

A 1-MHz Bandwidth 3.6-GHz 0.18- μm CMOS Fractional-N Synthesizer Utilizing a Hybrid PFD/DAC Structure for Reduced Broadband Phase Noise

Scott E. Meninger, *Member, IEEE*, and Michael H. Perrott, *Member, IEEE*

Abstract—A frequency synthesizer architecture capable of simultaneously achieving high closed-loop bandwidth and low output phase noise is presented. The proposed topology uses a mismatch compensated, hybrid phase/frequency detector and digital-to-analog converter (PFD/DAC) circuit to perform active cancellation of fractional-N quantization noise. When compared to a classical second-order $\Sigma\Delta$ synthesizer, the prototype PFD/DAC synthesizer demonstrates >29 dB quantization noise suppression, without calibration, resulting in a fractional-N synthesizer with 1-MHz closed-loop bandwidth and -155 dBc/Hz phase noise at 20-MHz offset for a 3.6-GHz output. An on-chip band select divider allows the synthesizer to be configured as a dual-band (900 MHz/1.8 GHz) direct modulated transmitter capable of transmitting 271-kb/s GMSK data with less than 3 degrees of rms phase error.

Index Terms—Fractional-N, frequency synthesis, noise cancellation, phase noise, quantization noise.

I. INTRODUCTION

FRACTIONAL-N frequency synthesizers provide high speed frequency sources that can be accurately set with very high resolution, which is of high value to many communication systems. Fig. 1 illustrates a classical fractional-N synthesizer, which consists of a phase-frequency detector (PFD), charge pump, loop filter, voltage-controlled oscillator (VCO), and a frequency divider that is dithered between integer values to achieve fractional divide ratios. In this classical case, the dithering is accomplished by the carry output of a digital accumulator whose input corresponds to the fractional portion of the divide value.

Unfortunately, the dithering action of the digital accumulator results in quantization noise with very high spurious content, which in turn corrupts the output noise performance of the synthesizer. The first efforts to correct this issue employed phase

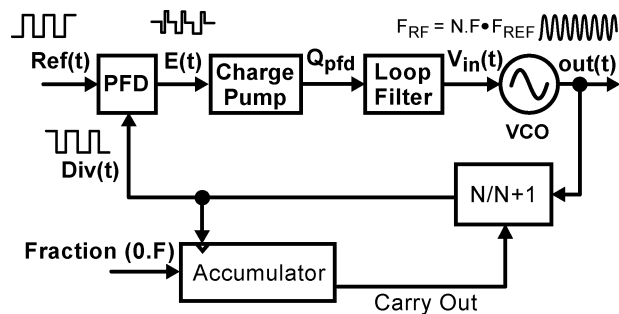


Fig. 1. Classical fractional-N synthesizer

interpolation [1], which was a method that attempted *cancellation* of the quantization noise. Unfortunately, this approach was limited by analog matching requirements, resulting in unacceptable spurious noise levels for many applications. This issue prevented the use of fractional-N frequency synthesizers in most commercial RF applications.

In seeking to mitigate the problem of spurious noise creation in the fractional-N approach, a major breakthrough was made when it was realized that the dithering operation could be done in a more sophisticated manner that largely avoids production of spurious content in the resulting quantization noise [2]–[4]. Riley *et al.* were the first to directly note that the digital accumulator could be viewed as a first-order digital $\Sigma\Delta$ modulator, and that the employment of a higher order $\Sigma\Delta$ modulator to perform dithering leads to substantially reduced spurious content in the quantization noise [4]. In addition, $\Sigma\Delta$ modulation shapes the quantization noise to high-frequency offsets so that it can be effectively *filtered* by the PLL dynamics [4]. These observations opened the door to the development of fractional-N synthesizers for commercial RF applications capable of achieving excellent noise performance.

While employing $\Sigma\Delta$ techniques and filtering the shaped quantization noise has been a key advance for enabling low noise fractional-N synthesis, it comes at the cost of restricted PLL bandwidth. Fig. 2 illustrates the tradeoff between noise performance and PLL bandwidth—as the PLL bandwidth is increased, increased levels of quantization noise pass through the PLL dynamics so that the output phase noise performance is degraded. Higher PLL bandwidth is highly desirable in order to achieve faster settling times and high data rates when directly modulating the frequency/phase of the synthesizer.

Manuscript received September 5, 2005; revised December 19, 2005. This work was supported in part by the MARCO Focus Center for Circuits and Systems Solutions (C2S2, www.c2s2.org), under Contract 2003-CT-888.

S. E. Meninger was with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA. He is now with Cavium Networks, Marlborough, MA 01752 USA (e-mail: meninger@alum.mit.edu).

M. H. Perrott is with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: perrott@mit.edu).

Digital Object Identifier 10.1109/JSSC.2006.870894

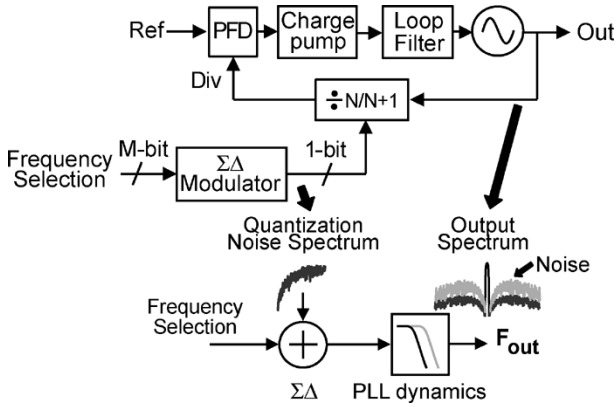


Fig. 2. Classical $\Sigma\Delta$ fractional-N synthesizer showing noise-bandwidth tradeoff.

Rather than simply relying on filtering, recent research efforts have begun to reconsider cancellation as a method to substantially reduce quantization noise [5]–[7]. Cancellation of quantization noise offers the potential of increasing PLL bandwidth while still maintaining excellent noise performance. Recent work has focused on simply adding separate cancellation circuitry to a $\Sigma\Delta$ fractional-N synthesizer which employs a higher order $\Sigma\Delta$ modulator for dithering [6], [7]. While this approach theoretically leverages both noise shaping and cancellation, it has thus far been limited to 16 dB improvement in noise due to imprecise matching of the cancellation circuitry to the error waveform produced by the quantization noise. This mismatch is a direct result of the fact that the cancellation DAC and PFD/charge-pump phase noise signals are summed at the charge-pump output and therefore have separate paths with different gains.

Rather than simply adding separate cancellation circuitry to a $\Sigma\Delta$ fractional-N frequency synthesizer, we have proposed a method of combining the cancellation circuitry with the synthesizer PFD in order to achieve an *inherent* match between the cancellation signal and error waveform. The proposed mismatch compensated PFD/DAC structure thereby leverages a circuit architecture approach to achieve *self-aligned* cancellation of quantization noise [5], [8]. This approach, which is illustrated in Fig. 3, is analogous to MASH $\Sigma\Delta$ modulation in which several low order $\Sigma\Delta$ stages are cascaded, and later stages cancel the noise of previous stages [9]. In particular, instead of using a higher order $\Sigma\Delta$ to perform divider dithering, a simple accumulator (i.e., first-order $\Sigma\Delta$ modulator) is used and its noise is canceled within the PFD/DAC structure. We have shown that this approach achieves measured quantization noise reduction of better than 29 dB with 1-MHz PLL bandwidth, has comparable spurious performance to its rivals, and requires no calibration or tuning of component parameters [10].

The focus of this paper is to present key circuit techniques that enable implementation of the PFD/DAC to achieve substantial reduction (>29 dB) of quantization noise. We begin by providing relevant background of classical fractional-N synthesis and then reviewing the key operating principles of the proposed PFD/DAC approach. With this background in place, we then describe the key circuit issues that must be addressed

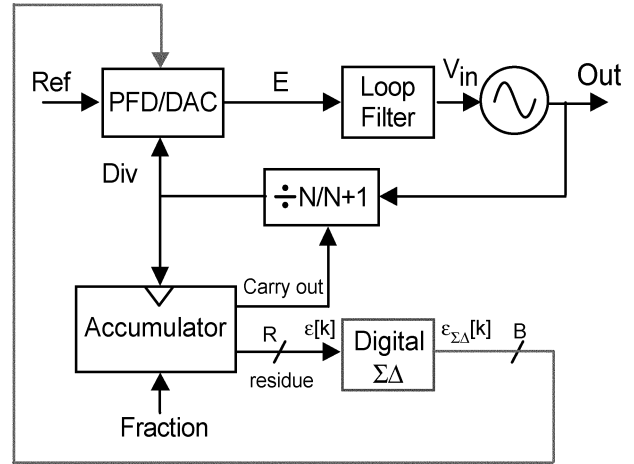


Fig. 3. Proposed mismatch compensated PFD/DAC synthesizer.

to achieve accurate self-alignment of the PFD/DAC structure. We then present enabling circuit approaches to meet the needs of divider retiming, timing mismatch compensation, charge pump noise reduction, and sample-and-hold functionality. The effectiveness of these circuit techniques are then demonstrated through measured results, and the paper ends with relevant conclusions.

II. BACKGROUND

We now review key characteristics of classical fractional-N frequency synthesis in order to provide better understanding of the proposed PFD/DAC approach. This will be done through example by illustrating the waveforms seen at the output of the PFD and charge pump, and then pointing out the noise problems introduced by such waveforms.

Fig. 4 illustrates our choice of PFD structure for the upcoming example, which we refer to as the offset tri-state PFD [11]. This PFD approach offers a highly linear phase detector characteristic by restricting the influence of divider edge variations to down current pulses only [1]. In contrast, the commonly used tri-state PFD approach activates both up and down current pulses as the divider edges vary, which leads to severe nonlinearity in the detector characteristic if the up and down current pulses are not perfectly matched [11]. The key downside of using the offset tri-state PFD is that it must operate with a nonzero steady-state phase error that is large enough to accommodate the full range of divider edge variations. The resulting phase offset reduces the effective phase detector range and also increases the influence of noise by the charge pump. In addition, jitter present in the delay circuit used to create the phase offset, which is labeled as t_{del} in the figure, directly degrades the noise performance of the synthesizer [11]. However, the high linearity achieved by the offset tri-state PFD structure avoids problems of quantization noise folding and spur generation that otherwise severely undermines synthesizer noise performance, and is therefore often worth the costs outlined above.

Assuming that the commonly employed type II PLL topology is used [1], the loop filter will contain an integrator which forces the average charge output by the charge pump to become zero under steady-state conditions. Fig. 4 illustrates the resulting

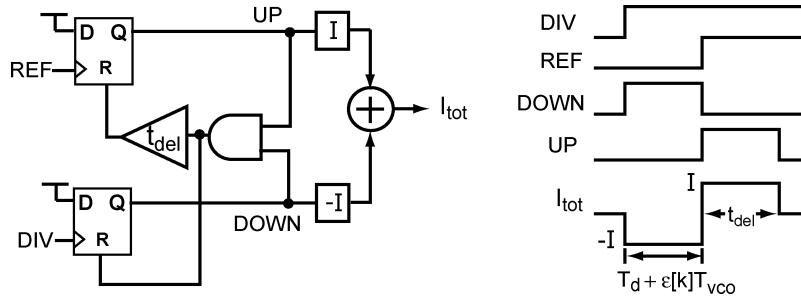


Fig. 4. Offset tri-state PFD and waveforms for fractional-N synthesizer.

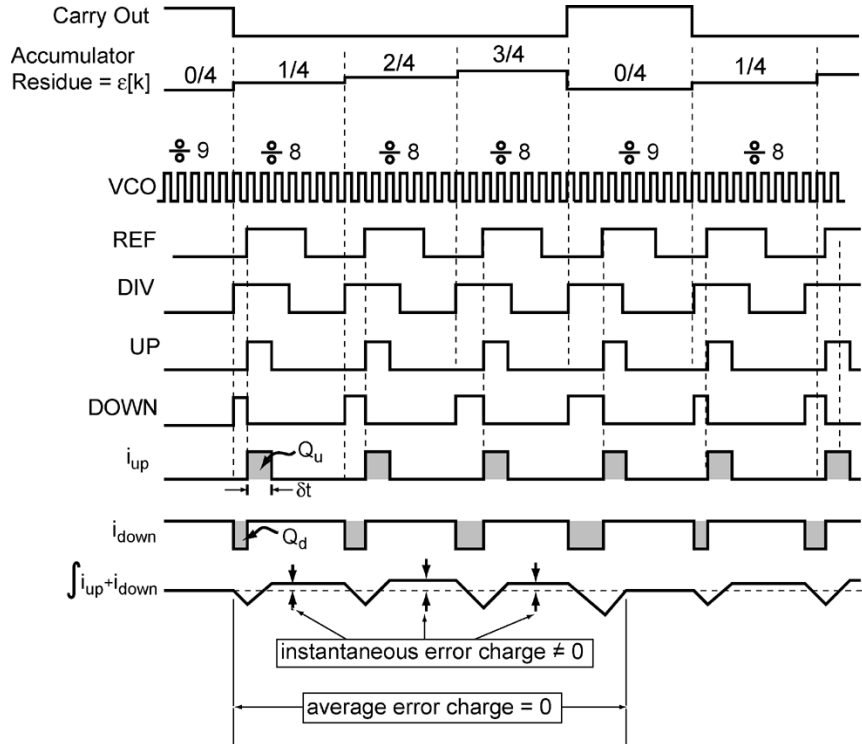


Fig. 5. Classical fractional-N synthesizer example waveforms for $N = 8.25$.

steady-state current waveform produced by the charge pump. We see that the negative pulse varies in width according to the instantaneous divider edge variation as described by the expression $\epsilon[k]T_{vco}$, where $\epsilon[k]$ corresponds to the fractional phase error produced by divider dithering, and T_{vco} corresponds to the period of the VCO output. Note that $\epsilon[k]$ is a discrete-time signal for which $0 \leq \epsilon[k] \leq 1$. In contrast, the up pulse has constant width of t_{del} as set by the delay circuit shown in Fig. 4. In order to achieve zero average charge transfer from the charge pump, the PLL feedback will adjust the average width of the down pulse such that

$$\text{average}(-I \cdot (T_d + \epsilon[k]T_{vco}) + I \cdot t_{del}) = 0 \quad (1)$$

where T_d is defined to be the minimum pulse-width of the down pulse and is constant under steady-state conditions.

Given the above PFD structure, Fig. 5 illustrates the key synthesizer waveforms under steady-state conditions when the divide value is chosen as $N = 8.25$. We see that, while the in-

tegrated current is zero on average, it instantaneously varies in time in a periodic manner according to the fractional value chosen. This periodicity leads to the fractional spur behavior that plagues the classical fractional-N approach. In this case, a fractional value of $1/4$ leads to the integrated current waveform repeating every four reference periods, so that fractional spurs at $1/4$ the reference frequency are produced. It is straightforward to show that the fractional spurs will change in frequency according to the fractional value chosen, which makes their cancellation particularly challenging.

III. PROPOSED PFD/DAC STRUCTURE

Given the above background on classical fractional-N frequency synthesis, we now review the key characteristics of the proposed PFD/DAC structure that performs self-aligned quantization noise cancellation. We do so through modification of the previous example.

Fig. 6 illustrates the proposed PFD/DAC structure, which consists of an offset tri-state PFD with two divider inputs and

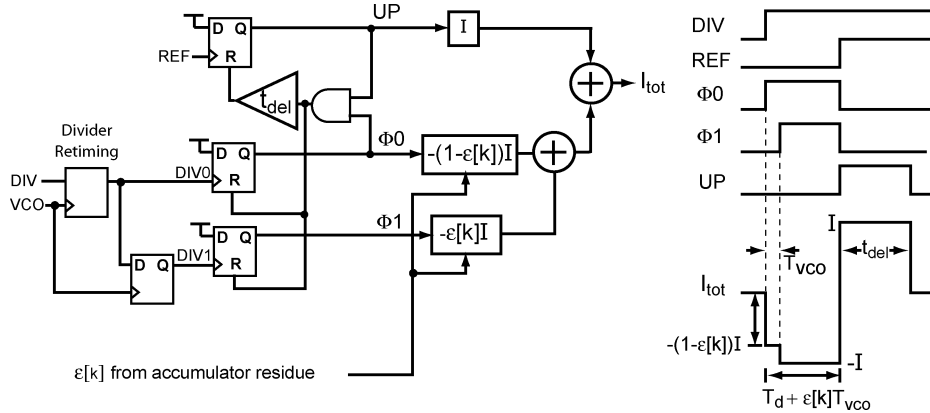


Fig. 6. Offset tri-state PFD and waveforms for PFD/DAC synthesizer.

one reference input. The two divider signals correspond to a retimed version of the divider output, Div0, and a one VCO cycle delayed version of the retimed divider output, Div1. The two PFD outputs feed into two charge pumps that implement the DAC portion of the structure. To achieve the DAC functionality, the two charge pumps associated with the divider signals are implemented as a shared DAC current source which is controlled with the predicted value of the fractional phase error, $\epsilon[k]$. The value of $\epsilon[k]$ is taken directly from the residue of the accumulator used to dither the divide value, as is commonly done in classical phase interpolation approaches [1].

Fig. 6 illustrates the resulting steady-state current waveforms produced by the charge pump in the PFD/DAC case. In contrast to the classical fractional-N approach, the charge pump waveform now has an extra degree of freedom to adjust the area of each individual down pulse. In particular, the initial portion of the down pulse, which is active for one VCO period in time (i.e., T_{VCO}), takes on the value $-(1+\epsilon[k])I$ rather than $-I$. To achieve an overall area of zero for the combined up and down pulses, we must satisfy the equation

$$-I \cdot (T_d + \epsilon[k]T_{VCO} - T_{VCO}) - I \cdot (1 - \epsilon[k])T_{VCO} + I \cdot t_{del} = 0. \quad (2)$$

In fact, a bit of algebraic manipulation of the above equation reveals that it reduces to

$$-I \cdot T_d + I \cdot t_{del} = 0. \quad (3)$$

Under steady-state conditions, T_d will be automatically adjusted by the PLL feedback dynamics to achieve the above relationship.

The last equation reveals that the combined area of the up and down pulses is zero *every* cycle, which is in contrast to the classical fractional-N case in which it could only be made zero on *average*. Said a different way, *as fractional-N dithering changes the pulse-width of the down current, the amount of charge delivered in a one VCO period wide window is varied to compensate, such that charge balance is maintained every period*. A key characteristic of the proposed PFD/DAC circuit is that it achieves this charge balance in a self-aligned manner such

that no calibration is required—we will discuss this point further below.

Fig. 7 illustrates the key synthesizer waveforms under steady-state conditions when the divide value is chosen as $N = 8.25$ and the PFD/DAC structure is employed. We see that the PFD/DAC operates to maintain equal and opposite charge for each of the up and down current pulses so that zero net charge is transferred to the loop filter and VCO input. However, a closer inspection of the figure reveals that the *shape* of the integrated up/down current pulses varies periodically in time. To prevent the time-varying shape from inducing fractional spurs, the charge pump can be followed by a sample-and-hold circuit that gates the charge to the loop filter such that it is only transferred during the off times of the up/down waveforms. The combination of the PFD/DAC and sample-and-hold provides for a high degree of suppression of fractional-N spurs. We examine the sample-and-hold circuit in detail later in this paper.

The ability of the PFD/DAC technique to completely cancel quantization noise is primarily determined by the accuracy of generating the one VCO period wide, variable current window at the beginning of the down pulse shown in Fig. 6. We will refer to this variable current window as a charge-box since the integral of current over time is charge. In order to generate an accurate charge-box, we need to set its time duration to *precisely* one VCO period, and we need set the current magnitude to *precisely* $I \cdot (1 - \epsilon[k])$. In [5], we proposed architectural tradeoffs that can be made to reduce the impact of mismatch on PFD/DAC performance. The following section proposes circuit techniques to achieve such accuracy in a self-aligned manner such that no calibration is required.

IV. CIRCUIT IMPLEMENTATION

Fig. 8 displays the key challenges in achieving an accurate charge-box for quantization noise cancellation within the PFD/DAC structure. Ideally, the charge-box has a time duration of exactly one VCO period, and is broken up into equally spaced current increments of $I/2^B$, where a B -bit PFD/DAC is desired. In reality, the circuits generating the charge box will be prone to both *timing* mismatch and *magnitude* mismatch. Timing mismatch corresponds to having a time duration for the charge-box that is different from one VCO cycle by a time offset of Δt , and is caused by mismatches between the registers




Fig. 7. PFD/DAC synthesizer example waveforms for $N = 8.25$.




Fig. 8. Charge-box mismatch.

and routing paths associated with signals Div0 and Div1 in Fig. 6. Magnitude mismatch corresponds to unequal division of current values within the charge box, and is caused by mismatch between the DAC current elements used to implement the scaling operations $(1 - \epsilon[k]) \cdot I$ and $\epsilon[k] \cdot I$ in Fig. 6. A previous PFD/DAC proposal suggested in [8] suffers from such issues, and is therefore limited in its ability to achieve high levels of quantization noise cancellation.

To counteract timing and magnitude mismatch, we propose the mismatch compensated PFD/DAC architecture depicted in Fig. 9. Here, the PFD/DAC proposal in [8] is augmented with several circuits that leverage dynamic element matching techniques to substantially eliminate mismatch. A key aspect of such techniques is that they do not depend on calibration to be effective—rather, the charge-box is effectively self-aligned *on average* with respect to achieving a time duration of one VCO

cycle and having equally spaced current increments. We describe these circuit techniques in the following four subsections.

The first two subsections focus on proposed circuit approaches to counteract timing mismatch. Referring to Fig. 9, the first subsection describes a divider retiming circuit that aligns the Div0 signal to the VCO output—the challenge in doing so is to reliably avoid meta-stability while using an asynchronous divider circuit. Given that Div0 is aligned to the VCO output, we can delay this signal by one VCO cycle (and thus create Div1) by simply passing it through another register clocked by the VCO. However, mismatches in the registers and layout paths associated with these signals leads to residual timing mismatch (Δt). The second section describes the timing mismatch compensation circuit indicated in Fig. 9—this circuit leverages dynamic element matching to mitigate such mismatch by effectively swapping the paths that Div0 and Div1 take through the following PFD and charge pump circuits.

The third subsection focuses on proposed circuit approaches to counteract magnitude mismatch and achieve high speed operation and low noise performance for the charge pump. As indicated in Fig. 9, a DAC mismatch shaping algorithm is used to scramble the mapping between the $\Sigma\Delta$ residue and DAC current elements such that mismatch-induced noise is shaped to high frequencies. The scrambling action is important in order to avoid noise folding of the first-order shaped quantization noise generated by the digital $\Sigma\Delta$ modulator shown in Fig. 3, which is used to increase the effective resolution of the PFD/DAC [5]. The charge pump speed and noise performance are improved through several circuit techniques, including $1/f$ noise reduction through resistor degeneration [12].

The final subsection focuses on implementation of the sample-and-hold circuit that is used to eliminate the influence of the charge pump current pulses on the loop filter and VCO. As discussed in Section III, this technique is a simple and

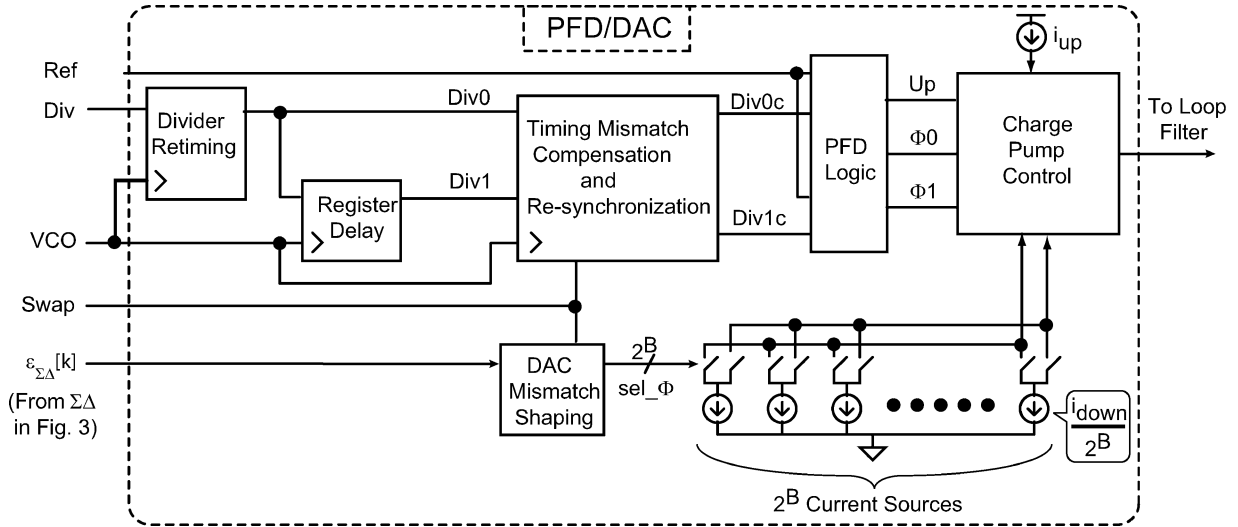


Fig. 9. Mismatch compensated PFD/DAC.

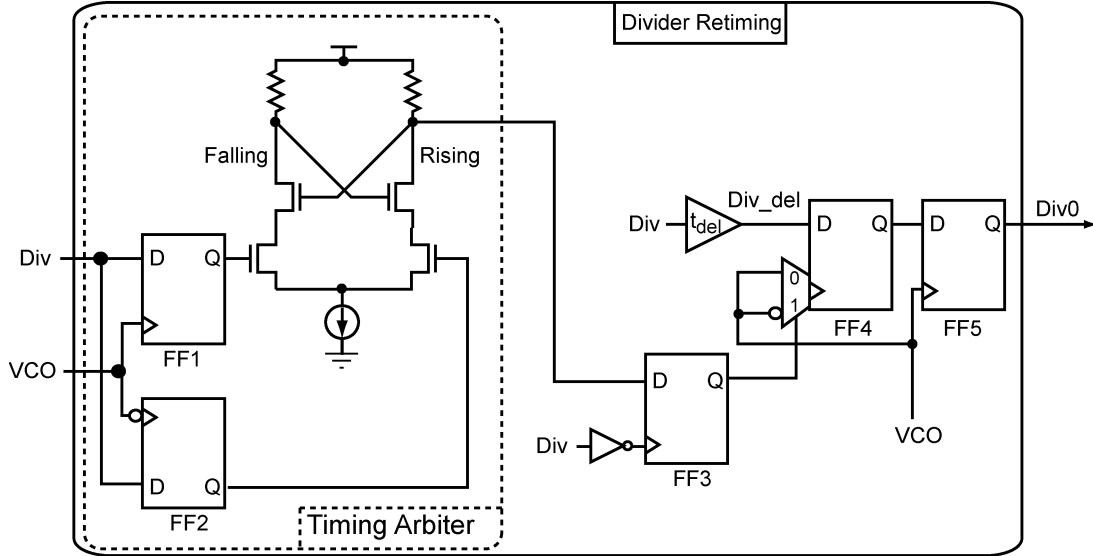


Fig. 10. Divider retimer.

effective method of eliminating fractional spurs due to periodic changes in the *shape* of the current pulses, and of also substantially reducing the magnitude of the reference spur.

A. Divider Retimer

Multi-modulus dividers used in high-speed synthesizer designs are typically asynchronous in nature in order to minimize power dissipation [13]. This creates a situation where the relative phase between the VCO and divider output varies substantially as a function of process, temperature, and divide value variations. If one attempts to directly retime the divider output with a register clocked by the VCO output, meta-stability problems can occur if the register setup or hold times are violated. Since a retimed divider output is required to achieve accurate timing duration in the PFD/DAC generated charge-box, this issue must be addressed to achieve robust operation of the proposed PFD/DAC.

In order to avoid such meta-stability problems, we propose the divider retimer circuit depicted in Fig. 10. Here we use a timing arbiter circuit to select *either* the rising or falling edges of the VCO output as the retiming signal according to which one minimizes the probability of incurring meta-stability in the re-timing register FF4. In contrast to the divider retiming approaches presented in [6], [14], which require information to be propagated through the various stages of the asynchronous divider, the circuit in Fig. 10 directly determines the likelihood of a meta-stable event and re-times accordingly.

Due to the high speed of the VCO signal, all logic in the divider retimer in our system was implemented using resistively loaded, source coupled logic (SCL) (which is also known as current mode logic (CML)). When the circuit is active, a high-speed differential arbiter evaluates whether a rising edge (FF1) or falling edge (FF2) triggered flip-flop generates a valid output level first. A simple, low-speed FSM (FF3) controls a retiming flip-flop (FF4) whose input is a delayed version of the divider

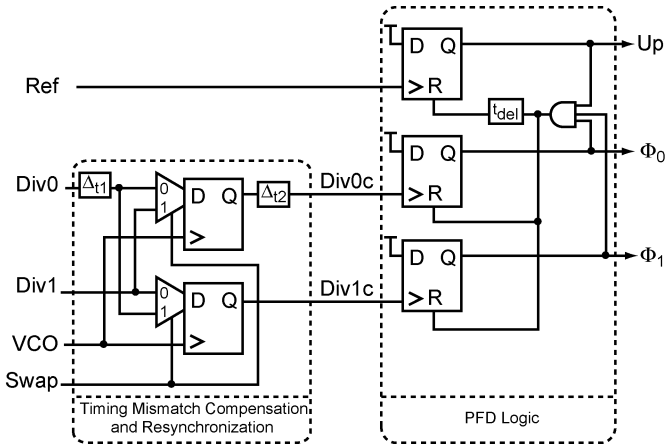


Fig. 11. Timing compensation and PFD logic.

output. The delay is designed to be slightly more than a setup-and-hold time to give some margin when retiming. The retiming flip-flop (FF4) is clocked by the opposite of the VCO phase output from the arbiter. The combination of delaying the divider and controlling the retiming with the opposite VCO phase at FF4 ensures valid retiming for the different possible meta-stable conditions at the arbiter. A detailed, phase-space explanation of the retiming circuit can be found in [11]. The retimer output is taken from FF5, which is always clocked on the VCO rising edge. In this way we establish that the retimer output, and therefore the charge-box, is always referred to the VCO rising edge.

We propose locking the selection of the VCO edge used by FF4 during times in which the noise performance of the synthesizer is critical. By doing so, we avoid possible variation in the retimer output that could occur if the timing arbiter is operating at the edge of a given selection region. This strategy fits in well with burst-mode communication schemes in which the synthesizer output is not required continuously. The time delay t_{del} is designed to be long enough to accommodate environmental drift during times that the choice of VCO edge used by FF4 is locked.

B. Timing Mismatch Compensation and PFD Logic

The time duration of the PFD/DAC charge-box is set by the timing difference between the divider retimer output, Div0, and the one VCO period delayed signal, Div1. Any mismatch in delay between these two paths corrupts the charge box, and results in incomplete quantization noise suppression.

Fig. 11 displays our proposed timing mismatch compensation circuit to reduce the impact of such timing mismatch. This circuit achieves matching delay paths for Div0 and Div1 by dynamically swapping their routing through the PFD logic in a pseudo-random fashion such that they both use each path for the same amount of time *on average*. Note that the residue error processed by the DAC current sources must be swapped as well, as shown in Fig. 9. This approach achieves a precisely set *average* time duration for the PFD/DAC charge box of one VCO cycle *without calibration*.

The cost of using the swapping approach is that noise will be generated as the PFD/DAC charge box *instantaneously* varies in time duration according to which swap path is chosen. The level

of this timing-mismatch induced noise depends on the magnitude of the timing mismatch. In Fig. 11, we separate timing mismatch into two components, Δ_{t1} and Δ_{t2} , which correspond to mismatch delay before and after the swapping registers, respectively. Since these registers align their outputs to their clock input, the impact of Δ_{t1} is negligible. However, Δ_{t2} directly has influence on the charge box time duration, and therefore must be minimized in order to reduce the amount of noise generated in the swapping process.

Pseudo-random control of the *Swap* signal results in a white power spectral density (PSD) for the timing-mismatch induced noise. In this design, we use a 23 register LFSR to produce a randomized *Swap* signal that has an average duty cycle of 0.5. The impact of this noise on overall synthesizer phase noise performance can be calculated based on known PLL parameters and an estimate of the residual time mismatch Δ_{t2} [11]. In future explorations of the PFD/DAC synthesizer, it would be desirable to develop a phase swapping technique that shapes the mismatch noise to minimize its in-band impact.

We now provide more details on the circuitry of the timing compensation and PFD logic block. SCL logic is used in order to generate fast edge rates, which is key for establishing a well defined charge-box at high frequencies (3.6 GHz for our prototype). As shown in Fig. 12, the phase swapping muxes shown in Fig. 11 are directly embedded in the flip-flop input latch stages in order to save power and area and to increase speed. In many designs, only transistors M1 and M4 are used in the first latch stage if an input multiplexer function is desired. However, without transistors M2, M3, M5, M6, there will be a change in loading at the internal nodes *latch* and *latchb* as the de-selected phase input signals switch. Transistor pairs M2,M3 and M5,M6 isolate the first stage latch nodes from the de-selected input pair, and eliminate this input state dependent mismatch, reducing residual timing error Δ_{t2} . Measurements show that Δ_{t2} ultimately limits low-frequency PFD/DAC output phase noise, so any means of reducing its magnitude is key to achieving best performance [11].

C. Magnitude Mismatch Compensation and Charge-Pump Unit Element Design

Thus far, we have proposed techniques to mitigate timing mismatch. Now, we focus on magnitude mismatch caused by variations between the DAC unit elements that create the variable current levels within the charge-box. We also propose several techniques to achieve high speed operation and low noise performance for the charge pump.

Unit element current mismatch results in nonlinearity in the DAC characteristic, which acts to fold quantization noise produced by the $\Sigma\Delta$ modulator shown in Fig. 3 and induce fractional spurs due to the periodic components of the accumulator residue present in its output, $\epsilon_{\Sigma\Delta}[k]$. To counteract such issues, the DAC mismatch shaping block shown in Fig. 9 scrambles the mapping from the DAC input to the DAC unit elements such that the DAC nonlinearity is removed on average and the mismatch noise is scrambled (to remove the fractional spurs) and shaped to high frequencies (to reduce its impact on in-band PLL noise). This block consists of a thermometer decoder and data weighted averager (DWA) [15], which performs a modified barrel shift of

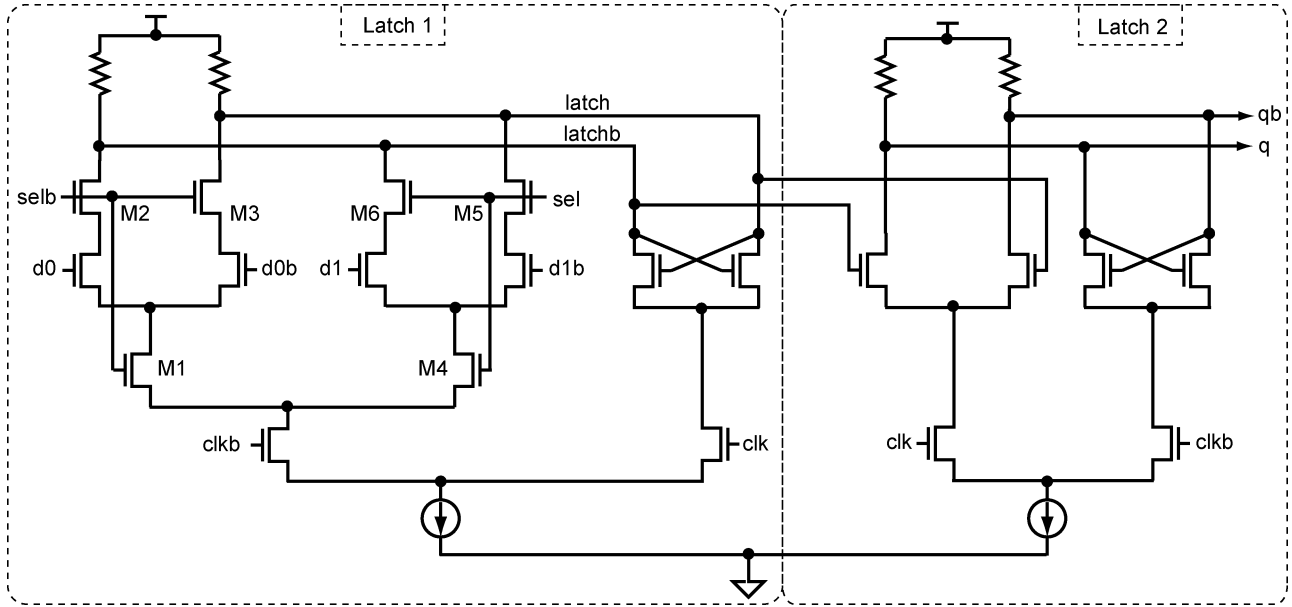


Fig. 12. High-speed differential flip-flop with mixed input stage and state mismatch compensation.




Fig. 13. Charge-pump unit element.

the unit elements as they are utilized by the system over successive periods.

The charge-pump unit element depicted by Fig. 13, which should be compared against Fig. 9 to properly understand the context of its role, uses several techniques to improve performance. A differential structure is used to enable fast switching. Phase select signals sel_{Φ_0} and sel_{Φ_1} are used to choose which PFD output, Φ_0 or Φ_1 , is used to direct the current of the given charge pump cell. Sharing of the same charge pump circuitry controlled by the Φ_0 and Φ_1 PFD outputs results in *intrinsic* matching between them.

One should note that the NMOS-based current sources are cascoded, while the PMOS-based current sources are not. Based on detailed behavioral simulations, it was determined that the NMOS-based current sources should be cascoded in order to maximize their output impedance. These current sources control the variable current within the charge-box, and their performance is critical to achieve accurate noise cancellation. However, the PMOS-based current sources simply need to provide a constant charge packet each reference period as set by the reference edge and delay cell within the PFD, so that their output impedance is not critical. Therefore, the pMOS current sources are left uncascoded, sacrificing output impedance for improved headroom.

Resistive degeneration is employed by both pMOS and nMOS current sources in order to reduce the impact of charge-pump noise, particularly $1/f$ noise [12]. MOS drain current noise is lowered according to

$$\overline{i_{ntot}^2} = \left| \frac{1}{1 + gm_1 R_n + \frac{R_n}{r_{o1}}} \right|^2 \cdot \overline{i_{nd}^2} \quad \text{A}^2/\text{Hz} \quad (4)$$

where $\overline{i_{nd}^2}$ is the current source device's (M1's) drain current noise PSD, r_{o1} is M1's output impedance. Note that the contribution of the degeneration resistor must be accounted for when calculating the total charge-pump noise since it will appear directly at the charge-pump output. In order to accommodate the extra headroom required to support the voltage dropped across the degeneration resistor, the tail nMOS device is sized to be long and wide, whereas the cascode transistor is scaled down to minimize capacitance at the current source output node. Hspice simulations show that, for reasonable values of gm , R_n , and r_{o1} , the circuit can achieve 20–40 dB of noise attenuation [11].

The overall charge-pump output is single ended—the unused side is connected to a “dump node” with controlled voltage,

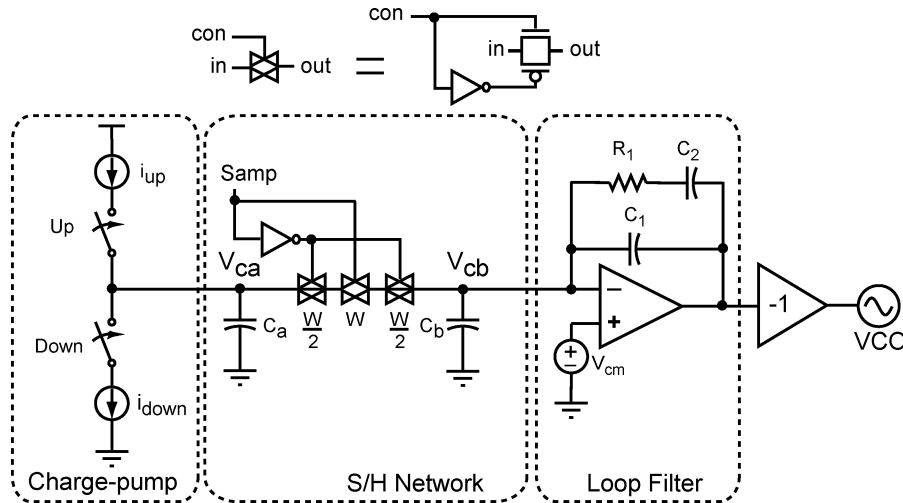


Fig. 14. Sample-and-hold loop filter.

V_{dump} , that is set to the same value as the op-amp common mode voltage, V_{cm} , used to implement the active loop filter depicted in Fig. 14. By keeping these voltages at the same value, the switching performance of the charge pump is improved since voltage transients that could cause a change in unit element output current are minimized.

D. Sample-and-Hold Loop Filter

As mentioned in Section III, the *shape* of the charge-pump output waveform changes periodically during steady-state operation, so that it contains some residual amount of energy at the fractional spur frequency [5]. In addition, there is significant spurious content at the reference frequency. If the charge-pump output is passed directly to the loop filter, residual fractional spurs and a significant reference spur will result.

In the literature, the reference spur is typically reduced by introducing many additional high-frequency poles to the loop filter. Reference spur performance may improve, but at the cost of added complexity and decreased closed-loop stability as more poles are added. Additionally, these poles will not reduce low frequency fractional spurs caused by the PFD/DAC output wave-shape. We therefore propose using the sample-and-hold loop filter depicted in Fig. 14 as a means to improve spurious performance [16].

The sample-and-hold loop filter can be understood by examining the operation of its switch. When the PFD/DAC output currents are active, the sample switch is open and the current sources charge or discharge C_a . When the PFD/DAC completes its operation, *Samp* goes high, and the op-amp summing junction is connected to C_a . Under steady-state PFD/DAC operation, zero net charge is transferred to capacitor C_a from the current sources over each reference period, assuming we ignore noise. By sampling C_a after the PFD/DAC completes its operation, no charge is transferred to the loop filter in steady-state. The VCO sees no disturbance on its control voltage, and *all* spurs can theoretically be eliminated!

Since the op-amp positive terminal is set to V_{cm} Volts, the minus terminal is also nominally at V_{cm} (plus or minus any input offset in the op-amp), and so the nominal voltage at the




Fig. 15. Differential-to-single-ended converter.

charge-pump output is also V_{cm} volts. Note, however, that the voltage at the charge-pump output (node V_{ca}) will make an excursion below V_{cm} during normal operation as indicated in Fig. 7. The output impedance of the charge-pump is therefore a possible concern since i_{up} and i_{down} will both vary from their nominal values according to their output impedance. Capacitor C_a is chosen to be large enough to constrain the voltage swing at the charge-pump output so that current source output impedance does not adversely impact performance. Capacitor C_b does not play a key role in the operation of the sample network, but does act as an intermediate charge-transfer reservoir during transient events when a step in phase error causes the error charge magnitude to exceed the output drive capability of the op-amp.

The sampling operation is performed using complementary transmission gate switches with charge-balancing dummy devices. *Samp* is a full-swing, single-ended signal derived from differential control logic in the PFD/DAC. Because the node voltages V_{ca} and V_{cb} both settle to V_{cm} every period before sampling is performed, the circuit acts as a constant V_{gs} sampling network, thereby minimizing nonlinear effects associated with variable channel resistance in sample switches.

The differential-to-single-ended converter circuit used by the sample-and-hold loop filter is depicted in Fig. 15. Dynamic in nature, it does not dissipate any static power and has the additional benefit that coincident complementary full-swing

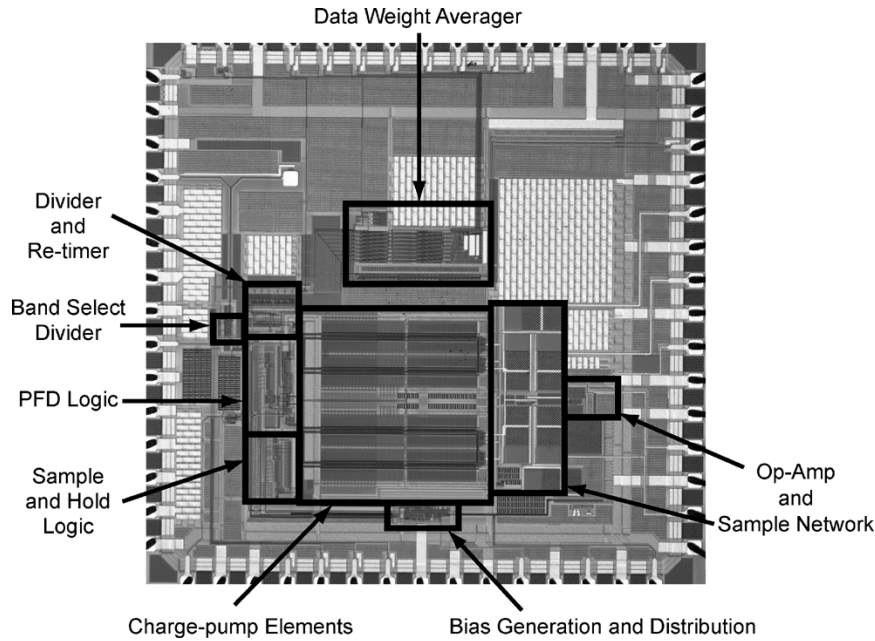


Fig. 16. Chip microphotograph.

output signals are generated. This is important for charge injection cancellation, where an extra inverter delay between *out* and *outb* would cause a phase difference between the overlap charge packets delivered through the nMOS and pMOS devices in the transmission gate. Coincident switching means that device overlap capacitance charge injection is in phase for nMOS and pMOS devices.

While the sample-and-hold loop filter theoretically eliminates both the reference spur and the residual fractional spur associated with the shape of the PFD/DAC output waveform, imperfect cancellation of charge-injection in the sampling network and on-chip coupling limit the amount of suppression obtainable. Measurements of a prototype synthesizer using the sample-and-hold loop filter demonstrate a reduction of the output reference spur from -55 dBc to -74 dBc for a 1-MHz bandwidth type II synthesizer that employs a single additional pole at 2.5 MHz. This 19-dB improvement clearly demonstrates the ability of the sample-and-hold loop filter to dramatically improve spurious performance.

Finally, we mention that circuit design was performed as an iterative process, where Hspice simulation results were used to update a detailed behavioral model. In this way, circuit nonidealities can quickly be evaluated in the context of the overall synthesizer system, which is impractical to simulate for noise performance at the SPICE level. Details of the behavioral modeling techniques are described in [11], and a behavioral model with an associated tutorial for the proposed architecture is freely available at [17] for further exploration.

V. MEASURED RESULTS

The PFD/DAC offers the potential to dramatically reduce fractional-N quantization induced phase noise. In order to determine the practical performance using this approach, a prototype PFD/DAC fractional-N synthesizer was fabricated using National Semiconductor's $0.18\text{-}\mu\text{m}$ CMOS process. Fig. 16 is

a chip microphotograph that has key circuit blocks labeled. The chip is pad limited due to the use of many pads for enhanced programmability. Its overall dimensions are $2.7\text{ mm} \times 2.7\text{ mm}$, with $1.8\text{ mm} \times 1.5\text{ mm}$ active area. The active area is dominated by the 7-bit PFD/DAC charge-pump unit elements. Fig. 17 shows the synthesizer IC system test configuration.

To achieve maximum flexibility in testing, the divider control accumulator and DAC control digital $\Sigma\Delta$ modulator are left off chip. It is important to have digital switching noise represented on-chip so that noise measurements will include realistic amounts of digital-to-analog coupling. The on-chip thermometer decoder and DWA circuitry used by the DAC Mismatch shaping block and the 23 register LFSR used to control the timing mismatch block *Swap* signal consist of significantly more digital gates than the functions implemented on the FPGA, so on-chip digital switching noise is adequately represented.

The synthesizer utilizes a second-order loop filter, consisting of C_1 , C_2 , and R_1 configured for a 1-MHz closed-loop bandwidth. An additional loop filter pole was placed at 2.5 MHz using C_p and R_p . All component values were calculated using the PLL Design Assistant design tool [18]. A ZCOMM 3905 3.6-GHz VCO was employed in the prototype system. The 50-MHz reference used by the system was derived on-chip from an off-chip low-noise 100-MHz OCXO.

Fig. 18 shows the measured noise performance of the unmodulated PFD/DAC synthesizer with calculated and fitted noise subcomponents superimposed. At low-frequency offsets, $1/f$ noise from the reference input dominates—this noise is manifested as jitter in the reference buffer depicted in Fig. 17. At intermediate frequency offsets spanning 100 kHz to 10 MHz, noise generated from swapping out the PFD/DAC timing mismatch (Δt_2) dominates. Based on the measured noise in this range, the value of timing mismatch was determined to be 10.9 ps using an analytical model for the synthesizer and noise variance [11]. At frequency offsets above 10 MHz, VCO noise





Fig. 17. PFD/DAC synthesizer system diagram.

Fig. 18. Measured phase noise and comparison to calculated and fitted noise sub-components ($N = 71.3$).

dominates. The calculated value of the PFD/DAC quantization noise reveals that this noise source is not dominant over any band. Charge-pump noise is not shown explicitly because it is estimated to be significantly lower than the timing mismatch noise magnitude. Overall, we see that this prototype PFD/DAC synthesizer exhibits high closed-loop bandwidth (1 MHz) and excellent phase noise performance (-155 dBc/Hz @ 20-MHz offset for 3.6-GHz output). For more details on the noise fitting and calculations performed for this system, please see [11].

Fig. 19 presents measured results for the unmodulated PFD/DAC synthesizer with all proposed techniques enabled compared to the case where the *same* system (with the same circuit parameters) is configured as a second-order MASH $\Sigma\Delta$ synthesizer. This comparison demonstrates that the PFD/DAC synthesizer achieves >29 dB of quantization noise suppression. This is significantly higher noise suppression than achieved by prior work that employs active noise cancellation techniques [6], [7], which achieve 16-dB and 15-dB suppression, respectively. At lower frequency offsets, the PFD/DAC synthesizer

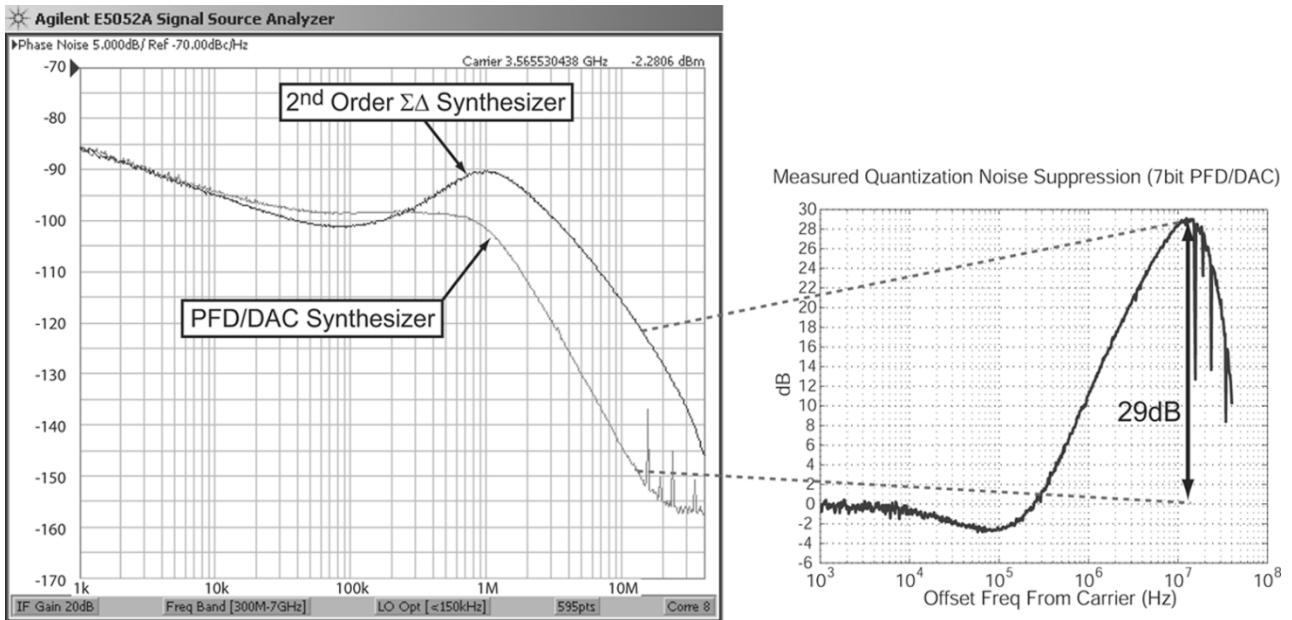

 Fig. 19. 1-MHz bandwidth PFD/DAC synthesizer versus second-order $\Sigma\Delta$ synthesizer ($N = 71.3$).

 TABLE I
 COMPARISON OF SYNTHESIZERS EMPLOYING ACTIVE QUANTIZATION NOISE CANCELLATION
 (ALL ARE IMPLEMENTED IN 0.18- μm CMOS)

	[6]	[7]	This Work
Div Control	2 nd Order $\Sigma\Delta$	3 rd Order $\Sigma\Delta$	1 st Order $\Sigma\Delta$
DAC Control	3 rd Order $\Sigma\Delta$	2 nd Order $\Sigma\Delta$	1 st Order $\Sigma\Delta$
Bandwidth	460kHz	700kHz	1MHz
Output Frequency	2.4GHz	2.1GHz	3.6GHz
Phase Noise @ 10MHz (normalized to 2.1GHz output)	-133dBc/Hz	-135dBc/Hz	-151dBc/Hz
Largest In-band Spur	-45dBc	-55dBc	-45dBc
Noise Suppression	16dB	15dB	29dB
Core Power	61mW	28mW	110mW

suffers from 2-dB increased noise compared to when it is configured as a second-order $\Sigma\Delta$ synthesizer. The increased noise is due to the randomized timing mismatch noise caused by swapping—the second-order $\Sigma\Delta$ synthesizer does not require such swapping. For more detailed measurements that explore performance of each of the proposed noise reduction techniques, the reader is referred to [11].

As Table I shows, the prototype PFD/DAC synthesizer achieves higher bandwidth, greater noise suppression, simpler modulator control, and comparable spurious performance to prior work. Although the current implementation consumes more power than its competitors, we believe that this issue can be largely remedied by lowering the number of bits in the PFD/DAC and using a more recent fabrication process than 0.18- μm CMOS. We discuss such issues in more detail in the following section. Table II summarizes the unmodulated performance of the PFD/DAC synthesizer.

We now examine measured results for the synthesizer modulated by a GMSK data stream. Direct modulation of fractional-N synthesizers offers the advantage over traditional modulation schemes in that mixers and DAC's can be eliminated from the transmit path, with the penalty that data passing

 TABLE II
 SUMMARY OF PFD/DAC SYNTHESIZER/TRANSMITTER PERFORMANCE

Specification	Value
Technology	0.18 μm CMOS (National Semiconductor)
Data Rate	Up to 1Mb/s GMSK Measured (900MHz/1.8GHz Bands)
Reference Frequency	50MHz
Bandwidth	1MHz
Phase Noise @ 100kHz (3.6GHz output)	-98dBc/Hz
Phase Noise @ 20MHz (3.6GHz output)	-155dBc/Hz
Charge-pump Current	6.6mA
PFD/DAC Resolution	7-bit
Largest In-band Spur (3.6GHz output)	-45dBc
Quantization Noise Suppression	>29dB
Core Power	110mW(1.8V)
Digital Power	5.4mW(1.5V)
I/O Buffer Power	66mW(1.8V)
Total Area	2.7mmX2.7mm
Active Area	1.8mmX1.5mm

through the synthesizer is low-pass filtered by the PLL dynamics. The PFD/DAC synthesizer allows substantial increase

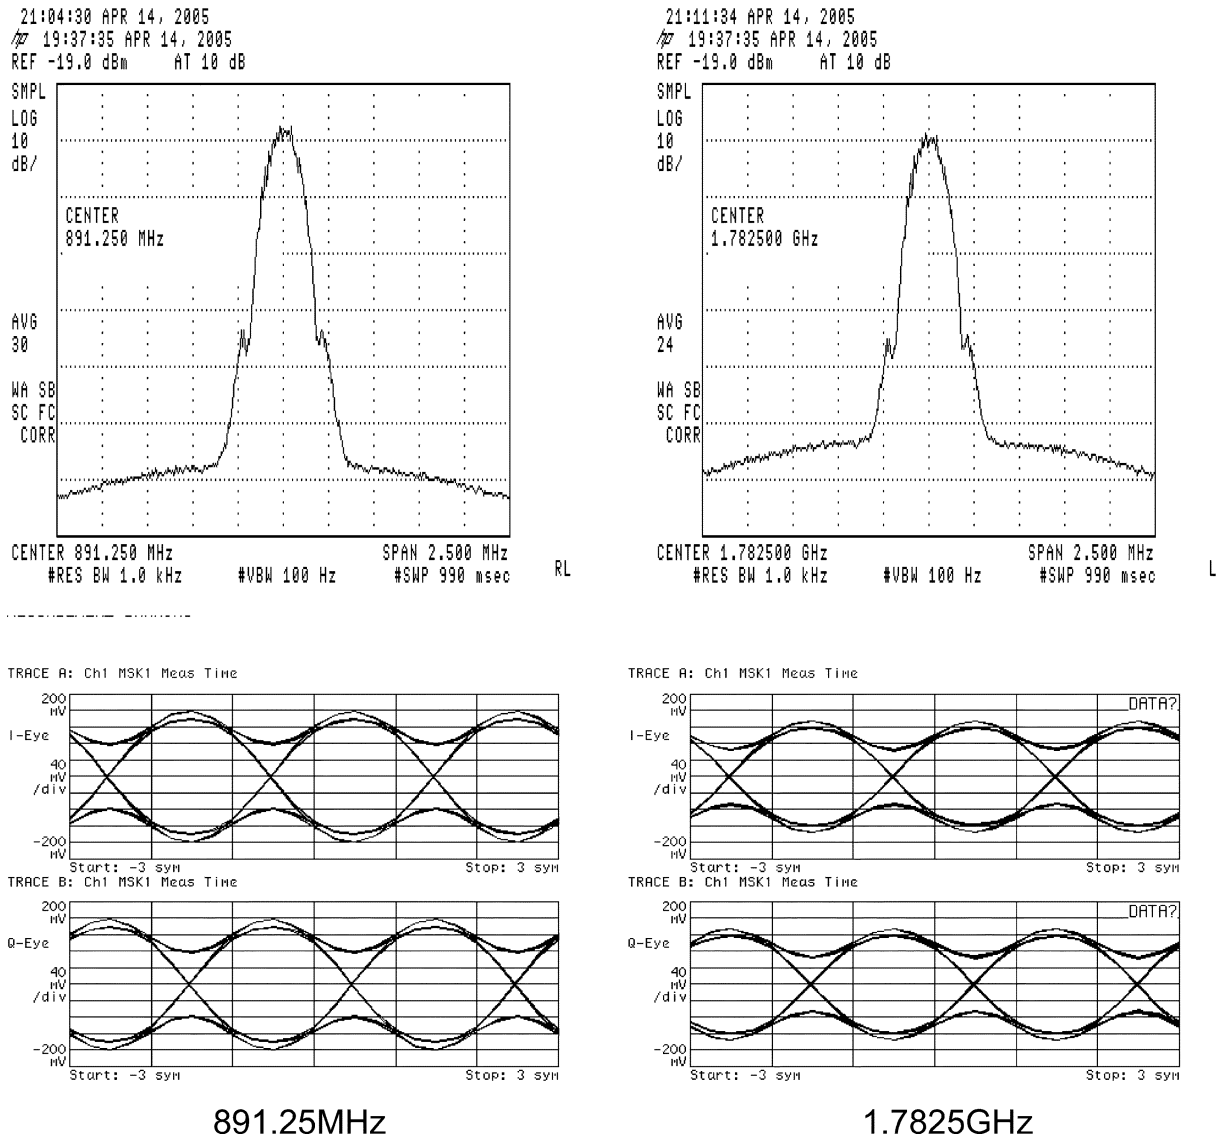


Fig. 20. PFD/DAC synthesizer measured 271 kb/s GMSK spectra and eye diagrams.

of the PLL bandwidth while still achieving low phase noise, thereby allowing higher data rates to be achieved.

To test the prototype system as a direct modulated transmitter, a GMSK filtered random data sequence was generated on a PC and downloaded into a pair of 256-kb FIFO memories. This data was then input to the FPGA containing the $\Sigma\Delta$ modulators used by the PFD/DAC synthesizer. Fig. 20 presents measured results for the PFD/DAC synthesizer modulated by 271-kb/s GMSK data for both 900-MHz and 1.8-GHz output bands. These bands are generated using an on-chip band select divider to divide down the 3.6-GHz VCO output. The demodulated eye diagrams, which were measured on a Hewlett Packard 89440A Vector Signal Analyzer, are wide open and have measured rms phase error magnitudes of 2.8 degrees and 2.7 degrees for the 900-MHz and 1.8-GHz bands, respectively. While we have not targeted a specific transmission standard, we report that the modulated response of Fig. 20 meets the GSM spectral mask at all frequencies except 400 kHz and 600 kHz, where it is 3 dB too high. This is attributed to the timing mismatch noise, which dominates low-frequency synthesizer noise performance. In

Section VI, we discuss ways of reducing timing mismatch in future designs.

Higher data rates with good performance are achievable using the PFD/DAC synthesizer because of its high closed-loop bandwidth. RMS phase error for 500-kb/s GMSK data was measured as 6.2 degrees and 6.5 degrees for the 900-MHz and 1.8-GHz bands, respectively. The increased phase error is attributed to inter-symbol interference (ISI) resulting from filtering of the data by the PLL closed-loop dynamics. As the data rate was increased further, increased levels of ISI were observed. A 757-kb/s data rate results in rms phase errors of 11 degrees and 10.7 degrees for the 900-MHz and 1.8-GHz bands. Measurements for data rates up to 1 Mb/s [11] suggest that even higher data rates with low rms phase errors can be achieved by employing the pre-emphasis filtering technique proposed in [19].

VI. FUTURE IMPROVEMENTS

While significant quantization noise reduction has been demonstrated with the proposed approach, it is worthwhile to

mention potential techniques to lower the power dissipation and further improve the noise performance of future designs.

Focusing initially on lower power dissipation, we suggest lowering the PFD/DAC resolution in order to reduce the driving requirements of the PFD logic, which is the major power draw in the prototype system. In this prototype, we sought 7 bits (128 unit elements) in order to test the limitations of the method, but comparable overall synthesizer noise performance could have been achieved with only 5 bits (32 unit elements). In addition, we suggest utilizing a fabrication process that allows fast enough speed to support implementation of the high speed blocks with full-swing logic rather than the SCL logic employed here. The power savings in such case would be dramatic.

Improvement of noise should be focused on lowering residual fractional spurs, reducing the influence of timing mismatch, and lowering the $1/f$ noise floor at low frequencies. To determine the worst case spur performance, the PFD/DAC synthesizer 3.6-GHz output was varied over a range of 125 channels spaced at 800 kHz increments between 3.56 and 3.66 GHz. The measured channel spacing, when divided down by the band select divider, results in the 200-kHz channel spacing used by 900-MHz band GSM transmitters. In addition, five channels that result in large low-frequency fractional spurs were measured to ensure that in-band spurious performance was adequately measured. These measurements, combined with detailed behavioral simulations, suggest that the worst case measured -45 -dBc in-band spurs observed for the prototype (3.6-GHz output) are caused by on-chip coupling. On-chip power regulators could be used to isolate supply bond-wires from spurious energy generated by on-chip circuitry, as well as to lower on-chip supply voltages of the digital circuits and thereby reduce the energy associated with switching events. It may also be possible to use spread spectrum clocking techniques on the digital logic in order to spread spurious energy into broadband noise. The impact of timing mismatch could be greatly reduced if the swapping method were altered such that noise shaping techniques were applied to push the associated noise to high frequencies. Finally, $1/f$ noise in the current prototype could have been substantially improved by more careful design of the reference buffer. In addition, when using the offset tri-state PFD, it is critical to minimize jitter in its delay cell since it will directly contribute to low frequency noise.

VII. CONCLUSION

In this paper, we have proposed utilizing a mismatch compensated PFD/DAC structure to dramatically reduce the impact of fractional-N dithering noise on synthesizer performance. The proposed architecture utilizes several compensation techniques to alleviate the presence of mismatch sources that would otherwise adversely affect noise cancellation.

Circuits capable of generating an accurate charge-box through the use of dynamic element matching techniques have been presented. Using these techniques, mismatch that would otherwise result in incomplete noise cancellation and large fractional spurs is converted to broadband noise that is then filtered by the PLL. The prototype PFD/DAC synthesizer exhibits

>29 -dB noise suppression when compared to the same system configured as a second-order $\Sigma\Delta$ synthesizer, demonstrating a simultaneous achievement of excellent noise performance and high synthesizer bandwidth, without the use of calibration. Measurements of the prototype system suggest that future designs of the PFD/DAC synthesizer should focus on reducing, and perhaps noise shaping, the timing mismatch between the two phase paths within the PFD/DAC, since this noise source limits noise performance. Measurements further suggest that on-chip coupling limits fractional spur levels to -45 dBc for the 3.6-GHz output, implying that this is a primary issue of concern for future designs. We mentioned a few possible methods of alleviating this problem in the last section.

When configured as a direct modulated GMSK transmitter capable of producing 1.8-GHz or 900-MHz output bands, the system exhibits excellent eye openings and rms phase error at 271 kb/s. The data rate can be raised at the cost of some increased rms phase error, which can be offset by applying some pre-emphasis in the digital filter used to generate the GMSK transmit filter. Both unmodulated and modulated synthesizer data suggest that the PFD/DAC technique is a viable synthesizer architecture for applications which require high bandwidth and low phase noise.

ACKNOWLEDGMENT

The authors would like to acknowledge National Semiconductor (particularly P. Holloway and S. Buddhiraju) for fabrication services, M. Park for help with the DWA circuitry, and the High Speed Circuits and Systems group for technical discussions. They would also like to acknowledge Teradyne for supporting Dr. Meninger as a Teradyne Fellow in 2004.

REFERENCES

- [1] W. Egan, *Frequency Synthesis by Phase Lock*. New York: Wiley Press, 2000.
- [2] J. Wells, "Frequency synthesizer," U.S. Patent 4,609,881, Sep. 2, 1986.
- [3] B. Miller and B. Conley, "A multiple modulator fractional divider," in *Proc. 44th Annu. Symp. Frequency Control*, May 1990, pp. 559–567.
- [4] T. A. D. Riley, M. A. Copeland, and T. A. Kwasniewski, "Delta-sigma modulation in fractional-N frequency synthesis," *IEEE J. Solid-State Circuits*, vol. 28, no. 5, pp. 553–559, May 1993.
- [5] S. E. Meninger and M. H. Perrott, "A fractional-N frequency synthesizer architecture utilizing a mismatch compensated PFD/DAC structure for reduced quantization-induced phase noise," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 50, no. 11, pp. 839–849, Nov. 2003.
- [6] S. Pamarti, L. Jansson, and I. Galton, "A wideband 2.4 GHz delta-sigma fractional-N PLL with 1 Mb/s in-loop modulation," *IEEE J. Solid-State Circuits*, vol. 39, no. 1, pp. 49–62, Jan. 2004.
- [7] E. Temporiti, G. Albasini, I. Bietti, R. Castello, and M. Colombo, "A 700 kHz bandwidth $\Sigma\Delta$ fractional synthesizer with spurs compensation and linearization technique for WCDMA applications," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1446–1454, Sep. 2004.
- [8] Y. Dufour, "Method and apparatus for performing fractional division charge compensation in a frequency synthesizer," U.S. Patent 6,130,561, Oct. 10, 2000.
- [9] S. E. Meninger and M. H. Perrott, "Bandwidth extension of low noise fractional-N synthesizers," in *Proc. IEEE Radio Frequency Integrated Circuits (RFIC) Symp.*, Jun. 2005, pp. 211–214.
- [10] —, "A dual band 1.8 GHz/900 MHz, 750 kb/s GMSK transmitter utilizing a hybrid PFD/DAC structure for reduced broadband phase noise," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 2005, pp. 394–397.

- [11] S. E. Meninger, "Low phase noise, high bandwidth frequency synthesis techniques," Ph.D. dissertation, Mass. Inst. Technol., Cambridge, MA, 2005.
- [12] T. Riley and J. Kostamovaara, "A hybrid $\Delta\Sigma$ fractional-N frequency synthesizer," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 50, no. 4, pp. 176–180, Apr. 2003.
- [13] C. Vaucher, I. Ferencic, K. Locher, S. Sedvallson, U. Voegeli, and Z. Wang, "A family of low-power truly modular programmable dividers in standard 0.35 μm CMOS technology," *IEEE J. Solid-State Circuits*, vol. 35, no. 7, pp. 1035–1045, Jul. 2000.
- [14] H. Chien, T. Lin, B. Ibrahim, L. Zhang, M. Rofougaran, and A. Rofougaran, "A 4 GHz fractional-N synthesizer for IEEE 802.11a," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2004, pp. 46–49.
- [15] R. T. Baird and T. S. Fiez, "Linearity enhancement of multibit $\Delta\Sigma$ A/D and D/A converters using data weighted averaging," *IEEE Trans. Circuits Syst. II, Analog Digit. Signal Process.*, vol. 42, no. 12, pp. 753–762, Dec. 1995.
- [16] J. A. Crawford, *Frequency Synthesizer Handbook*. Norwood, MA: Artech House, 1984.
- [17] S. E. Meninger, Behavioral simulation of a wideband fractional-N synthesizer using the CppSim program. [Online]. Available: <http://www-mtl.mit.edu/researchgroups/perrottgroup/tools.html>
- [18] M. H. Perrott, PLL design using the PLL Design Assistant Program. [Online]. Available: <http://www-mtl.mit.edu/researchgroups/perrottgroup/tools.html>
- [19] M. H. Perrott, T. L. Tewksbury, III, and C. G. Sodini, "A 27-mW fractional-N synthesizer using digital compensation for 2.5-Mb/s GFSK modulation," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 2048–2059, Dec. 1997.



Scott E. Meninger (M'96) received the B.S. degree in electrical engineering and computer science from Boston University, Boston, MA, in 1996, and the S.M. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, in 1999 and 2005, respectively.

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 Laboratories Design Center, Salem, NH, where he designed high-accuracy digital temperature sensors. Since May 2005, he has been with Cavium Networks, Marlborough, MA, where he designs high-performance mixed-signal circuits for use in network microprocessors.



Michael H. Perrott (M'91) 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 the Massachusetts Institute of Technology, 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 Technology, and taught a course on the theory and implementation of frequency synthesizers. From 1999 to 2001, he was with Silicon Laboratories, Austin, TX, and developed circuit and signal processing techniques to achieve high-performance clock and data recovery circuits. He is currently an Assistant Professor in electrical engineering and computer science at the Massachusetts Institute of Technology, and focuses on high-speed circuit and signal processing techniques for data links and wireless applications.