = \frac{1}{12M^2}, S_{ee}(e^{j\omega}) = \frac{1}{12M^2}\left|2sin\left(\frac{\omega}{2}\right)\right|^{2L}. \tag{4}$$

These properties are henceforth referred to as "desired properties".

III. OVERVIEW OF DELTA-SIGMA MODULATOR BASED FREQUENCY SYNTHESIS

A block diagram of the frequency synthesizer is shown in Fig. 3(a). A low noise reference oscillator generates a periodic binary signal, $v_{ref}(t)$, of frequency f_{ref}. In the following, n is used as a running index that counts the number of cycles of $v_{ref}(t)$. The DDSM quantizes the given sequence, $x[n] = \alpha + m[n]$, where $m[n]$ is a zero-mean sequence and $0 \leq \alpha, x[n] \leq 1.0$, to $v[n]$, just as described in Section II.

The voltage controlled oscillator (VCO) generates M periodic signals, $s_k(t), k = 0, 1, ..., M-1$, all of which have

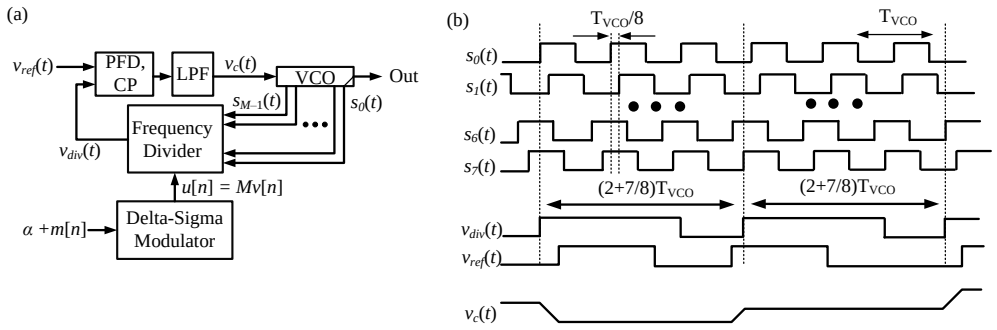

Fig. 3. (a) Block diagram of a digital delta-sigma modulator based frequency synthesizer. (b) Example timing diagram.

the same instantaneous frequency that is proportional to the control voltage, $v_c(t)$, but a different phase offset:

$$s_k(t) = cos(2\pi f_0 t + 2\pi \int_0^t K_v v_c(\tau) d\tau + 2\pi k/M),$$

where f_0 and K_v are constants characteristic of the VCO.

The frequency divider circuit generates a binary output signal, $v_{div}(t)$. Suppose that the $(n-1)^{th}$ positive going transition of $v_{div}(t)$ is aligned with some positive going transition of $s_{k[n-1]}(t)$, where $k[n-1] \in 0,1,...,M-1$ is an arbitrary integer. The frequency divider ensures that there are exactly $(N+v[n])$ VCO cycles (including fractional cycles) between the $(n-1)^{th}$ and n^{th} positive going transitions of $v_{div}(t)$. It does so by counting $(N+\lfloor v[n] \rfloor)$ full cycles of $s_{k[n-1]}(t)$ and then "switching" to $s_{k[n]}(t)$ i.e. aligns its n^{th} positive going output transition with the very next positive going transition of $s_{k[n]}(t)$ where,

$$k[n] = (k[n-1] + Mv[n]) mod\ M. \tag{5}$$

The frequency divider operation is illustrated in Fig. 3(b) for the example case of $N = 2$, $M = 8$, $v[n] = 7/8$, and $v_c(t) = c$, a constant. Effectively, the divider output signal, $v_{div}(t)$, has a frequency that is $(N+v[n]) = (2+7/8)$ times lower than f_0.

The phase frequency detector (PFD) compares the times of occurrence of positive going transitions of $v_{div}(t)$ and $v_{ref}(t)$ and accordingly directs the charge pump (CP) and low-pass loop filter (LPF) circuits to raise (or lower) the control voltage, $v_c(t)$, by an amount that is proportional to the difference between the times of occurrence. If the positive going transition of $v_{div}(t)$ occurs later than that of $v_{ref}(t)$, then the control voltage, $v_c(t)$ is raised, thereby increasing the instantaneous frequency of the VCO outputs, and hastening subsequent positive going transitions of the divider output, $v_{div}(t)$. The resulting negative feedback operation ensures that, in steady state, the outputs of the divider and the reference oscillator are aligned. Consequently, the control voltage, $v_c(t)$, "settles" to such a value that the output, $s_0(t)$, is:

$$s_0(t) = cos\left(2\pi\left[N + y(t)\right] f_{ref} t\right), \tag{6}$$

where

$$y(t) = \alpha + \sum_{n=0}^{\infty} \left[\frac{m[n]}{M} + \frac{e_{\Delta\Sigma}[n]}{M}\right] h(t - nT_{ref}), \tag{7}$$

$T_{ref} = 1/f_{ref}$, and $h(t)$ is a causal function determined by the dynamics of the negative feedback loop. The actual form of $h(t)$ is unimportant for this discussion. It suffices to know that it suppresses the components of $y(t)$ in the frequency band, $[f_{loop}, \infty)$ i.e. it behaves like a low pass filter with a bandwidth of f_{loop}.

The DDSM based frequency synthesizer can be used to generate a periodic signal with a fixed frequency, $(N+\alpha) f_{ref}$, simply by setting $m[n] = 0$. Furthermore, its frequency, and hence its phase, can be modulated accurately according to $m[n] f_{ref}$ as long as $f_{ref}/2R < f_{loop}$. In either role, it is an important component in communication transceiver circuitry. The DDSM enables precise frequency synthesis. If not for the DDSM, the high resolution frequency synthesis would have been impossible: practical circuits force M to be small limiting the granularity of frequency resolution. In fact, in most cases called fractional-N phase locked loops (FNPLLs), M = 1, implying that the VCO generates only a single periodic signal, the divider only counts full cycles, and the frequency synthesis granularity is f_{ref}. So, modulation not possible; furthermore, the granularity is too coarse for most practical channelized wireless systems such as GSM. The DDSM enables these applications by quantizing $\alpha + m[n]$ to $v[n]$.

Just like in the DDSM based DAC showed in Fig. 1, the DDSM quantization noise degrades performance. Specifically, it causes phase noise in the output, $s_0(t)$. It can be shown that the PSD of the phase noise of $s_0(t)$, if its exists, is:

$$S_\phi(f_m) = \frac{\pi^2}{3M^2 f_{ref}} \left| 2sin\left(\frac{\pi f_m}{f_{ref}}\right) \right|^{2(L-1)} |H(j2\pi f|^2 \text{ dBc/Hz},$$

where $H(s)$ is the Laplace transform of $h(t)$, f_m is offset frequency from the target fixed frequency, $(N+\alpha) f_{ref}$, and dBc/Hz refers to power measured in a 1Hz band and normalized to the variance of $s_0(t)$. Since the DDSM suppresses quantization noise within the modulation sequence's bandwidth, B, high precision frequency synthesis and frequency modulation with high signal-to-noise ratio are achieved.

Since the synthesizer circuits (represented by $H(s)$) suppress the quantization noise outside the modulation signal's real bandwidth, $Bf_{ref}/2\pi$, so called "out-of-band" spurious emissions are suppressed as well.

However, practical DDSM based frequency synthesizers suffer from several problems that result in elevated phase noise and strong undesirable spikes in the power spectral density of the phase noise. Most of these problems can be traced to two effects: (a) the quantization error, $q[n]$, not possessing the "desired properties" mentioned at the end of Section II, and (b) the effect of transistor non-linearity and inevitable circuit fabrication errors on the residual shaped quantization noise. These problems and recent DDSM quantization noise conditioning techniques that address them are discussed in the following sections.

IV. PROBLEMS WITH QUANTIZATION NOISE

A. Quantizer Overload

In many DDSM based frequency synthesizers, the goal is to aggressively reject quantization noise within the signal band, B, while restricting the range of $v[n]$ to only a few values. For example, the aforementioned DDSM based fractional-N phase locked loops (FNPLLs) use $M = 1$ and restrict $v[n]$ to 0 or 1. Such restrictions relax the requirements on the frequency divider, the phase frequency detector, and the charge pump circuits and are hence desirable. Since the frequency divider has to accommodate only two possible division modulii, N or $N+1$, divider power consumption can be significantly reduced. A dual modulus frequency divider is also insensitive to systematic delays in the divider circuits that depend on $v[n]$. As described in the next section, if $v[n]$ were a multi-level signal, these systematic delays would act like non-linearity and cause elevated noise and undesirable spikes in the PSD of $s_0(t)$. A small range of allowable $v[n]$ also reduces the maximum instantaneous errors between the times of occurrences of positive going transitions of $v_{ref}(t)$ and $v_{div}(t)$, which can be shown to be,

$$|\tau[n]| \leq T_{vco} \left| \sum_{m=0}^{n-1} (v[m] - \alpha) \right|.$$

Smaller instantaneous errors reduce the circuit noise (thermal noise from transistors and resistors) contributions of the phase frequency detector and the charge pump.

While a small range of $v[n]$ is desirable from a circuit implementation point of view, DDSMs with a limited output range suffer from quantizer overload (see Section II) and several attendant problems. Quantizer overload leads to $q[n]$ not being independent of $x[m]$ for at least some integers m,n or not being independent of $q[m]$ for at least some integers $m \neq n$. In turn, the DDSM may not exhibit the expected quantization noise shaping offered by the chosen filters, $F(z)$ and $G(z)$, i.e. (4) may not be valid anymore.

In general, for DDSMs with quantizer overload, the effective quantization noise suppression achieved within the signal band, B, may vary significantly depending on the magnitude of the input sequence, $x[n]$ e.g., the value of α whem $m[n] = 0$. In other words, the estimate of the PSD of the shaped quantization noise shaping in (4) is inaccurate. A few researchers [1] have employed alternative models of the quantizer that are more sophisticated than the additive model shown in (2) to better predict achievable quantization noise shaping. For example, [1] uses a minimum mean square error model for the quantizer:

$$q[n] = v[n] - (k_x r_x[n] + k_q r_q[n]), \tag{8}$$

where $r_x[n]$ depends on the input sequence, $x[n]$, $r_q[n]$ is noise uncorrelated with $x[n]$, and k_x and k_q are real numbers. The minimum mean square error fit ensures that the residual quantization error, $q[n]$, is uncorrelated with $r[n]$ allowing linear analysis and making predictions about the PSD of the shaped quantization noise. The DDSM is simulated extensively for special cases of constant inputs, $x[n] = \alpha$, or sinusoids, to empirically estimate the gain terms, k_x and k_q, and the variance of the quantization error, $q[n]$, namely σ_q^2, as functions of α. This allows predicting the PSD of the shaped quantization noise as a function of α:

$$S_{ee}(e^{j\omega}) = \frac{\sigma_q^2(\alpha)}{\left|1 + k_q(\alpha)F(e^{j\omega})G(e^{j\omega})\right|^2}. \tag{9}$$

The method does predict signal-to-noise ratio reasonably well for constant and sinusoidal input sequences. However, the situation with more general input sequences in unclear. The extensive DDSM simulations that are required make the method difficult to use in practical integrated circuit design.

Even with better prediction of the effective quantization noise shaping, quantizer overload causes a bigger problem - DDSM instability. For a given range of $v[n]$ and depending on the chosen noise transfer function, $NTF(z)$, there is only a small range of values of $x[n]$ over which the DDSM is stable and the PSD of the shaped quantization noise can be predicted by (9). Beyond this range, the DDSM is termed unstable, doesn't offer any useful quantization noise shaping for use in the frequency synthesizer. Note that the more aggressive the quantization noise shaping e.g. higher L, or the more limited the range of $v[n]$, the lower the range of $x[n]$ for which the DDSM remains stable. For the case of the DDSM based fractional-N PLLs with $v[n]$ limited to 0 or 1, the DDSM stability problem is most aggravated.

B. Limit Cycles

Delta-sigma modulators pose another serious challenge in the form of limit cycles. It is well known that the delta-sigma modulator output and its quantization error, $q[n]$, become periodic for constant rational inputs. Analog delta-sigma modulators, in which the samples of input sequence, $x[n]$, are real numbers and not necessarily rational numbers, are not very susceptible to limit cycles. This is because, in practical implementations of such analog delta-sigma modulators, $x[n]$ represents an analog voltage or current signal and invariably contains some random noise component that may be sufficient to eliminate limit cycles in most cases [2].

In contrast, the input of a DDSM is a digital number that is always rational (see Section II) and has no circuit noise. So,

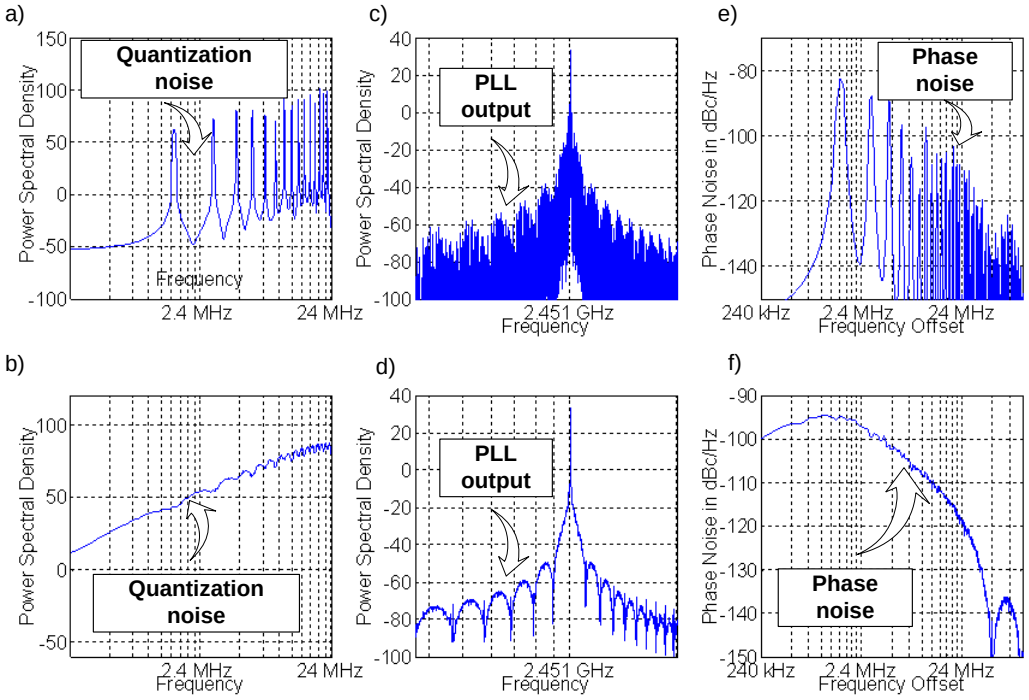

Fig. 4. Simulated power spectral densities of shaped quantization noise (a) without and (b) LSB dithering. Simulated power spectral densities of frequency synthesizer output (c) without and (d) with LSB dithering. (e) Zoomed-in version of (c). (f) Zoomed-in version of (d).

for every constant input, the DDSM output and quantization error are both periodic resulting in spikes in their PSDs. Consequently, limit cycles are problematic in DDSMs: they result in strong spikes in the PSD of the phase noise of $s_0(t)$. Spikes in the PSD are particularly degrading to most wireless and consumer electronics applications such as cell phones, audio/video equipment.

In spite of the paramount importance of eliminating limit cycles in DDSMs, the problem has attracted only scant attention to the task so far. The author's prior work [3] has investigated so-called "LSB dithering" techniques towards this purpose: a random sequence, $d[n] \in 0,1$, called "dither", added to the least significant bit of a binary representation of $x[n]$ e.g.,

$$x_{new}[n] = x[n] + d[n]2^{-p},$$

for a DDSM with a p-bit input, is sufficient to eliminate spikes in the PSD of $q[n]$ and $v[n]$ for a large class of DDSMs including DDSMs of 2^{nd} order or higher. Specifically, consider DDMSs for which $f[n]$ and $g[n]$ are integer valued, the qantizer does not overload, and $K = 2^p/2^Q > 2$. If one of the following conditions is satisfied for every positive integer, m: for each integer pair $(k_1,k_2) \neq 0,0)$ and $0 \leq k_1,k_2 < K$,

1) the sequence, $(k_1 f[n] + k_2 f[n+m])\, mod(K)$ does not converge to 0 as $n \to \infty$,
2) a non-negative integer, $s = s(m) \neq m$ exists such that $(k_1 f[s] + k_2 f[s+m])\, mod(K) = K/2$,
3) a non-negative integer, $g = g(m) < m$ exists such that $(k_2 f[g])\, mod(K) = K/2$

then, the quantization error, $q[n]$, has the following properties:

1) the conditional random variable $q[n]|x[n-m]$ converges in distribution to U,
2) $(q[n], q[n-m])$ converges in distribution to (U,V),

where U and V are independent random variables that are uniformly distributed over the no overload range of $q[n]$. These properties imply not only the elimination of spikes in the PSD of the quantization noise, but also that $q[n]$ asymptotically acquires the aforementioned "desired properties" and approaches a white random process with a flat PSD, $S_{qq}(e^{j\omega}) = 1/12M^2$, thereby validating the PSD predictions in (4). The sufficient conditions have been shown in [3] to be satisfied by DDSMs of 2^{nd} order or higher.

Simulations in MATLAB have confirmed these results, as shown in Fig. 4. The simulated PSD of the shaped quantization noise, $e_{\Delta\Sigma}[n]$, without and with LSB dithering, in a 2^{nd} order DDSM are plotted in Figs. 4(a) and 4(b) respectively. Note that the x-axis shows an absolute frequency scale when the unit sampling rate is assumed to f_{ref} = 48 MHz. The 2^{nd} order DDSM used $M = 1$, and $p = 16$ and applied a constant input. The corresponding PSDs of $s_0(t)$ obtained when the DDSM was used in a frequency synthesizer with $N = 51$, f_{ref} = 48 MHz are plotted in Figs. 4(c) and 4(d): zoomed-in plots showing the PSD on one side of the central spike are shown in Figs. 4(e) and 4(f).

Effectively, LSB dithering alters the properties of the quantization error of the DDSM: the two-level additive random sequence, $d[n]$, eliminates limit cycles and imparted

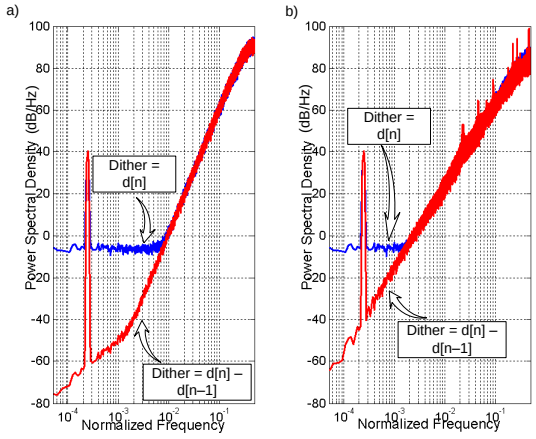

Fig. 5. Simulated power spectral densities of shaped quantization noise (a) with unfiltered LSB dither, and (b) with filtered LSB dither, in a 3^{rd} order digital delta-sigma modulator. Simulated power spectral densities of shaped quantization noise (a) with unfiltered LSB dither, and (b) with filtered LSB dither, in a 2^{nd} order digital delta-sigma modulator.

desired properties of whiteness on the quantization error, $q[n]$. Consequently, a circuit engineer trying to design a DDSM based frequency synthesizer could simply use a 2^{nd} or higher order DDSM, make sure that the quantizer does not overload, and employ LSB dithering: he/she is guaranteed to avoid PSD spikes due to DDSM limit cycles, and can exactly predict the PSD of the phase noise of $s_0(t)$.

However, this quantization noise conditioning comes with a minor penalty: elevated phase noise caused by $d[n]$, particularly when p is small. Fortunately, further conditioning can be employed to suppress the noise contribution within the signal band, B, by filtering $d[n]$ prior to addition [4]. As shown in [4], as long as $d[n]$ undergoes net two or more integrations before it reaches the quantizer, it ensures that the quantization error, $q[n]$ has the aforementioned desired properties of independence and whiteness. Again, the quantization noise conditioning using dither is illustrated in Fig.5 when the dither sequence is filtered using a $(1-z^{-1}$ filter i.e. when $d[n]-d[n-1]$ is added to the DDSM input. In a 3^{rd} order DDSM, even with such filtering, the random sequence, $d[n]$, sees net two integrations so, limit cycles are eliminated, as is evident from Fig. 5(a). In contrast, in a 2^{nd} order DDSM, with filtering, $d[n]$ sees only net integrations so, limit cycles are not eliminated, as is evident from Fig. 5. Note that in all cases, the simulation employed a sinusoidal input sequence.

Note that there is scope and necessity for further improvement. For instance, the LSB dithering technique is not successful when employed in DDSMs with quantizer overload: the aforementioned conditions are not satisfied. However, as described in the previous subsection, DDSMs with limited range of $v[n]$, particularly those that limit $v[n]$ to only two values, are desirable owing to circuit implementation conditions. It is conceivable that changing the statistical properties of $d[n]$ may allow the elimination of limit cycles in DDSMs with overloaded quantizers as well.

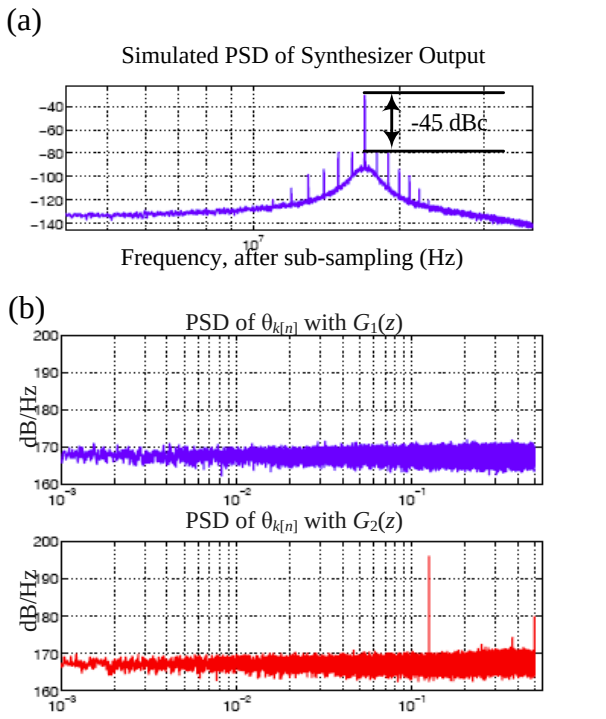

Fig. 6. (a) Simulated power spectral density of synthesizer output in the presence of phase errors. (b) Simulated power spectral density of $\theta_{k[n]}$ using $G_d(z)=G_1(z)$ and $G_d(z)=G_2(z)$.

V. EFFECTS OF CIRCUIT ERRORS ON QUANTIZATION NOISE

In most DDSM based frequency synthesizers, performance is limited by inevitable transistor fabrication errors and circuit non-linearity. Specifically, undesired spikes in the PSD of the output and elevated noise in the desired signal band, $Bf_{ref}/2\pi$, result from these errors. While there are several sources of such errors, only two important ones are discussed here for the sake of simplicity. In each case, the problem is described and a recent quantization noise conditioning technique that addresses it is briefly described.

A. Phase Errors

One of the main problems with the DDSM based frequency synthesizer described in Section III is caused by errors in the phases of the M periodic signals, $s_k(t), k = 0,1,...,M-1$. Due to transistor sizing mismatches and systematic circuit errors in the VCO, the M periodic signals include small phase errors:

$$s_k(t) = cos(2\pi(N+y(t))f_{ref}t + \frac{2\pi k}{M} + \theta_k),$$

where $f_{ref}(t)$ is given by (7).

The deleterious effects of these VCO phase errors are illustrated in Fig. 6(a). The figure shows plots of the simulated PSD of the phase noise and the output, $s_0(t)$, of a DDSM based frequency synthesizer with $M=16$, f_{ref} = 48 MHz, $\alpha = 1/3$, $N=51$, a 3^{rd} order DDSM. The θ_k are chosen randomly from a zero mean, Gaussian distribution

with a root mean square (r.m.s.) value of 10% of the nominal phase step between adjacent VCO output signals, $2\pi/M$. As is evident from the figure, even such small phase errors result in very strong pikes making this frequency synthesizer unsuitable for use even in applications that have relatively relaxed requirements such as Bluetooth.

The origin of the strong spikes is best understood by referring to Fig. 7(b). As described in Section III, the n^{th} positive going output transition of $v_{div}(t)$ is aligned with a positive going transition of $s_{k[n]}(t)$, where $k[n]$ is given by (5). Whenever this happens, a time error, $\tau_{err}[n]$ proportional to $\theta_{k[n]}$ is introduced into the negative feedback loop. These dynamic time errors add to the phase noise in $s_0(t)$ caused by the DDSM's shaped quantization noise. Specifically,

$$s_0(t) = cos\left(2\pi\left[N+y(t)\right]f_{ref}t + \sum_{m=0}^{\infty}\theta_{k[n]}g(t-nT_{ref})\right),$$

where $g(t)$ is a causal function determined by the dynamics of the synthesizer's negative feedback loop. Its actual form is not important for the present discussion.

Apart from a scaling factor, the effect of these phase errors is that of a one-to-one mapping from $k[n]$ to $\theta_{k[n]}$ i.e. as a memoryless lookup table. Closer observation of (5) reveals that it describes the behavior of a 1^{st} order DDSM that quantizes a sequence of fractions, $Mv[n]$, to a sequence of integers, and that $k[n]$ is its quantization error. It is well known that the samples of the quantization error of a 1^{st} order DDSM are definitely not independent of the DDSM input sequence, or white. Consequently, it is not surprising that the non-linear mapping, $\theta_{k[n]}$, causes strong spikes. Traditional circuit techniques to address such a problem would involve careful design and/or elaborate calibration and correction efforts to eliminate the errors, θ_k. However, they are too small to reliably correct in either approach. The author of this paper and colleagues have reported a recent quantization noise conditioning technique [5] that effectively addresses this technique very simply, as described below.

The technique filters a binary random sequence, $\eta[n] \in 0,1$, with a finite impulse response filter, $G_d(z)$, and adds the result to the calculation of $k[n]$:

$$k[n] = (k[n-1] + v[n] + \eta[n] * g_d[n]) mod M, \qquad (10)$$

where $g_d[n]$ is a sequence of integers whose Z-transform is $G_d(z)$. It was proved in [5] that if $\forall i = 0,1,...,Q-1, \exists$ integer $n = n(i)$ such that $|g_d[n]| = 2^i$, then, the PSD of the sequence, $\theta_{k[n]}$, does not have any spikes as long as the mapping is not one-to-many i.e. $\theta_{k1} \neq \theta_{k2}\ \forall k1 \neq k2$. Since the phase errors occur randomly, the likelihood of getting equal errors is very low. Readers interested in the details of the mathematical proof are referred to [5].

Essentially, the filtered dither alters the statistical properties of the 1^{st} order DDSM's quantization noise, $k[n]/M$, such that it becomes insensitive to an arbitrary non-linear, one-to-one mapping. The technique is illustrated in Fig. 6(b)

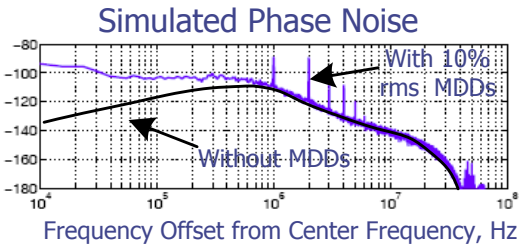

Fig. 7. Simulated power spectral density of synthesizer output in the presence of modulus dependet divider delays.

by plotting the simulated PSD of $\theta_{k[n]}$ for two cases of $G_d(z)$,

$$G_1(z) = 1 - 2z^{-1} + 4z^{-2} - 8z^{-3} + 8z^{-4} - 4z^{-5} + 2z^{-6} - z^{-7}$$
$$G_2(z) = 1 - 3z^{-1} + 5z^{-2} - 5z^{-3} + 3z^{-4} - z^{-5}.$$

Note that $M = 16$; while $G_1(z)$ satisfies the conditions, $G_2(z)$ does not. The same experiment was repeated for several randomly chosen sets of phase errors, $theta_k$, with similar results; only one of them is shown here. As is evident, if the conditions on $G_d(z)$ are not satisfied, stong spikes in the PSD are observed; if they are satisfied, spikes are completely eliminated.

The technique, general and powerful as it is, comes with a penalty: the filtered dither adds significant amount of phase noise. Fortunately, the dither filter $G_d(z)$ can be chosen such that the contribution is suppressed aggressively within the signals band, $Bf_{ref}/2\pi$. Outside this band, the dynamics of the synthesizer, given by $H(s)$, suppress it significantly.

B. Modulus Dependent Delays

The effects of such circuit errors are illustrated below in the context of the DDSM based FNPLL for which $N = 51$, $M = 1$, $\alpha = 1/48$, $m[n] = 0$, f_{ref} = 48 MHz, and a 2^{nd} order DDSM is used to generate $v[n]$. Note that LSB dithering is employed to render $q[n]$ white as described in the previous section. As described in Section III, ideally, $s_0(t)$ would have a frequency of 2.489 GHz; the phase noise resulting DDSM quantization noise would result in error in the FNPLL output's phase whose (simulated) PSD would be as shown in Fig. 4(f).

Inevitable circuit errors result in delays through the divider circuit blocks that depend on the value of $v[n]$. These errors, called modulus dependent delays (MDDs), can be represented by a non-linear mapping, $\Delta_{v[n]}$. They disturb the phase of the FNPLL's output to the tune of a non-linear function of $v[n]$. Specifically,

$$s_0(t) = cos\left(2\pi\left[N+y(t)\right]f_{ref}t + \sum_{m=0}^{\infty}\Delta_{v[n]}g(t-nT_{ref})\right),$$

where $y(t)$ is given by (7). Fig. 7 plots the simulated PSD of the phase noise of the resultant $s_0(t)$; MDDs were randomly chosen according to a Gaussian distribution with a standard deviation of just 10% of the VCO period. As is evident from the figure, the MDDs cause both the discrete spikes to reappear and a significant increase in the phase noise for

low frequency offsets. This effect is commonly referred to as quantization noise "folding over" and is difficult to analyze and/or quantitatively predict.

While circuit errors e.g., MDDs cause the discrete spikes, ultimately, the properties of $q[n]$ are to blame. The reasons are not surprising. Consider the output of the 2^{nd} order DDSM for $m[n] = 0$:

$$v[n] = \alpha + q[n] - 2q[n-1] + q[n-2]$$

. If the table look-up non-linearity were represented by a simple mapping function, $g(v[n]) = (v[n])^2$, then

$$g(v[n]) = \alpha^2 + 2\alpha \sum_{m=0}^{2} e[n-m] + \sum_{m=0}^{2}\sum_{l=0}^{2} e[n-m]e[n-l]$$

Note that LSB dithering renders $q[n]$ and $q[m]$ jointly independent (even if asymptotically as $n \to \infty$). However, that is not sufficient to ensure that the autocorrelation of the product terms in the double summation is not periodic. Consequently, in spite of LSB dithering, even simple non-linearities such as squaring can expose underlying correlations and cause spikes in the PSD of overall phase noise.

A recent technique called "successive requantization" employs a Markov chain based quantizer instead of a DDSM based quantizer to make the synthesizer insensitive to modulus dependent divider delays [6]. The technique is best explained by considering the simple case where $M = 2^p/2$. For simplicity, assume that,

$$x[n] \in \{-2^{p-1}, ..., 0, ..., 2^{p-1} - 1\}/2^p.$$

Then, the technique quantizes as follows:

$$v[n] = \frac{1}{2}\left(x[n] + s[n]\right),$$

where

$$2^p s[n] = \begin{cases} 2l[n], & \text{if } 2^p x[n] \text{ is even,} \\ 2l[n]+1, & \text{if } 2^p x[n] \text{ is odd,} \end{cases}$$

where $l[n]$ is an integer. Note that quantization by only one-bit is achieved. The technique chooses $2^p s[n]$ according to the three additional criterion:

1) $s[n]$ is chosen with the goal of minimizing the running sum of $2^p s[n]$, namely $t[n]$, where

$$t[n] = \sum_{m=0}^{n-1} s[m].$$

   For example, if $t[n]$ is positive, the technique tries to choose a negative value of $s[n]$.
2) $s[n]$ is chosen randomly whenever possible.
3) $s[n]$ is chosen such that the $E\{s[n]^2\}$ is a constant $\forall n$.

Since the first criterion bounds the first sum of $2^p s[n]$, it can be shown that the power spectral density of $s[n]$ has a zero at dc [6]. The second criterion ensures that the autocorrelation of $s[n]$ and hence that of $v[n]$ is not periodic thereby eliminating spikes in their PSDs. To appreciate the third criterion, consider the effect of a square non-linearity on $v[n]$, for the simple case of $x[n] = \alpha$, a constant:

$$v[n]^2 = \frac{1}{4}\left(\alpha^2 + 2\alpha s[n] + s[n]^2\right).$$

By ensuring that $E\{s[n]^2\}$ is a constant $\forall n$, the PSD of $v[n]^2$ will not exhibit any spikes in its PSD.

Essentially, $s[n]$ is chosen according to the state value, $t[n]$, of a finite state machine. The transition probabilities of the finite state machine may vary with n depending on whether $2^p x[n]$ is even or odd, implying that the underlying Markov chain is potentially non-homogenous. However, as long as the transition probabilities are chosen such that $E\{s[n]^2\}$ is a constant $\forall n$, $v[n]$ is insensitive to square non-linearity. The technique quantizes to large values of M by applying the procedure recursively (on $v[n]$, on the result of quantizing $v[n]$ by one-bit, and so on). Furthermore, the technique also can ensure insensitivity to higher order non-linearities as well.

Note that the shaped quantization noise in this technique is significantly higher than in the DDSM's case. Furthermore, it is not aggressively shaped: the shape of its PSD is similar to that of a 1st order DDSM, but higher. However, the technique represents a promising approach to quantization that could be particularly useful to practical integrated DAC design.

VI. Conclusions

Digital delta-sigma modulators enable high resolution digital-to-analog data conversion from simple circuits. However, limitations of quantization noise shaping, and circuit errors limit the performance of DDSM based D/A converters. Recent quantization noise conditioning techniques such as dithering and have successfully addressed these problems. Such problems and example conditioning techniques were described in the context of a delta-sigma modulator based frequency synthesizer. A few unresolved problems that could benefit from theoretical advances in the study of quantization were also identified.

References

[1] S. H. Ardalan, and J. J. Paulos, "An Analysis of Non-Linear Behavior in Delta-Sigma Modulators," *IEEE Transactions on Circuits and Systems*, vol. CAS-34, no. 6, June 1987, pp. 593-603.

[2] I. Galton, "Granular quantization noise in a class of delta-sigma modulators," *IEEE Transactions on Information Theory*, Vol. 40, No. 3, 1994, pp. 848-859.

[3] S. Pamarti, J. Welz, I. Galton, "Statistics of the Quantization Noise in One-Bit Dithered Single Quantizer Delta-Sigma Modulators," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 54, no. 3, March 2007, pp. 492-503.

[4] S. Pamarti, I. Galton, "LSB Dithering in MASH Delta-Sigma D/A Converters," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 54, no. 4, April 2007, pp. 779-790.

[5] S. Delshadpour and S. Pamarti, "A Spur Elimination Technique for Phase Interpolation Based Fractional-N PLLs," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 55, no. 6, July 2008, pp. 1639-47.

[6] A. Swaminathan, et al, "A digital quantizer with shaped quantization noise that remains well behaved after nonlinear distortion," *IEEE Trans. Signal Process.*, vol. 55, pp. 5382, Nov. 2007.