

A Fractional- N Frequency Synthesizer Architecture Utilizing a Mismatch Compensated PFD/DAC Structure for Reduced Quantization-Induced Phase Noise

Scott E. Meninger and Michael H. Perrott

Abstract—Techniques are proposed to dramatically reduce the impact of quantization noise in $\Sigma\Delta$ fractional- N synthesizers, thereby improving the existing tradeoff between phase noise and bandwidth that exists in these systems. The key innovation is the introduction of new techniques to overcome nonidealities in a phase-frequency detector (PFD)/digital-to-analog converter (DAC) structure, which combines the functionality of both phase detector and cancellation DAC into a single element. The proposed architecture achieves better gain matching between the phase-error signal and cancellation DAC than offered by previous approaches. Dynamic element matching techniques are introduced to mitigate the effects of PFD/DAC unit element and timing mismatch on synthesizer phase noise performance. We present behavioral simulations of an example application of this technique that demonstrates 36-dB reduction in broad-band quantization-induced phase noise with the use of a 7-b PFD/DAC. Simulations further demonstrate that fractional spurs are rejected to levels < -90 dBc when a low-cost, low-overhead digital gain correction technique is employed.

Index Terms—Frequency synthesizers, phased-locked loops (PLLs), phase noise, sigma-delta modulation.

I. INTRODUCTION

FREQUENCY synthesis is an essential technique employed in RF systems to achieve local oscillator (LO) generation or direct modulation transmission. Fractional- N synthesis offers the advantage over integer- N based systems of decoupling the choice of synthesizer resolution from bandwidth. Fast-settling, high-resolution synthesis becomes possible, giving greater design flexibility at the system level. Fractional- N synthesis can be separated into two categories: classical fractional- N synthesis and $\Sigma\Delta$ fractional- N synthesis.

The classical approach to fractional- N synthesizer design employs dithering and phase interpolation, as depicted in Fig. 1 [1]. An accumulator carry-out signal is used to dither the control input to a multimodulus divider such that a fractional average divide value is obtained from a divider that supports integer values. A digital-to-analog converter (DAC) is used in conjunction with a phase accumulation register to cancel out periodicities in the phase-error signal E . The main performance lim-




Fig. 1. Classic fractional- N synthesis with phase interpolation.

itation of this approach centers around the difficulty in creating a precise match between the DAC output and the phase-error signal. Spurious performance of synthesizers based on this architecture is typically around -60 dBc, which is too high for use in the most aggressive RF LO applications.

In $\Sigma\Delta$ fractional- N synthesis [2]–[7], the spurious performance is improved through $\Sigma\Delta$ modulation of the divider control. The quantization noise introduced by dithering the divide value is whitened and shaped to high frequencies, such that it is substantially filtered by the synthesizer dynamics. In order to obtain sufficient randomization to reduce spurs to negligible levels, $\Sigma\Delta$ modulators of order 3 or higher (often employing LSB dithering) are required, necessitating a higher order loop-filter to counteract increased noise slope. The shaped quantization noise often dominates at high offset frequencies, introducing a noise-bandwidth tradeoff, which translates to low-closed loop bandwidths for low phase noise synthesizers. This tradeoff somewhat negates the central idea behind fractional- N synthesis, which is to increase synthesizer bandwidth.

Two approaches have emerged to reduce the impact of the noise-bandwidth tradeoff. The first involves reducing the quantization step-size of the divide value dithering action through the use of multiple voltage-controlled oscillator (VCO) (or divider) phases [8]–[10]. While introducing multiple VCO or divider phases is the ideal means by which to reduce the quantization step-size, in practice the number of phases possible is limited. The phase resolution is often set by a gate delay, which for high-frequency outputs, can be a significant fraction of the VCO period. Additionally, to generate multiple phases, either ring

Manuscript received May 1, 2003; revised July 2003. This paper was recommended by Guest Editor M. Perrott.

The authors are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: coyote@mh.mit.edu).

Digital Object Identifier 10.1109/TCSII.2003.819114

oscillators must be used, which have inherently poorer phase noise performance than LC oscillators, or a delayed-locked loop (DLL) with all of the associated overhead must be employed [10]. Mismatch between the phases appears and must be carefully dealt with. A calibration scheme included in [8] resulted in -55 -dBc fractional spur levels.

Very recently, a modification has been proposed to the selection logic used in the multiphase approach of [8]. In this case, very high-order $\Sigma\Delta$ modulation is applied to the phase selection mux in order to make each tap equi-probable in a histogram sense [11]. For the simulated DC synthesizer input presented, < -80 dBc was obtained via seventh-order modulation of the selects to a 16-input multiphase mux. Noise shaping of this order requires large numbers of extra poles in the loop-filter to counteract the increased noise slope. Additionally, the multiphase approach is ultimately limited in its ability to reduce broad-band phase noise by limitations in creating the delay.

The second approach to reduce the noise-bandwidth tradeoff, uses a DAC to cancel the error signal [12]. This method builds on the idea behind phase interpolation, but introduces $\Sigma\Delta$ design techniques to reduce the impact of DAC nonlinearity. The main limitations with this architecture center around the achievement of good matching between the DAC output and phase-error signal, which is difficult because the two are processed by separate circuits whose outputs are summed. Also, a high-resolution DAC is required to achieve desired performance. In [12], a 16-b coarse DAC and 16-b fine DAC were used, resulting in 16-dB improvement in broad-band phase noise and -60 -dBc fractional spur levels.

An alternative approach that utilizes a DAC to reduce quantization-induced phase noise is proposed in [13]. The separate phase detector and DAC circuit elements are replaced by a hybrid structure. By embedding the two functions into one circuit, an intrinsically better gain match between the phase-error and DAC cancellation signals is obtained. However, the architecture presented in [13] does not address the issue of mismatch between unit elements of the DAC, or between the timing signals in the phase detector, which will result in incomplete phase-error signal cancellation and spurious feed-through. While both of the new techniques succeed in reducing wideband phase-noise levels by effectively reducing the phase quantization step-size, spurious performance is on the same order as that reported in classical fractional- N architectures.

The proposed architecture, shown in Fig. 2, leverages advances in noise-shaping DAC design to ease the requirements on the cancellation DAC in the traditional fractional- N approach in a manner similar to [12]. It utilizes a PFD/DAC structure [13] to obtain a good intrinsic gain match between the phase-error and DAC cancellation signal. However, this work makes the key contribution of introducing techniques to minimize the impact that PFD/DAC unit element and timing mismatch sources have on phase noise performance. Indeed, matching issues create the bottleneck in previous approaches since they result in error feed-through that is manifested in the phase noise spectrum as large spurs, or increased broad-band phase noise. The proposed architecture incorporates several digital signal processing techniques to reduce the impact of nonidealities that occur in the PFD/DAC such as unit element mismatch, timing mismatch, and




Fig. 2. Proposed architecture.

any residual gain mismatch occurring between the PFD/DAC output and phase-error signal.

II. PROPOSED APPROACH

A key issue with prior fractional- N synthesizer implementations is that the cancellation DAC output and phase-error signal are poorly matched. This error is a direct result of the fact that separate circuits have been used to implement the two required blocks. Fig. 3 depicts the proposed PFD/DAC structure, which greatly reduces mismatch between the two signal paths. The proposed PFD/DAC differs from a prior implementation of a hybrid phase detector/cancellation DAC scheme [13] in that it compensates for mismatches within the PFD/DAC structure itself. As will be discussed, mismatch between magnitudes of the phase-error and cancellation signal and timing mismatch between signal paths in processing phase information is a key consideration for achieving a high quality overall gain match.

A brief description of the PFD/DAC circuit architecture is as follows. A register based delay cell is used to create a delayed divider phase. These signals then pass through a timing mismatch compensation and resynchronization block, which accounts for timing mismatches in the two critical phase information paths. The two output phases from the timing mismatch compensation and resynchronization block are compared to a reference via two phase detectors that control the charge-pump. The output of the digital processor from Fig. 2 is input to the DAC mismatch shaping block which accounts for errors between the unit elements which feed the charge-pump. The output of the DAC mismatch shaping block steers a bank of current source DAC elements to the phase detector controlled charge-pump. The output of the charge-pump is sent to the loop-filter. While a charge-pump based PLL will be presented, the PFD/DAC could be implemented with switched voltage sources as well.

Fig. 4 offers an intuitive explanation of the PFD/DAC approach. The top sequence of waveforms depicts the output of the charge-pump when a multiphase divider is used to reduce the quantization step-size, as in [8] and [9]. The resolution of the multiphase divider has been set to four for this example. This approach will be referred to as vertical slicing, since the VCO period is being quantized in time/phase. The bottom waveforms represent the outputs of the PFD/DAC at the same quantization level. Since in this case the magnitude of the charge-pump




Fig. 3. Proposed PFD/DAC.




Fig. 4. Vertical and horizontal resolution of the VCO period.

output inside the dashed box is quantized, this approach will be referred to as horizontal slicing. The variable ϵ represents the LSB outputs of the accumulator, which is a measure of the phase-error. The dashed box represents the possible edge locations of the divider output if it were quantized to two locations one VCO period apart (T_{vco} in the figure), and the magnitude of the charge-pump output if it were quantized to two levels. In other words, the dashed box represents the possible values of charge-pump current and locations of divider edge in a classical phase interpolation based synthesizer with a quantization step-size equal to one VCO period.

Vertical slicing changes the amount of charge that can be delivered inside the dashed box, thereby reducing the quantization step-size. The particular phase chosen is determined by the accumulator residue ϵ . Horizontal slicing accomplishes the same goal of reducing the quantization step-size in an orthogonal manner. Rather than choosing different divider phases, the charge-pump current is delivered in two steps. Residue ϵ is used to control how much current is delivered on the first phase $[(1 - \epsilon) \cdot I_{chp}]$ is delivered], with the remainder being added in on the second phase. The net charge inside the dashed box




Fig. 5. Operation of the extended range PFD/DAC.

is the same for both cases. It is much easier to increase resolution in a DAC by adding more unit elements than to resolve finer and finer time-steps in a multiphase approach. Horizontal resolution allows for a much higher resolution solution, a much greater achievable reduction in quantization step-size and, therefore, a much greater obtainable reduction in broad-band quantization-induced phase noise. Processing the phase-error signal created by the charge-pump in the same circuitry that generates the cancellation signal via the DAC elements results in an inherent match between these signals. Well understood mismatch shaping techniques are applied via the DAC mismatch shaping block to alleviate any mismatches that occur in the unit elements making up the PFD/DAC [14].

In practice, the $\Sigma\Delta$ current steering control will vary between zero and full-scale. For certain implementations of the current source and charge-pump structure, debiasing the current steering circuitry internal to the charge-pump at either extreme may become an issue. A solution to this problem is to implement the register based delay cell as two registers in series, creating a total delay of two VCO cycles. In order to maintain bias current in each current steering path, one fourth of the full scale current is directed to the early path and one fourth to the late path. The remaining half full-scale current is controlled by the digital $\Sigma\Delta$ modulator. This situation is depicted in Fig. 5, where the total charge-pump current has been normalized to one. The bias and $\Sigma\Delta$ controlled currents add so that, in each current steering path, a minimum of one-fourth full-scale current is always present. The unit elements that are summed to generate the one-fourth full-scale bias currents are shuffled in combination with the $\Sigma\Delta$ controlled unit elements to minimize the impact of mismatch.

The number of unit elements required for the extended range structure is double that of the nonextended range PFD/DAC to achieve the same effective resolution. An extra bit of full-scale resolution is therefore required to obtain the same level of broad-band noise reduction. The extended range structure uses half of the full scale current for biasing purposes. Therefore, the effective resolution, or number of bits that are controlled by the digital $\Sigma\Delta$ modulator, is decreased by 1 b when compared to the nonextended range case. For reasonable implementations of the PFD/DAC (<10 b) the increased area penalty required to achieve a desired effective resolution may be acceptable to improve charge-pump linearity and to ease its design.

A. Digital Gain Error Compensation

While vertical slicing is limited in resolution for practical reasons, it is the preferred approach in an ideal sense because the divider phase quantization is reduced directly. Although the net charge enclosed by the dashed boxes of Fig. 4 are the same for vertical and horizontal slicing, their *shapes* are different. This suggests that horizontal slicing will achieve a very good DC match between the phase-error and cancellation signals, but some frequency dependent gain error will result.

Fig. 6 illustrates the systematic frequency dependent gain error that results from the shape mismatch between the horizontally and vertically sliced waveforms. Time-domain behavior, as depicted in the left plot, is represented by the number of steps being resolved for the two techniques, and the corresponding charge transferred during the resolved VCO period. To simplify analysis, the charge-pump magnitude is normalized to one. A look at the Fourier transforms for each shows that at dc the difference between the spectra is zero, as expected since the shaded regions in the time domain plot have the same area. As frequency increases, the difference between the spectra increases and will be manifested in the phase noise power spectrum by imperfect fractional spur cancellation. This behavior places a limit on the ability of the horizontally resolved system to exactly cancel the phase-error waveform in the absence of a correction scheme.

There are two ways to compensate for the error introduced by the shape of the horizontally resolved waveform. The first is to use well-known sample-and-hold techniques in a similar manner as applied in classical phase interpolation [1]. The shape mismatch could theoretically be eliminated, resulting in no fractional spur feed-through. Additionally, reference feed-through is mitigated. We are currently investigating this analog, circuit-focused technique, but present here an alternative method that compensates this issue using a completely digital, rather than analog, approach.

Our proposed alternative to traditional sample-and-hold techniques becomes evident upon analysis of the error resulting from horizontal slicing. The Fourier transform for the charge enclosed inside the dashed box for the vertical slicing case is

$$Q_{\text{vert}}(j\omega) = \frac{1}{j\omega} (e^{-j\omega\epsilon_v} - e^{-j\omega T_{\text{VCO}}}) \quad (1)$$

and for the horizontal case is

$$Q_{\text{horiz}}(j\omega) = \frac{\epsilon_h}{j\omega} (1 - e^{j\omega T_{\text{VCO}}}). \quad (2)$$




Fig. 6. Vertical versus horizontal resolution. (a) Time domain charge behavior. (b) Frequency domain behavior.

By expanding these expressions using Taylor series, keeping up to the second order terms, and subtracting to obtain the error, we arrive at

$$Q_{\text{err}}(j\omega) \approx -\frac{j\omega}{2} \epsilon(1 - \epsilon) \quad (3)$$

where ϵ corresponds to the accumulator residue.

Having arrived at a simple closed form approximation for the expected error, we can build a digital gain compensation block to correct for it. Fig. 7 depicts the implementation of the gain compensation block. Some portion of the accumulator residue bits are used to address a look up table (LUT), which may be implemented as read-only memory (ROM) or random-access memory (RAM). The output of the LUT is differentiated and summed with the residue. This is then sent to the digital $\Sigma\Delta$ modulator controlling the PFD/DAC.

There are two points to note about the digital gain compensation block. The first point is that in practice the LUT has finite input and output resolution. In Section III-B, the impact of




Fig. 7. Digital gain compensation block.

changing the values of X and Y , the LUT input and output resolutions, respectively, will be presented. It will be shown that only 1 Kb of ROM (with $R = 20$, $X = 6$, and $Y = 4$) is required to achieve an 18 dB improvement in fractional spur rejection. This additional rejection, coupled with the improved gain match due to the PFD/DAC, achieves overall fractional spur levels of < -90 dBc in detailed behavioral level simulations.

The second point relates to the variable N_{nom} used to calculate the compensated accumulator output. The factor of $1/N_{\text{nom}}$, where N_{nom} is the nominal divide value, stems from the fact that the compensator is clocked at the reference (or divider) frequency, whereas the phase-error is referred to the VCO period. Ideally, this renormalization factor needs to vary with the instantaneous divide value, requiring a full digital divider. Simulations have shown that using a static $1/N_{\text{nom}}$ value yields good results. As long as the approximation error is less than the desired level of compensation, the approximation is acceptable. For a 20 dB improvement, the VCO frequency can change by $\pm 5\%$ from the nominal value while maintaining a valid approximation. If the synthesizer is employed in a system with requirements exceeding this range, multiple LUTs can be employed to keep the approximation error acceptable, or a full-digital divider can be implemented in the compensation block. Since a 1 Kb LUT occupies very little on-chip area, multiple LUTs appear to be the best solution. The compact nature of memory in modern processes coupled with the small memory size required to achieve high levels of compensation translates into a very low area penalty for a large degree of design flexibility.

B. Timing Mismatch Compensation

Thus far, we have examined element mismatch in the PFD/DAC which may be dealt with using well understood mismatch shaping techniques [14], and a frequency dependent gain error which results from the pulse-shape of horizontal slicing that can be compensated using sample-and-hold techniques or a digital signal processing block. An additional source of mismatch in the system stems from signal skew inside the PFD/DAC. This issue is depicted in Fig. 8. Mismatch in physical layout, loading, and device gradients results in a propagation delay difference between the signal paths for Φ_0 and Φ_1 . The value of time resolved in Fig. 6, therefore, will not equal T_{VCO} in practice. This timing mismatch results in a gain error and fractional spur feed-through.

We can apply dynamic element matching techniques to correct for timing mismatch. As shown in Fig. 8, use of retiming flip-flops limits the skew between the two phase paths to differences between the flip-flop clk-to-q times and PFD circuit paths. The muxes are toggled by a phase swap signal so that the two phase paths see each PFD, on average, the same amount of

time. The current steering control bits from the $\Sigma\Delta$ modulator are selectively inverted to maintain correct functionality. It is possible to embed the muxes into the flip flops, and to embed the charge-pump into the phase-detector structure, greatly enhancing intrinsic matching.

The consequences of introducing the phase swapping process can most easily be understood via a straightforward time analysis. From Fig. 8, we see that the retiming flip-flops reset the timing error. We lump all of remaining mismatch between the two paths into a variable, Δ_t , which is referenced to the output of one of the flip-flops. In the example shown, Δ_t is referenced to the lower flop. Following through the time evolution of the phase paths, we see that the path for Φ_0 experiences an average delay of

$$t_{\text{del}\Phi_0} = (1 - D) \cdot 0 + D \cdot \Delta_t \quad (4)$$

while the path for Φ_1 sees

$$t_{\text{del}\Phi_1} = (1 - D) \cdot \Delta_t + D \cdot 0 \quad (5)$$

where D is the duty cycle of the swap control. Clearly, if the duty cycle is set to be 0.5, each path will see the same average delay and the timing mismatch is eliminated.

A subtlety to note is that the swapping process can be represented by a multiplication between the phase information in each path (Φ_0 and Φ_1) and an error signal whose amplitude is varied between zero and Δ_t by the swap control signal. Since both Φ_0 and Φ_1 contain spurious content, the swap control should be made nonperiodic in order to avoid the generation of mixing products. Therefore, two constraints are placed on the characteristics of the phase swapping control signal. First, it must have an average value very near 0.5 to ensure that both phase signals see the same average propagation times. Second, the swap signal must contain little or no spurious energy. For this reason, two possible implementations arise for control of the swapping operation: a pseudorandom linear feedback shift register (LFSR) or a single bit output $\Sigma\Delta$ modulator with a sufficient order to ensure that the output spectrum is random. We have found the LFSR to be the better solution.

III. RESULTS

Since the most aggressive synthesizer performance requirements often exist in LO applications, unmodulated synthesizer performance will be presented at the behavioral simulation level. The proposed technique can also be used with modulated synthesizers. In the discussion to follow, all design calculations for the synthesizer dynamics were performed using the PLL Design Assistant tool described in [15]. Simulations were done using the CppSim behavioral level C++ simulator presented in [16], which is based on area conservation principles and has been verified to show excellent agreement with measured results [17]. Both tools can be downloaded.¹

The two goals of the proposed approach are: 1) to reduce the broad-band quantization-induced phase noise so that synthesizer bandwidth can be extended and 2) achieve sufficiently high levels of fractional spur rejection that the synthesizer is useful

¹Online. Available: <http://www-mtl.mit.edu/research/perrottgrouptools.html>




Fig. 8. Timing mismatch compensation.

in a wide variety of applications. These performance goals will each be presented in turn.

A. Predicted versus Simulated Broad-Band $L(\Delta f)$

Fig. 9 shows the *calculated* performance for a synthesizer utilizing a 20-b accumulator. The proposed architecture with a 7-b, first-order $\Sigma\Delta$ controlled extended range PFD/DAC is compared to a classical second-order $\Sigma\Delta$ frequency synthesizer. The simulated reference frequency is 50 MHz, and the dual modulus divider is a divide-by-100/101, resulting in a 5-GHz output frequency. Sources of phase noise are as follows. The charge pump noise model is set so that its simulated low-frequency contribution to phase noise is -105 dBc/Hz. The simulation model of the VCO is set so that it exhibits -152 dBc/Hz noise at a 20-MHz offset. These numbers represent state-of-the-art performance and are used as baseline number in an aggressive synthesizer design—the VCO phase noise for an output frequency of 5 GHz as simulated here is equivalent to the requirement for a 900 MHz GSM VCO. The synthesizer bandwidth is set to 1 MHz, which is a value 10X higher than typical bandwidths reported in the literature. Additional poles at 4 and 5 MHz are added to reduce the effect of the reference spur. This is a standard approach in the literature. Finally, the accumulator input LSB is always set high to aid in generating a more randomized residue. Setting the accumulator LSB is done in place of introducing a separate dither source due to its simplicity.

As Fig. 9 demonstrates, the second-order $\Sigma\Delta$ synthesizer shaped quantization noise is dominant with such a high bandwidth when using the classical approach. By contrast, the $\Sigma\Delta$ noise is reduced by 36 dB using the new approach, and is no longer dominant at any frequency. The broad-band phase noise performance is therefore determined solely by the charge-pump and VCO!

Fig. 10 is a behavioral level *simulated* phase noise spectrum for a synthesizer using the new approach, utilizing a 7-b PFD/DAC. The unit elements in the PFD/DAC are mismatched according to a Gaussian profile with $\sigma = 2\%$. Data weight




Fig. 9. Calculated improvement using the proposed approach. (a) Classic second-order $\Sigma\Delta$ synthesizer. (b) Proposed architecture.

averaging [14] is used to randomize selection of the unit elements. Phase swapping controlled by a 28-register LFSR is





Fig. 10. Simulated results of the proposed approach.

used to ameliorate a 2 ps timing skew introduced between the two phase paths in the PFD/DAC. The PFD/DAC itself utilizes the extended range approach described in Section II. The PLL is a type-II second-order loop with a Butterworth response.


As can be seen, the simulated spectrum matches the predicted spectrum extremely well, and demonstrates that the quantization noise has indeed been reduced by 36 dB! Larger reductions in broad-band quantization noise can be made by increasing the number of elements in the PFD/DAC. It should be noted that the PFD/DAC resolution sets the broad-band quantization-induced phase noise reduction independently of the synthesizer frequency resolution, which is set by the number of bits in the divider control accumulator. The example synthesizer presented here aims for 36 dB since a 7-b extended range PFD/DAC is a reasonable design goal given results reported in the literature for noise-shaping DACs.

Fig. 11 presents simulation results for the example 20-b synthesizer employing a 7-b PFD/DAC with the sources of mismatch varied. In order to isolate the effect of the mismatch on the quantization noise spectrum, the VCO and charge-pump noise sources are set to zero. Fig. 11(a) shows that the randomization of DAC elements using data weight averaging is very effective. Simulations reveal that for the Gaussian mismatch profile simulated, values of σ up to 0.05 do not impact the quantization noise spectrum. Values above this begin to exhibit some small levels of tone-feed-through at high offset frequencies. It is reasonable to assume that $\sigma < 0.02$ in practice if proper layout techniques are applied. The timing mismatch is set to zero for these simulations, since the phase swapping process results in broad-band, unshaped noise, which would mask the PFD/DAC element mismatch noise. The Gaussian mismatch generator seed was varied and several sets of simulation run to verify that the randomization scheme is robust. The plots presented in Fig. 11 are typical results for various values of σ .

Fig. 11(b) presents simulation results for the example synthesizer with PFD/DAC element mismatch set to zero, and timing mismatch, Δ_t (which is defined as the relative mismatch in time between the two PFD/DAC phase paths as depicted in Fig. 8), varied. It is apparent that the phase mismatch is converted to broad-band phase noise. There are two important points to note



(a)



(b)

Fig. 11. Effects of mismatch on broad-band noise performance 7-b PFD/DAC. (a) Unit element mismatch simulation. (b) Timing mismatch simulation.

in Fig. 11. The first is that only the quantization-induced phase noise is plotted in Fig. 11. The second point is that at low offset frequencies, charge-pump noise normally dominates the total phase noise profile. While the low-frequency quantization-induced phase noise in Fig. 11(b) increases as Δ_t increases, its magnitude will not overtake that of the -105 dBc/Hz charge-pump used by the example synthesizer presented in this paper until $\Delta_t \sim 5$ ps. The quantization-induced phase noise at intermediate and high offset frequencies in Fig. 11(b) does not change with increasing Δ_t . Since the quantization-induced phase noise is normally dominant over this frequency range, we can see that by maintaining the same level of performance for varying Δ_t , the phase swapping technique is very effective.

In order to evaluate the intrinsic limitations of the dynamic element matching techniques, a suite of simulations was run with the PFD/DAC resolution set to 14 b. Such a high resolution is useful for simulation as it sets the quantization noise level inherent to the PFD/DAC well below that of the mismatch sources. Fig. 12 presents the results of these simulations. For element mismatch simulations, the Gaussian mismatch generator




Fig. 12. Effects of mismatch on broad-band noise performance 14-b PFD/DAC. (a) Unit element mismatch simulation. (b) Timing mismatch simulation.

seed was varied over several values in the same manner as for the 7-b PFD/DAC. The broad-band noise is so greatly reduced that, at low levels of mismatch, the fractional spur for the dc accumulator input simulated is not masked by broad-band noise. It should be noted that, since the resolution bandwidth of the FFT used to calculate the power spectral densities in Fig. 12 is not 1 Hz, the spur level cannot be read directly from the plot. In order to obtain an accurate spur magnitude in dBc, the spur power has to be normalized to the VCO output power. This normalization methodology is followed in Section III-B, where spurious performance is presented. Focusing on the broad-band noise, the simulations suggest that PFD/DAC element mismatch noise has much less impact on the quantization noise spectrum than timing error. Both simulation runs demonstrate that the proposed technique is very robust in the face of mismatch errors.

B. Spurious Performance

Having established the ability of the proposed approach to reduce broad-band phase noise, its spurious performance is now

evaluated. In evaluating spurious performance, the VCO and charge-pump noise sources are set to zero. This is so that the broad-band noise associated with these system elements does not mask the presence of tones. Additionally, simulations have revealed that, if the number of bits in the PFD/DAC is greater than five (as is the case for the 7-b PFD/DAC being presented here), while the broad-band noise is reduced with increasing resolution, the spurious performance does not noticeably improve. This result is intuitive if we realize that the two factors determining the overall quality of fractional spur cancellation are the number of levels available in the PFD/DAC and the shape of the cancellation signal. Once the number of PFD/DAC levels increases past a critical value, any inability to completely cancel the fractional spur is dominated by the nonideal shape of the horizontal waveform. Therefore, in order that the broad-band PFD/DAC quantization noise not mask spurs, an infinite resolution (real number based) PFD/DAC is used to evaluate spur performance. Finally, having shown that phase swapping converts timing error in the PFD/DAC to broad-band noise, which could cover up spurs in the output spectrum, Δ_t is also set to zero for spurious simulations.

By removing all sources of broad-band noise, the true spur rejection capabilities of the proposed approach can be evaluated. Comparison simulations were done to verify the assumption that turning off both PFD/DAC element and timing mismatch does not alter spurious performance. This is also an intuitive result, since, after the application of dynamic element matching techniques, both represent broad-band noise sources.

The goal of the digital compensation scheme is to reduce fractional spurs, but a few words are in order regarding the reference spur. Given the high bandwidth of the synthesizer and relatively low reference frequency, suppressing the reference spur is made more difficult than in prior work employing much lower bandwidths. In the example synthesizer discussed here, the reference spur at 50 MHz measures -60 dBc. In a target application such as GSM, the band select filter bandwidth is 20 MHz, and the reference spur would be further attenuated. Additionally, it is possible to use additional higher order poles to further suppress the reference spur, as is common practice in the literature. Finally, as previously discussed, sample-and-hold techniques similar to those used in classical fractional- N synthesis can be employed to reduce reference feed-through, and to potentially eliminate the shape mismatch error between vertical and horizontal slicing.

Focusing now on fractional spurs, Fig. 13 shows a simulation result for a particular input with and without the digital gain compensation enabled. As mentioned previously, the compensation LUT has finite input and output resolution. This is represented in the simulation results by denoting a synthesizer as being X/Y compensated, where X is the number of address bit to the LUT, and Y is the number of output bits. With 6/4 compensation, representing a 1 Kb LUT, the fractional spur is reduced by >15 dB for this example!

In order to examine rejection over a broad range of fractional spur values, detailed behavioral level simulations are performed over a wide range of accumulator input values. The methodology used, as well as simulation results, is depicted in Fig. 14. The upper trace of Fig. 14(a) shows the output spectrum for




Fig. 13. Example of digital compensation for improved spurious performance. (a) Compensated versus uncompensated spectrum and (b) zoom-in of (a).




Fig. 14. Spurious performance methodology and simulation results. (a) Spurious performance methodology. (b) Spurious performance simulation results.

a particular accumulator input. A tone-detection algorithm is used to detect any spurs in the spectrum. Simulations are run for 28 accumulator input codes, and the worst-case spurs at each frequency are determined across simulation runs. The bottom left trace of Fig. 14(a) shows how, once worst-case tones are determined, an envelope is used to represent the maximum spur levels. Fig. 14(b) shows results from a number of simulation runs with various levels of quantization in the digital compensation LUT. With a 20-b accumulator, it is extremely time intensive to simulate the entire 20-b input space so, to generate the plot, 28 simulations were done for each compensation level, with the accumulator input varied so that the fractional spurs generated would span a 20-MHz bandwidth. The steps in the plot are due to stepping the accumulator input through a set of values chosen to generate fractional spurs across a range of offset frequencies. The worst-case tones without any digital compensation are < -74 dBc and occur at very low offset frequencies. In order to more closely examine the low-frequency spur performance,

an additional 186 finely-spaced simulations are done with the accumulator input set to produce low-frequency fractional spurs, as shown in Fig. 15. (The envelope floor is set at -80 dBc for the uncompensated synthesizer and -100 dBc for the compensated synthesizer in Fig. 15). The figure shows that, once the 1-MHz bandwidth is exceeded, filtering by the PLL dynamics helps reduce spur feed-through.

Returning to the results presented in the right plot of Fig. 14, the -74 dBc raw performance is seen to be very good, and may be attributed to the enhanced gain match obtained by the PFD/DAC as compared to prior art. Once compensation is enabled, the spur performance improves dramatically. "Ideal compensation" means that the resolution on the LUT matches that of the input accumulator, and serves as a limiting case. The synthesizer simulated is assumed to have a 20-b input accumulator, corresponding to 20/20 compensation for the ideal case. With 10/10 compensation near ideal results are obtainable. However, 10/10 compensation requires a 1 Mb LUT, which is rather large. By contrast, with 6/4 compensation and 1-kb LUT, all tones are




Fig. 15. Close-in spur performance.




Fig. 16. Maximum spur levels for various levels of compensation.

kept below -92 dBc, an 18 dB improvement from the uncompensated case. Fig. 14 and 15 demonstrate that using the compensation scheme results in improved performance across frequencies.

Finally, Fig. 16 summarizes the envelope results presented in the right plot of Fig. 14. Maximum fractional spur level is contrasted between the uncompensated synthesizer (dashed line) and compensated synthesizer for various combinations of LUT input and output resolution. Best possible performance is achieved with a very high resolution LUT and corresponds to 22-dB improvement and a < -95 dBc maximum spur. The 6/4 compensation level used in the example synthesizer results in 18-dB improvement and < -92 dBc maximum spur levels.

IV. CONCLUSION

Techniques for reducing broad-band quantization phase noise in fractional- N synthesis while achieving high levels of both fractional spur rejection and mismatch error tolerance have been presented. The proposed PFD/DAC structure incorporates several dynamic element matching techniques to greatly reduce the

impact of both element mismatch and timing skew on synthesizer phase noise performance. For the example synthesizer presented, 36 dB reduction in quantization-induced phase noise is demonstrated via behavioral simulation results. 40 dB or more reduction may be possible, depending on the achievable levels of matching in the PFD/DAC. An analysis of a frequency dependent gain error that occurs in any attempt at horizontal cancellation of the fractional spur has lead to a low overhead, all digital solution capable of achieving fractional spur levels of < -92 dBc. Although this paper presents a simulation based evaluation of the proposed technique, a great deal of attention has been paid to issues regarding actual circuit implementation. In addition to mismatch, bias effects on linearity for possible realizations of the PFD/DAC are addressed through the introduction of an extended range PFD/DAC structure. An integrated circuit is currently being designed in order to compare measured results with the simulations and analyzes presented.

REFERENCES

- [1] U. L. Rohde, *Digital PLL Frequency Synthesizers, Theory and Design*. Englewood Cliffs, NJ: Prentice-Hall, 1983.
- [2] T. A. Riley, M. A. Copeland, and T. A. Kwasniewski, "Delta-Sigma modulation in fractional- N frequency synthesis," *IEEE J. Solid-State Circuits*, vol. 28, pp. 553–559, May 1993.
- [3] B. Miller and R. J. Conley, "A multiple modulator fractional divider," *IEEE Trans. Instrum. Meas.*, vol. 40, pp. 578–583, June 1991.
- [4] I. Galton, "Delta-Sigma data conversion in wireless transceivers," *IEEE Trans. Microwave Theory Tech.*, vol. 50, pp. 302–315, Jan. 2002.
- [5] M. Perrott, T. Tewksbury, and C. Sodini, "A 27 mW fractional- N synthesizer using digital compensation for 2.5 Mb/s GFSK modulation," *IEEE J. Solid-State Circuits*, vol. 44, pp. 2048–2059, Dec. 1997.
- [6] B. De Muer and M. Steyaert, "A CMOS monolithic $\Sigma\Delta$ -controlled fractional- N frequency synthesizer for DCS-1800," *IEEE J. Solid-State Circuits*, vol. 37, pp. 835–844, July 2002.
- [7] W. Rhee, B. Song, and A. Ali, "A 1.1 GHz CMOS fractional- N frequency synthesizer with 3-b third-order $\Sigma\Delta$ modulator," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1453–1460, Oct. 2000.
- [8] C. Park, O. Kim, and B. Kim, "A 1.8 GHz self-calibrated phase-locked loop with precise I/Q matching," *IEEE J. Solid-State Circuits*, vol. 36, pp. 777–783, May 2001.
- [9] K. Lee *et al.*, "A single chip 2.4 GHz direct-conversion CMOS receiver for wireless local loop using multiphase reduced frequency conversion technique," *IEEE J. Solid-State Circuits*, vol. 36, pp. 800–809, May 2001.
- [10] W. Rhee and A. Ali, "An on-chip compensation technique in fractional- N frequency synthesis," in *Proc. IEEE ISCAS*, vol. 3, July 1999, pp. 363–366.
- [11] T. A. Riley and J. Kostamovaara, "A hybrid $\Sigma\Delta$ fractional- N frequency synthesizer," *IEEE Trans. Circuits Syst. II*, vol. 50, pp. 176–180, Apr. 2003.
- [12] S. Pamarti and I. Galton, "Phase noise cancellation design tradeoffs in delta-sigma fractional- N PLLs," *IEEE Trans. Circuits Syst. II*, pp. 829–838, Nov. 2003.
- [13] Y. Dufour, "Method and Apparatus for Performing Fractional Division Charge Compensation in a Frequency Synthesizer," U.S. Patent 6 130 561, 2000.
- [14] S. Norsworthy, R. Schreier, and G. Temes, *Delta-Sigma Data Converters: Theory, Design, and Simulation*. Piscataway, NJ: IEEE Press, 1997.
- [15] C. Y. Lau and M. H. Perrott, "Fractional- N frequency synthesizer design at the transfer function level using a direct closed loop realization algorithm," in *Proc. IEEE Design Automation Conf.*, June 2003, pp. 526–531.
- [16] M. Perrott, "Fast and accurate behavioral simulation of fractional- N frequency synthesizers and other PLL/DLL circuits," in *Proc. IEEE 39th Annu. Design Automation Conf.*, 2002, pp. 498–503.
- [17] M. Perrott, M. Trott, and C. Sodini, "A modeling approach for $\Sigma\Delta$ fractional- N frequency synthesizers allowing straightforward noise analysis," *IEEE J. Solid-State Circuits*, vol. 37, pp. 1028–1038, Aug. 2002.



Scott E. Meninger received the B.S. degree in electrical engineering and computer science from Boston University (BU), Boston, MA, in 1996 and the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 1999. He is currently working toward the Ph.D. degree at MIT.

From 1996 to 1997, he designed clean-room test equipment for Vicor Corporation, Andover, MA. From 1999 to 2001, he was with National Semiconductor's East Coast Labs Design Center,

Salem, NH, as a Staff Designer working on high-accuracy digital temperature sensors. His research interests include low-power and high-speed mixed signal design, with an emphasis on utilizing digital signal processing techniques to ease system design.

Mr. Meninger held a Trustee Scholarship at BU from 1992 to 1996. He currently holds a Teradyne Fellowship at MIT.



Michael H. Perrott received the B.S. degree in electrical engineering from New Mexico State University, Las Cruces, in 1988, and the M.S. and Ph.D. degrees in electrical engineering and computer science from Massachusetts Institute of Technology (MIT), Cambridge, in 1992 and 1997, respectively.

From 1997 to 1998, he was with Hewlett-Packard Laboratories, Palo Alto, CA, working on high-speed circuit techniques for Sigma-Delta frequency synthesizers. In 1999, he was a Visiting Assistant Professor at the Hong Kong University of Science and Tech-

nology, where he taught a course on the theory and implementation of frequency synthesizers. From 1999 to 2001, he worked at Silicon Laboratories, Austin, TX, and developed circuit and signal processing techniques to achieve high-performance clock and data recovery circuits. Currently, he is an Assistant Professor in the Department of Electrical Engineering and Computer Science, MIT, where he focuses on high-speed circuit and signal processing techniques for datalinks and wireless applications.