A 2.4-GHz RF Fractional-N Synthesizer With BW $= 0.25 f_{\mathrm{REF}}$

Long Kong, *Member, IEEE*, and Behzad Razavi, *Fellow, IEEE*

***Abstract*— A fractional-N synthesizer architecture incorporates a 35-tap finite impulse response filter that suppresses the $\Sigma\Delta$ noise, but does not affect the loop bandwidth (BW). Employing a three-stage ring oscillator and operating with a 22.6-MHz reference frequency, the synthesizer achieves a BW of around 5.6 MHz with a power consumption of 10 mW. Realized in 45-nm digital CMOS technology, the synthesizer exhibits a phase noise of −121.4 dBc/Hz at 10-MHz offset and an integrated jitter of 1.5 $\mathrm{ps_{rms}}$.**

***Index Terms*— $\Sigma\Delta$ noise, finite impulse response (FIR) filter, fractional-N synthesizer, loop bandwidth (BW), phase-locked loop (PLL).**

I. Introduction

WIDEBAND RF synthesizers have been of great interest for decades as they suppress the oscillator phase noise and allow fast switching. In this spirit, fractional-N synthesis was conceived to use a high reference frequency and avoid the tradeoff between the loop bandwidth (BW) and the channel spacing [1], [2]. However, the use of $\Sigma\Delta$ modulators and the problem of their noise peaking have worked against this premise.

A greater synthesizer BW becomes even more attractive, if ring oscillators are used for frequency generation. With their wide tuning range and low sensitivity to substrate coupling, ring topologies provide more degrees of freedom in synthesizer design than doing their LC counterparts. The cost is a higher phase noise unless the loop BW can be increased.

This paper proposes a fractional-N synthesizer architecture that achieves a BW close to $f_{\mathrm{REF}}/4$, thereby allowing the use of a ring oscillator for 2.4-GHz RF applications [3]. This is accomplished through the use of a mostly digital noise filter having a much sharper roll-off than the analog counterparts. Section II provides the background for this paper, and Section III illustrates the proposed noise suppression technique. Section IV describes the synthesizer architecture, and Section V deals with several design considerations. Section VI presents the experimental results.

Manuscript received September 30, 2017; revised December 10, 2017 and January 15, 2018; accepted January 16, 2018. Date of publication February 9, 2018; date of current version May 24, 2018. This work was supported in part by the Qualcomm Innovation Fellowship Program, in part by the Broadcom Fellowship Program, and in part by Realtek Semiconductor. This paper was approved by Associate Editor Woogeun Rhee. *(Corresponding author: Long Kong.)*

The authors are with the Department of Electrical Engineering, University of California at Los Angeles, Los Angeles, CA 90095-1594 USA (e-mail: longkong@ucla.edu; razavi@ee.ucla.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2018.2796544

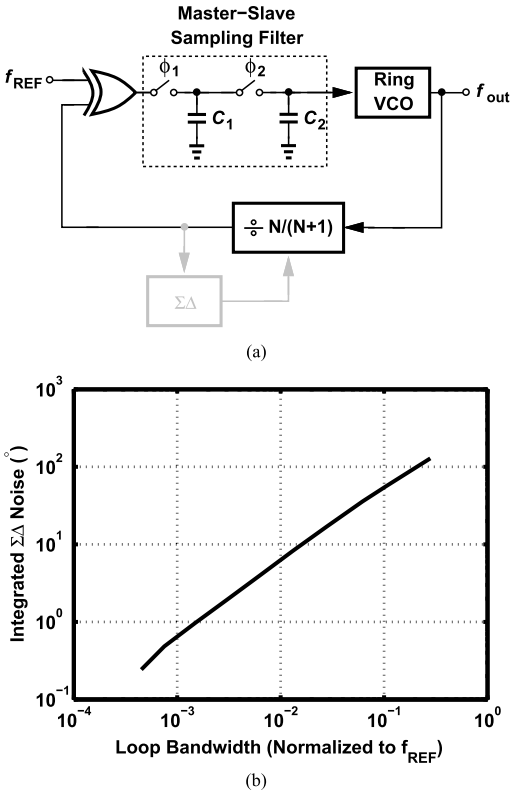

Fig. 1. (a) Synthesizer with MSSF. (b) Integrated noise of a third-order $\Sigma\Delta$ modulator at PLL output versus loop BW.

II. Background

An integer-N architecture achieving a BW of about $f_{\mathrm{REF}}/2$ has been proposed in [4]. Illustrated in Fig. 1(a), this approach employs an XOR phase detector (PD) to increase the BW and a master–slave sampling filter (MSSF) to suppress the ripple on the oscillator's control voltage. Note that, due to the finite loop gain, the XOR output exhibits the pulses of certain width, but the MSSF removes the ripple. One can readily add fractional-N operation to this system, but the $\Sigma\Delta$ modulator noise peak at $f_{\mathrm{REF}}/2$ would cause high phase noise. To quantify this tradeoff, we assume a third-order multi-stage noise shaping (MASH) $\Sigma\Delta$ topology, integrate the synthesizer output phase noise from a zero offset to f_{REF}, and plot the result as a function of the loop BW [Fig. 1(b)]. It can be seen that for the integrated noise to remain negligible (e.g., <0.45°), the BW cannot exceed $0.0007 f_{\mathrm{REF}}$.[1] On the other hand, for sufficient voltage-controlled oscillator (VCO) phase noise

[1]While not reflected here, higher order loop filters can be used as the loop BW becomes very small, allowing values greater than $0.0007 f_{\mathrm{REF}}$.

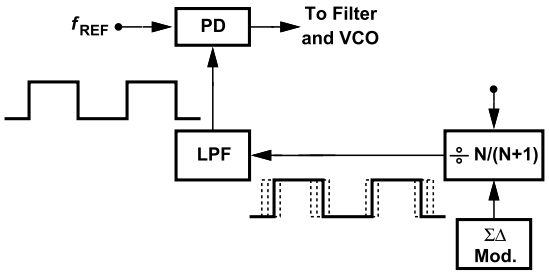

Fig. 2. Suppression of $\Sigma\Delta$ quantization noise by an LPF.

attenuation, we must increase the BW by about a factor of 360. This contention between the $\Sigma\Delta$ and VCO contributions can be avoided only if the total $\Sigma\Delta$ noise is reduced independently by another 50 dB.

In order to overcome the BW-noise tradeoff, a number of techniques have been introduced. For example, feedforward cancellation by means of a digital-to-analog converter (DAC) can reduce the $\Sigma\Delta$ noise by a maximum of 15–33 dB [5]–[8]. This concept has been extended to the time domain where a digital-to-time converter (DTC) is utilized to cancel the quantization noise [9]–[11], but gain error and nonlinearity issues encountered in DAC-based solutions exist here as well. Another approach is to raise the $\Sigma\Delta$ oversampling ratio and hence the reference frequency. This concept leads to cascaded phase-locked loop (PLL) topologies, wherein an integer-N loop multiplies the main reference and applies the result to a fractional-N loop [12], [13]. The use of two VCOs, however, is generally cumbersome and prone to unwanted couplings. We should also remark that two VCOs operating at high frequencies in a cascade always present a less favorable noise-power tradeoff than does a single VCO.

An alternative method is to filter—rather than simply cancel—the $\Sigma\Delta$ noise before it reaches the VCO. Such a filter must, of course, not limit the loop BW, an interesting dilemma. For example, [14] implements such a filter by inserting between the VCO and the feedback divider a phase interpolator (PI) driven by delayed copies of the $\Sigma\Delta$ modulator output, thus creating a finite impulse response (FIR) filter action. This design requires an eight-phase VCO and incorporates equal FIR tap coefficients, achieving a BW of $f_{REF}/10$ with a noise reduction of 34 dB [14]. As another method, [13] presents multiple notches to the $\Sigma\Delta$ noise transfer function by adding the divider output with its delayed version. This approach, however, filters out the $\Sigma\Delta$ noise only around the notch frequency. To achieve greater filtering, a low-pass filter (LPF) that suppresses the $\Sigma\Delta$ noise from low frequencies up to the loop BW is highly desired.

In this paper, we introduce a single-loop fractional-N synthesizer that performs the filtering action between the feedback divider and the PD. A mostly digital 35-tap FIR structure with arbitrary coefficients suppresses the $\Sigma\Delta$ noise at $f_{REF}/2$ by 62 dB and affords a loop BW of around $f_{REF}/4$.

III. Proposed Noise Suppression Method

A. Basic Idea

We wish to interpose an LPF between the feedback divider and the PD so as to remove the phase fluctuations arising from the $\Sigma\Delta$ action (Fig. 2). From another perspective, the filter receives a pulsewidth-modulated signal from the divider and must produce a "clean," periodic signal having a frequency equal to f_{REF}. If the average divide ratio is equal to $N + \alpha$, where α is the fractional part, then the VCO frequency is given by $(N + \alpha) f_{REF}$.

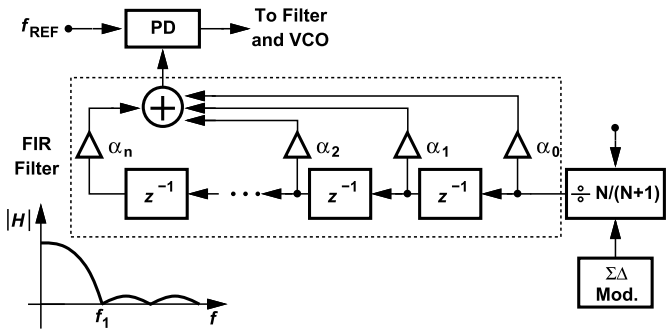

Fig. 3. Conceptual diagram of the proposed idea.

Operating on a phase (rather than a voltage) quantity, the filter can be more efficiently realized in the digital domain (Fig. 3), but, unlike standard digital FIR filters, it must process a 1-bit input. We thus face three questions: 1) How do we configure the filter to operate on phase quantities? 2) How do we choose the clock frequency of the filter? and 3) How do we implement the FIR tap coefficients, which may require a word length of, say, 10 bits, for a 1-bit digital signal?

Let us begin with the first question, and note that the divider output contains phase jumps equal to integer multiples of the VCO period, T_{VCO}. Conceptually, we surmise that if delayed copies of the divider output are created and properly summed, then the result can have less jitter. We must therefore return to Fig. 3 and determine how many delayed copies and with what weighting factors should be added. This is equivalent to defining the necessary frequency response of the FIR filter (Section III-C).

It is important to note that mismatches among the α coefficients simply alter the transfer function but do not cause nonlinearity, a critical advantage over DAC-based solutions.

We can consider voltage-domain implementation of the coefficients and the summation in Fig. 3, but we would face a serious difficulty: the residual $\Sigma\Delta$ noise in the summer output would be subjected to the PD's limiting action, experiencing folding to baseband. According to simulations, such an arrangement would raise the baseband noise by 30 dB. It is thus necessary to seek another method of scaling and combining the copies.

The next issue relates to the choice of the clock frequency and the FIR unit delay, i.e., z^{-1} in Fig. 3. Since the $\Sigma\Delta$ noise reaches a peak at $f_{REF}/2$, and since the plot in Fig. 1(b) dictates a filter BW well below this value, we expect that the first null of $|H|$ in Fig. 3, f_1, must be a small fraction of f_{REF}. That is, a unit delay on the order of $T_{REF} = 1/f_{REF}$ is necessary. For example, 35 such units provide zeros at the integer multiples of $f_{REF}/35$.

Unfortunately, the units cannot be clocked by f_{REF}: even if a single flip-flop clocked by f_{REF} follows the divider, the flip-flop's output carries no feedback information, prohibiting lock.

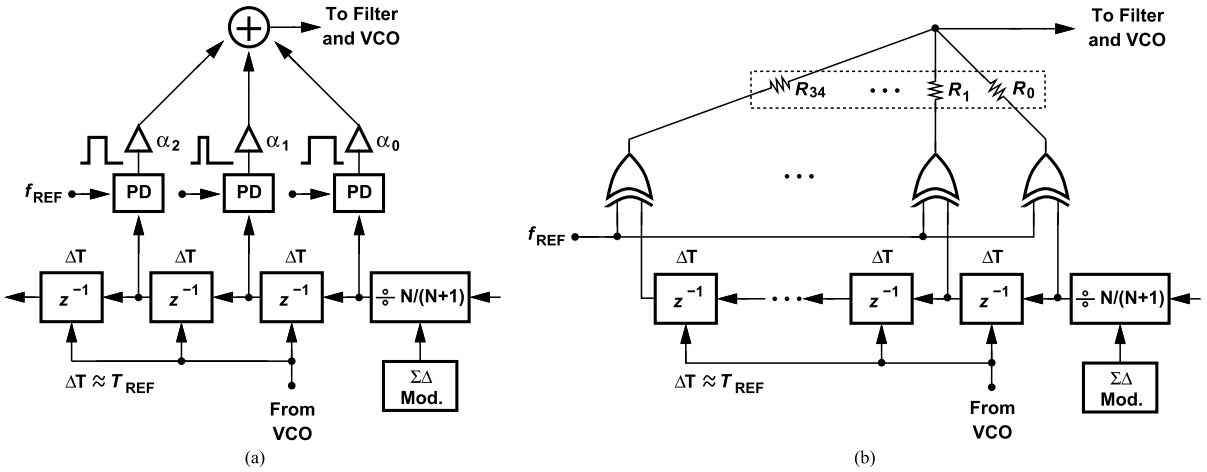

Fig. 4. FIR filter (a) with phase detection, multiplication, and summation and (b) implementation. (z^{-1} consists of approximately N flip-flops).

Our proposed solution clocks the units by the VCO output, and includes enough flip-flops in Fig. 3 to produce a total unit delay equal to T_{REF}. One naturally asks whether, instead, we can use any other low-frequency clock available in the PLL for this purpose so as to reduce the power consumption. We recognize that even if $f_{VCO}/2$ is utilized, the first flip-flop's output has a time resolution of $2T_{VCO}$ whereas its input exhibits phase jumps equal to $\pm T_{VCO}$, $\pm 2T_{VCO}$, and so on. This phase truncation would fold considerable $\Sigma\Delta$ noise to baseband.

The architecture of Fig. 3 offers two advantages over DAC feedforward cancellation. First, it is inherently a linear system and hence free from noise folding. (The loop imperfections are studied in Section V-A.) Second, its cancellation characteristics can be more precisely defined than the latter because it does not rely on matching between a DAC and a charge pump.

We should point out another property of this architecture. Since all of the flip-flops in the delay line are clocked by the VCO, any perturbation in the VCO phase readily appears at the end of the delay line; that is, the FIR structure does not filter the VCO output and hence does not limit the PLL BW [13].

B. FIR Implementation

We now deal with the implementation of the coefficients and the summation in Fig. 3. As shown in Fig. 4(a), we swap the order of summation and phase detection, i.e., we first compute the phase difference between the output of each delay stage and the reference and then combine the PDs' outputs with proper weighting factors.

Sensing the entire $\Sigma\Delta$-induced phase jumps, each PD must be linear enough not to fold noise to baseband. An XOR gate is a suitable candidate. According to transistor-level simulations, the rise in the baseband noise observed at the XOR output is negligible. The pulses generated by the PDs in Fig. 4(a) represent $\phi_{REF} - \phi_0$, $\phi_{REF} - \phi_1$, ..., $\phi_{REF} - \phi_n$, where ϕ_1, ..., ϕ_n denote the delay line output phases. Scaling the amplitudes of the pulses by the coefficients and adding the results, we obtain a quantity proportional to $\alpha_0(\phi_{REF} - \phi_0) + \alpha_1(\phi_{REF} - \phi_1) + \cdots + \alpha_n(\phi_{REF} - \phi_n) = (\alpha_0 + \alpha_1 + \cdots + \alpha_n)\phi_{REF} - (\alpha_0\phi_0 + \alpha_1\phi_1 + \cdots + \alpha_n\phi_n)$. The first sum models a constant pulsewidth, i.e., a static phase, and the second can be made nearly constant by the proper choice of α_0, ..., α_n. In other words, the tailoring of the FIR response in the frequency domain corresponds to minimizing the variation of this sum in the time domain. To visualize this point, we plot in Fig. 5 the eye diagram produced by a certain filter response (Section III-C). As explained in Section III-C, only the falling edges matter here and they exhibit a jitter of 7 ps_{pp}. The divider output, on the other hand, produces about 1.68 ns_{pp} of jitter.

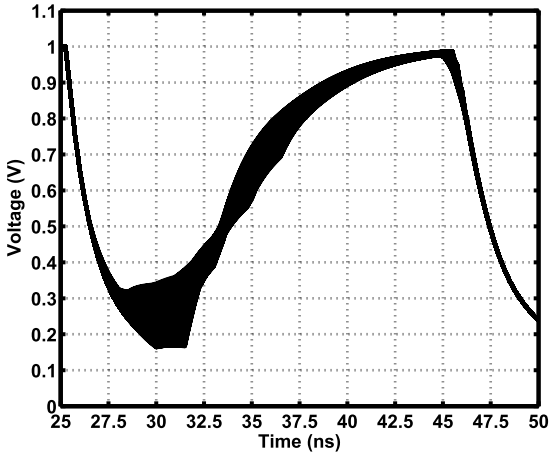

Fig. 5. Eye diagram of filtered output.

The actual implementation is shown in Fig. 4(b), where resistors perform the addition in the voltage domain, and $\alpha_j = (R_0 \parallel R_1 \parallel \cdots \parallel R_{34})/R_j$. The inverse dependence upon R_j means that we must first invert the coefficients and compute a convenient unit resistor value (Section III-C).

C. Filter Frequency Response

The length and the coefficients of the FIR filter must be chosen so as to create a frequency response that sufficiently

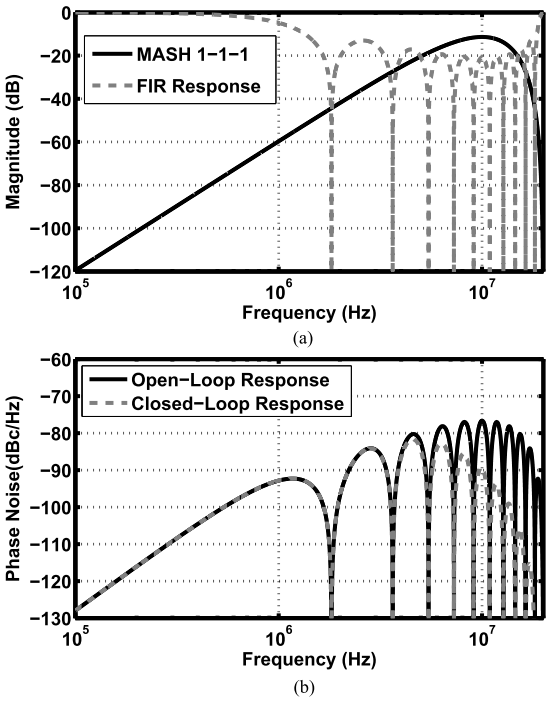

Fig. 6. (a) Output spectrum of MASH 1–1–1 modulator and example of magnitude response for an 11-tap FIR filter. (b) Resulting phase noise plots for open-loop and closed-loop synthesizers.

attenuates the $\Sigma\Delta$ noise. As a guideline, Fig. 1(b) shows that a suppression of more than 50 dB is necessary, if a loop BW near $f_{\rm REF}/4$ is desired. Thus, the FIR design should be tailored such that the area under the *product* of its response and the $\Sigma\Delta$ noise spectrum is small enough. Of course, since the delay elements are clocked at the VCO frequency, we wish to minimize their number.

Simulations can readily rule out certain FIR designs and point to more appropriate choices. Fig. 6(a) shows the output spectrum of a 1–1–1 MASH $\Sigma\Delta$ modulator clocked at $f_{\rm REF} = 20$ MHz. As an example, let us subject this profile to an integrator (representing the path from the modulus control to the divider output phase) and then to an FIR response with 11 equal taps [Fig. 6(a)]. We recognize that the rejection at $f_{\rm REF}/2 (= 10$ MHz) is equal to $20 \log 11$ because the transfer function, $(z^{-10} + \cdots + z^{-1} + 1)/11$, assumes a value of 1 at $f = 0$ and 1/11 at $f = f_{\rm REF}/2$. The resulting phase noise spectrum is shown in Fig. 6(b), for the open-loop and closed-loop synthesizers. The area under the closed-loop response is as high as -17.2 dBc from $f = 0$ to $f = f_{\rm REF}$. Moreover, the product has a local peak around $f = f_{\rm REF}/4.4$ because the FIR filter's 3-dB BW ($\approx$0.8 MHz) is excessively large.

The foregoing observations reveal that a greater number of taps must be chosen, both to achieve a higher rejection in the vicinity of $f_{\rm REF}/2$ and to reduce the BW (and hence avoid the local peak). For example, a 35-tap realization with equal coefficients yields a rejection of $20 \log 35 \approx 31$ dB at $f_{\rm REF}/2$ and a BW of approximately $f_{\rm REF}/79$. However, even such a design produces an integrated phase noise of 3.5°.

In order to alleviate the situation, we can further increase the number of taps. To achieve a 50-dB reduction in integrated noise, we require 285 taps if the coefficients are equal, which translates to tens of thousands of flip-flops. Alternatively, we can select the individual coefficients so as to tailor the response to our needs, a flexibility offered by the resistive combining scheme proposed in Fig. 4(b). Fig. 7(a) illustrates an example of such an endeavor: a 35-tap Kaiser–Bessel design [15] provides a 3-dB BW of 400 kHz and a rejection of 65 dB at $f_{\rm REF}/2$.[2] More importantly, the closed-loop output spectrum [Fig. 7(b)] exhibits an integrated phase noise of 0.38°. The coefficients are given by $\alpha_{34} = \alpha_0 = 0.0011, \alpha_{33} = \alpha_1 = 0.0024, \ldots, \alpha_{17} = 0.0575$ [15].

For the $\Sigma\Delta$ noise spectrum assumed here, a 35-tap Kaiser–Bessel filter does not necessarily provide the optimum response. In fact, it is possible to adjust the coefficients

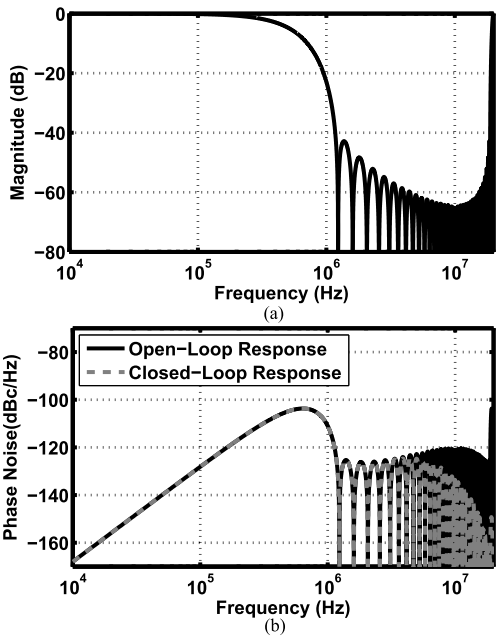

Fig. 7. (a) 35-tap Kaiser–Bessel FIR response. (b) Synthesizer output phase noise due to $\Sigma\Delta$ contribution.

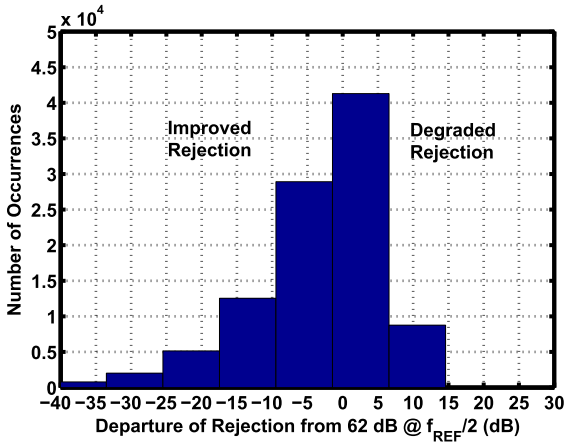

Fig. 8. Monte-Carlo results showing departure of rejection from 62 dB at $f_{\rm REF}/2$.

[2]The Kaiser window provides more flexibility in choosing the filter response than other windows.

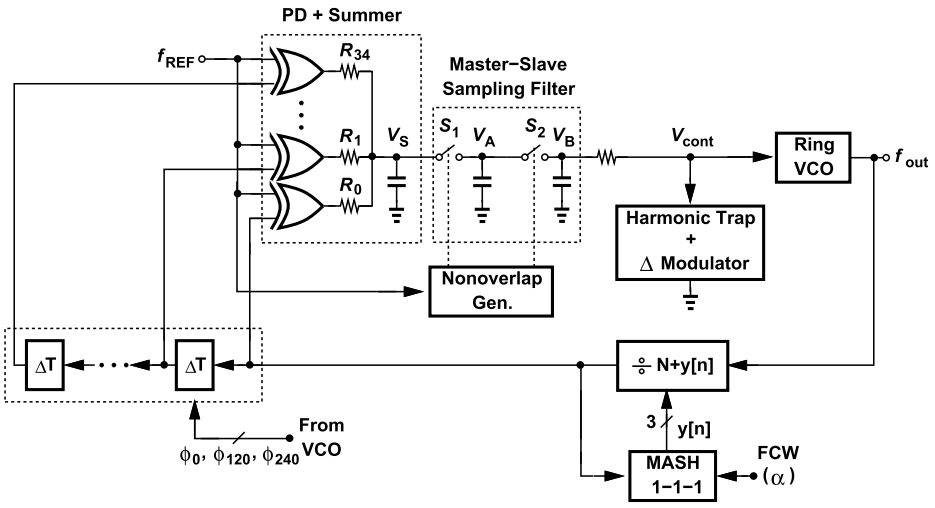

Fig. 9. Proposed synthesizer architecture.

manually so as to lower the BW but at the cost of rejection at $f_{REF}/2$ and vice versa. It may even be possible to achieve a slightly lower integrated phase noise through such adjustments. However, another important factor is the realizability of the coefficients, i.e., the "resistor spread" in Fig. 4(b) as well as the resolution necessary in implementing the resistor values. In the example described earlier, the spread is equal to $\alpha_{17}/\alpha_{34} = 52.3$, and the coefficient values are such that a unit resistor of 400 Ω can realize them with reasonable accuracy.

For the Kaiser–Bessel response, we have the following resistor values: $R_0 = R_{34} = 40$ kΩ, $R_1 = R_{33} = 18.5$ kΩ, ..., $R_{17} = 0.783$ kΩ. We must decide on a unit resistor, R_u, that can yield these values with acceptable accuracy. We choose $R_u = 400$ Ω and use various parallel and series combinations to form $R_0, \ldots, R_{34}$, inevitably incurring some error. The actual filter response has a rejection of 62 dB at $f_{REF}/2$.

D. Effect of Nonidealities

The XORs and the resistors in Fig. 4(b) suffer from two imperfections that can affect the performance. We first consider the effect of resistor mismatches. While introducing no nonlinearity, these mismatches do alter the transfer function, particularly the rejection at $f_{REF}/2$. Fortunately, it is much simpler to obtain a given level of matching between resistors than the transistors (as would be necessary in DACs, charge pumps, and PIs).

In a typical analog layout, resistor mismatches can be readily maintained below about 1%. We thus perform Monte-Carlo simulations to determine how much the rejection varies. The simulations assume that each unit resistor in the array of $R_0, \ldots, R_{34}$ can vary according to a Gaussian distribution with $\sigma = 1\%$. We also compute the mismatches of the XOR output resistances from cadence Monte-Carlo simulations and combine the results with the unit resistor mismatches. Fig. 8 shows the departure of the rejection from 62 dB. The distribution's asymmetry is due to the inverse dependence of the coefficients upon their corresponding resistor values. Here, the samples that fall below 0 dB have a greater rejection. Of the 100 000 samples, only about 5% experience a degradation of 10 dB or higher. (If even a tighter distribution is desired, a careful layout can yield resistor mismatches less than 0.5%, as is practical in A/D and D/A converter designs.)

The other imperfection relates to the finite output resistance of the XOR gates as it appears in series with the actual resistors. This resistance must be small enough to introduce negligible mismatches. The XOR design employs large transistors (W/L $= 4$ μm/40 nm) for low flicker noise, and its output resistance is still about 200 Ω, but it has negligible impact on the response (Appendix A).

IV. Proposed Synthesizer Architecture

In this section, we apply the $\Sigma\Delta$ noise suppression method described earlier to the synthesizer architecture introduced in Fig. 1(a). Shown in Fig. 9, the fractional-N loop consists of the FIR filter, including the XOR PDs and the resistive combiner, an MSSF, a ring VCO, a multi-modulus feedback divider, and an MASH 1–1–1 $\Sigma\Delta$ modulator. To reduce the ripple on the oscillator's control voltage, a harmonic trap establishes notches at f_{REF} and $2f_{REF}$, and a Δ modulator sets the notch frequencies by measuring and minimizing the ripple [4]. It is demonstrated in [4] that this architecture can achieve a BW close to $f_{REF}/2$ in the integer-N mode. Fig. 10 shows the synthesizer $\Sigma\Delta$-induced phase noise for BW $= f_{REF}/2$ and $f_{REF}/4$. Despite the suppression provided by the FIR action, the output exhibits 7 dB of peaking in the former case. The BW is thus reduced to $f_{REF}/4$ by lowering the PD gain, K_{PD}. (The sampling filter provides notches at mf_{REF} and only a slight attenuation below $f_{REF}/2$.)

The MSSF in Fig. 9 is clocked by nonoverlapping phases derived from f_{REF}, a point of contrast to the integer-N architecture in [4], where these phases are obtained from the feedback signal. This choice is made here to ensure that the MSSF clock itself is free from the $\Sigma\Delta$ noise and hence does not corrupt the sampled signal.

The use of the MSSF considerably widens the frequency capture range. Following the analysis in [4], we can express

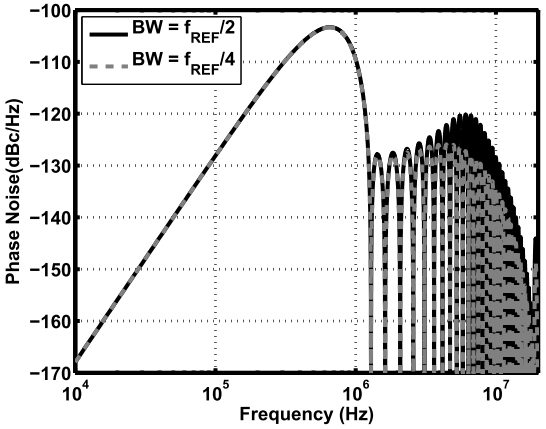

Fig. 10. Synthesizer $\Sigma\Delta$-induced phase noise for different loop BWs.

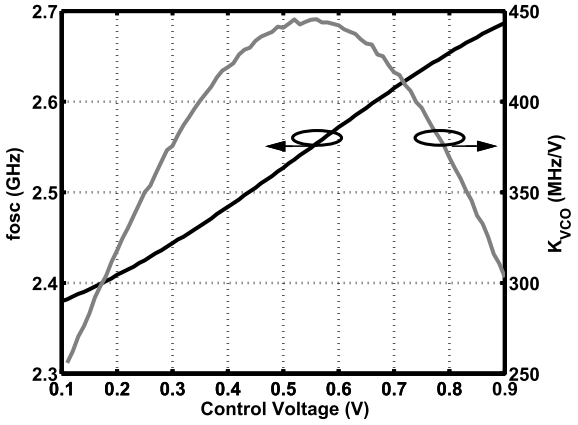

Fig. 11. Simulated VCO frequency and tuning gain versus control voltage.

the range for the present design as:

$$|f_{\text{REF}} - f_{\text{FB}}| < \frac{f_{\text{REF}}}{2} \tag{1}$$

where f_{FB} denotes the feedback signal frequency. Thus

$$\frac{f_{\text{REF}}}{2} < f_{\text{FB}} < \frac{3f_{\text{REF}}}{2}. \tag{2}$$

Translated to the VCO output, this range is around 1.2 to 3.6 GHz, much wider than our VCO's analog tuning range (320 MHz). In actual measurements, we first turn on only the tap with $\alpha_{17} = 0.0575$ to reach lock and then the remaining taps. According to simulations, the overall lock transient takes about $50T_{\text{REF}}$. This sequence can be automated through the use of a counter that initially counts about 50 reference cycles and subsequently enables the remaining taps.

The VCO in Fig. 9 is realized as a three-stage ring topology with both varactor tuning and digital control [4]. Plotted in Fig. 11 are the oscillation frequency and K_{VCO} versus the control voltage. The performance can be summarized as follows: overall tuning range: 2.31–3.05 GHz, $K_{\text{VCO}} \approx 400$ MHz/V; power consumption: 3 mW, phase noise at 5-MHz offset: -115.5 dBc/Hz.[3]

The MASH 1–1–1 $\Sigma\Delta$ modulator in Fig. 9 employs a word length of 16 bits for a frequency resolution of 345 Hz at 2.4 GHz. The feedback divider modulus changes in unity steps from 103 to 110 by the modulator. The FIR filter is clocked by three VCO phases for the reasons explained in Section V-C.

V. Design Considerations

A. Nonlinearity Issues

While reduced considerably by the FIR filter in Fig. 9, the $\Sigma\Delta$ noise can still experience folding, if the MSSF and VCO are excessively nonlinear. In order to formulate the MSSF nonlinearity, we consider the simplified cascade shown in Fig. 12(a), where the XOR and $R_{\text{eq}}(\approx 45\ \Omega)$ represent the parallel combination of the 35 XORs and their summing resistors in Fig. 9. Switches S_1 and S_2 are controlled by nonoverlapping phases derived from f_{REF}, and the effect of C_0 is ignored for now. Recall from Section III-C that the resistive summer output closely resembles a periodic waveform toggling between zero and V_{DD} (as in an integer-N architecture), with some residual $\Sigma\Delta$ noise on the transitions. This waveform must, therefore, be so filtered as to yield an acceptably small ripple on the VCO control line, a task performed by the MSSF in Fig. 9. (As shown in [4], the master–slave sampling action yields zeros in the transfer function located at mf_{REF}.) The operation of the cascade is illustrated by the waveforms in Fig. 12(a). Switch S_1 turns on as V_{REF} goes high (at $t = t_1$), allowing C_1 to charge to V_{DD}. At $t = t_2$, the PLL feedback signal, V_{FB}, rises, V_{XOR} falls, and C_1 begins to discharge. The phase difference between the falling edge of V_{REF} and the rising edge of V_{FB} is set by the synthesizer such that the discharge lowers V_A to the value required by the VCO, V_1. Finally, S_1 turns off and S_2 turns on, impressing V_1 on C_2. In the steady state, the nominal ripple amplitude is zero.

We wish to compute an equivalent PD gain for this cascade and examine its nonlinear behavior. Including the effect of $C_0(=40$ pF) (and neglecting droop due to leakage currents), we write the denominator of the transfer function from V_{XOR} to V_A as $R_{\text{eq}}R_{\text{on}}C_0C_1s^2 + (C_0R_{\text{eq}} + C_1R_{\text{on}} + C_1R_{\text{eq}})s + 1$, where R_{on} is the switch resistance, and calculate the poles. The falling transition of V_A can be expressed as

$$V_A(t) = -\frac{p_2}{p_1 - p_2}V_{\text{DD}}e^{p_1 t} + \frac{p_1}{p_1 - p_2}V_{\text{DD}}e^{p_2 t} \tag{3}$$

where p_1 and p_2 denote the poles. To calculate V_1 in Fig. 12(a), we set t to $[(\pi - \Delta\phi)/(2\pi)]T_{\text{REF}}$

$$V_1 = \frac{V_{\text{DD}}}{p_1 - p_2}\left[-p_2 \exp\frac{p_1(\pi - \Delta\phi)T_{\text{REF}}}{2\pi} + p_1 \exp\frac{p_2(\pi - \Delta\phi)T_{\text{REF}}}{2\pi}\right]. \tag{4}$$

This PD characteristic is shown in Fig. 12(b) for the values chosen in our work. Also shown is the case of $C_0 = 0$, exhibiting a higher nonlinearity. In Appendix B, we prove that, with $C_0 = 40$ pF, the XOR/MSSF cascade introduces negligible noise folding because the residual $\Sigma\Delta$ noise at the

[3]The simulated VCO tuning range including the discrete steps at TT 27 °C is 2.10–2.84 GHz, at SS 80 °C, 1.83–2.48 GHz, and at FF 0 °C, 2.39–3.23 GHz; K_{VCO} varies by 5% across the process and temperature.

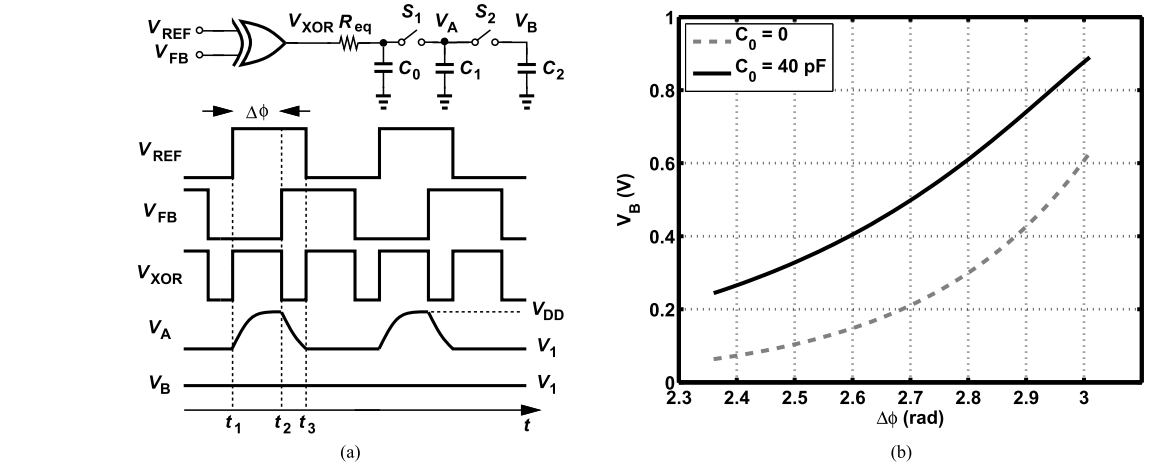

Fig. 12. (a) XOR PD driving MSSF. (b) Control voltage versus static phase error.

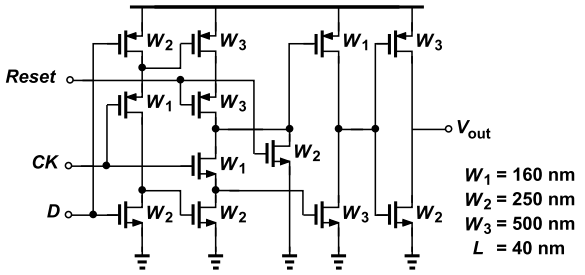

Fig. 13. TSPC flip-flop.

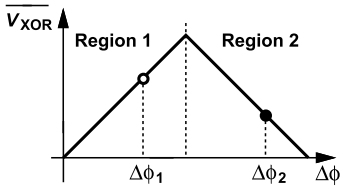

Fig. 14. Nonmonotonic XOR PD characteristic.

summer output is already small. Also, the width of S_1 in Fig. 9 is programmable so as to vary R_{on} with the output frequency and hence maintain a relatively constant K_{PD}.

The PD gain-nonlinearity tradeoff translates to one between the loop BW and noise folding. Based on simulations, the optimum values can be chosen so as to minimize the VCO phase noise and yet avoid significant folding.

B. Delay Line Implementation

The FIR filter in Fig. 9 consists of 34 unit delays equal to $T_{REF} = 1/f_{REF} \approx 50$ ns. In other words, the total delay presented to the divider output is 1.7 μs, a value extremely difficult to realize by the asynchronous stages such as inverters. We instead opt for flip-flops driven by the VCO, requiring about 4000 of them. (In our prototype, we employ 3672 flip-flops because our measurements use a crystal oscillator running at 22.6 MHz). The FIR filter data path operates at a low frequency, drawing negligible power. The clock path, on the other hand, is driven by the VCO and consumes a power given by $f_{VCO}C_{CK}V_{DD}^2$, where C_{CK} denotes the total capacitance in the clock path of the flip-flops. The true single-phase clock (TSPC) topology [16] proves useful here as it employs only two clocked transistors per flip-flop. Shown in Fig. 13 is the resettable design, consuming 1.5 μW at 2.4 GHz. The overall delay line thus dissipates 5.6 mW in 40-nm technology. This amount is expected to fall in more advanced nodes, making the proposed approach more power efficient. For example, according to simulations in 16-nm technology, the power reduces to 1.5 mW at 2.4 GHz, and the capacitance to 0.75 pF. Both of these values are reasonable for a 2.4-GHz synthesizer.

C. PD Nonmonotonicity

XOR gates operating as PDs exhibit a nonmonotonic behavior, if the phase error grows sufficiently large. As shown in Fig. 14, the PD gain changes the sign as the phase error goes from $\Delta\phi_1$ to $\Delta\phi_2$. We observe that this nonmonotonicity folds high-frequency noise, if the $\Sigma\Delta$-induced phase fluctuations force $\Delta\phi$ to cross from Regions 1 to 2. Indeed, with as many as 34 unit delays added to the phase of the feedback signal, some XORs in Fig. 9 can reside near this boundary and experience the random polarity reversals. (Recall that each XOR senses the reference and one delayed signal from the delay line). Noise folding can also be seen, if we approximate the XOR characteristic by a parabola and recognize that the $\Sigma\Delta$ noise is subjected to such nonlinearity.

To investigate this effect, we recall from Section III-A that the FIR unit delay, ΔT, is nominally equal to T_{REF}. However, the N flip-flops within each unit are clocked by the VCO, whose frequency may not be an integer multiple of f_{REF}. That is

$$\Delta T = NT_{VCO} \tag{5}$$

$$= \frac{N}{N+\alpha} T_{REF} \tag{6}$$

$$\approx \left(1 - \frac{\alpha}{N}\right) T_{REF}. \tag{7}$$

Since, the delay departs from $T_{\rm REF}$ by $(\alpha/N)T_{\rm REF}$, it appears as a phase offset and continues to accumulate through the delay line, producing a maximum error of $34(\alpha/N)T_{\rm REF}$ at the end. We can identify three groups of XORs: some fully operate in Region 1, some fully in Region 2, and some randomly cross the boundary, causing noise folding. But even the second group presents an issue: the negative gains carried by these XORs translate to negative FIR taps, distorting the Kaiser–Bessel response, and also partially cancel the average positive gains provided by the first group, reducing the loop gain. For these reasons, we employ a method that confines all of the XORs to Region 1 for any value of α.

To control and limit the amount of phase offset accumulation, we wish to adjust the FIR unit delay for different values of the $\Sigma\Delta$ fractional control word, α. Let us assume that the unit delay is changed from $NT_{\rm VCO}$ to $(N+k)T_{\rm VCO}$, allowing (5)–(7) to be written as

$$\begin{aligned}\Delta T &= (N+k)T_{\rm VCO}\\ &= \frac{N+k}{N+\alpha}T_{\rm REF}\\ &\approx \left(1+\frac{k}{N}\right)\left(1-\frac{\alpha}{N}\right)T_{\rm REF}\\ &\approx \left(1+\frac{k-\alpha}{N}\right)T_{\rm REF}. \end{aligned} \tag{8}$$

We recognize that, ideally, k is chosen equal to α so as to minimize the accumulation. But the synchronous delay line offers a resolution of $T_{\rm VCO}$ and hence a minimum k equal to 1, whereas α <1. Fortunately, the three-stage ring VCO provides three phases, improving the resolution to $T_{\rm VCO}/3$.

The foregoing thoughts lead to the delay control scheme depicted in Fig. 15, where three consecutive units are shown. For α <1/6, $k=0$, i.e., the unit delay is simply equal to $NT_{\rm VCO}=107T_{\rm VCO}$ [Fig. 15(a)]. For $1/6\le\alpha<1/2$, we insert one more FF in each unit and clock it with the 120°, 240°, and 0° VCO phases in every three consecutive units [Fig. 15(b)]. As a result, the unit delay rises to $(N+1/3)T_{\rm VCO}$, yielding a maximum $|k-\alpha|$ of 1/6. Similarly, for $1/2\le\alpha<5/6$, we seek a unit delay equal to $(N+2/3)T_{\rm VCO}$ and a maximum of 1/6 for $|k-\alpha|$. To this end, we insert another FF in every third unit [Fig. 15(c)] and swap the 0° and 120° phases. For $5/6\le\alpha$ <1, we have a similar arrangement as Fig. 15(a), but with one more FF in each unit [Fig. 15(d)].

With a maximum phase offset of $T_{\rm VCO}/6$ per unit, the accumulated delay reaches 34 × (69 ps) ≈ 2.36 ns. To this value, we must add the maximum $\Sigma\Delta$-induced phase fluctuation, about $2T_{\rm VCO}=0.83$ ns. Now we ensure that, with a departure equal to 2.36 ns + 0.83 ns = 3.19 ns, the XORs remain in Region 1.

D. Overall Phase Noise

With the $\Sigma\Delta$ noise suppressed by the FIR filter, the synthesizer output phase noise arises primarily from the VCO. The $f_{\rm REF}/4$ BW relaxes the VCO design, allowing a phase noise value as high as −115.5 dBc/Hz at 5-MHz offset. The overall

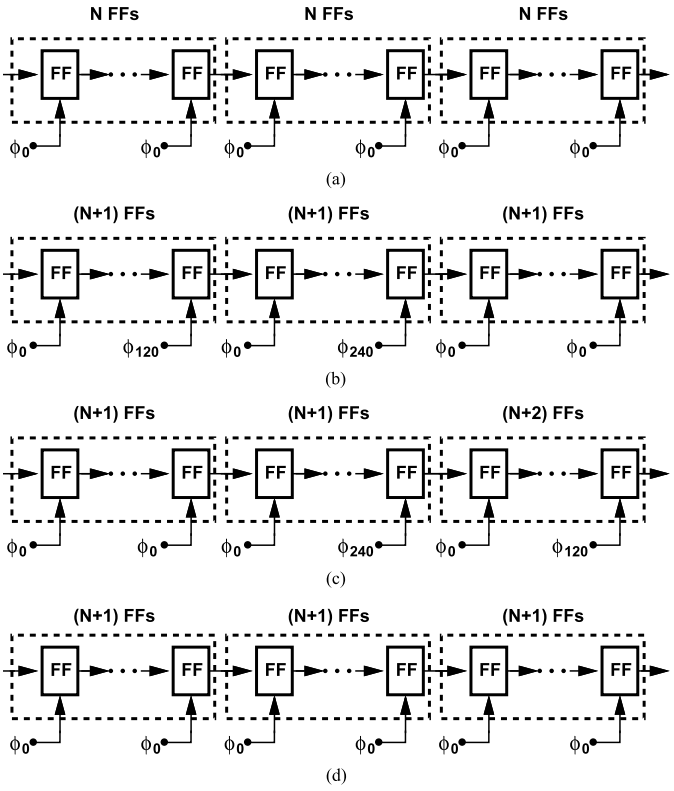

Fig. 15. Delay arrangement when (a) α < 1/6, (b) $1/6\le\alpha<1/2$, (c) $1/2\le\alpha<5/6$, and (d) $5/6\le\alpha<1$.

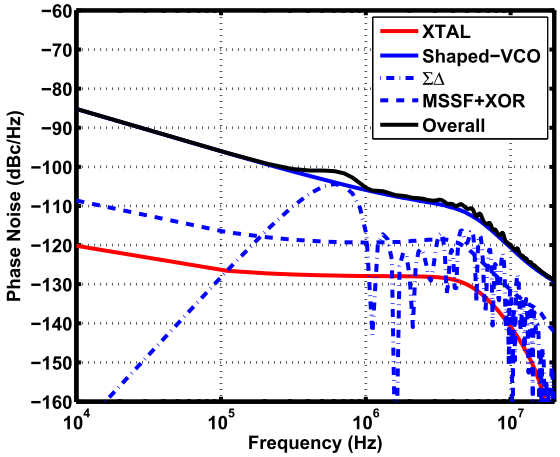

Fig. 16. Overall simulated PLL phase noise and main contributors, all referred to output.

simulated synthesizer phase noise along with the main contributors is shown in Fig. 16. The overall integrated jitter is equal to 1.44 $\rm ps_{rms}$, and despite a peak of −105 dBc/Hz, the $\Sigma\Delta$ noise translates to only 0.44 $\rm ps_{rms}$.

It is possible to further increase the loop BW, but at the cost of higher reference spurs. Also, as shown in Fig. 17, the reduction in the total jitter begins to diminish for BWs exceeding around 5 or 6 MHz.

VI. Experimental Results

The fractional-N synthesizer has been fabricated in TSMC's 45-nm digital CMOS technology and characterized with

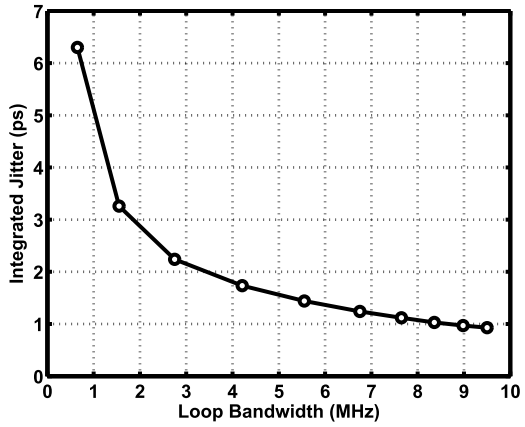

Fig. 17. Simulated integrated jitter versus loop BW.

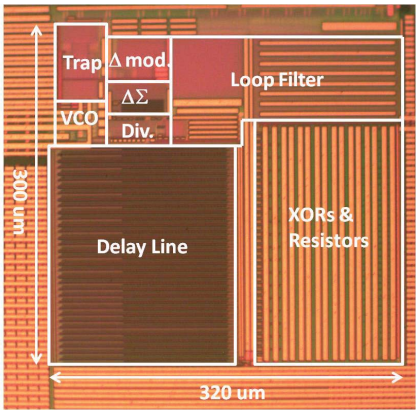

Fig. 18. Die micrograph.

a 22.6-MHz crystal oscillator acting as the reference. Operating from a 1-V supply,[4] the prototype draws 10 mW, with 3 mW dissipated in the VCO, 5.6 mW in the delay line, 0.4 mW in the divider, and 0.8 mW in the XOR PDs and MSSF. Shown in Fig. 18 is the die photograph, whose active area measures 300 μm $\times$320 μm.

Fig. 19 shows the measured output spectra in the fractional-N mode in two cases: only two taps of the FIR filter are enabled (top) and all 35 taps are enabled (bottom). The divide ratio is equal to $106 + 2^{-15} + 2^{-16}$. In the former case, the FIR filter provides some attenuation beyond $f_{REF}/2$ ($\approx$11 MHz), but the close-in phase noise remains very high. In the latter, on the other hand, the $\Sigma\Delta$ noise falls below the VCO contribution. The reference spurs are at -60 dBc.

Fig. 20 shows the measured phase noise in the integer-N mode, serving as a "baseline" for quantifying various contributions. In this mode, the MSSF, VCO, and crystal reference phase noise components are present. The in-band noise is about -107 dBc/Hz with a BW of around 5.6 MHz. Integrated from 10-kHz to 50-MHz offset, the rms jitter amounts to 1.15 ps. Fig. 21 plots the measured output phase noise in the fractional-N mode for two-tap and 35-tap FIR configurations. In the former, the phase noise remains as high as -76 dBc/Hz

[4]The VCO and harmonic trap use a 1-V analog supply while the remaining blocks share a 1-V digital supply.

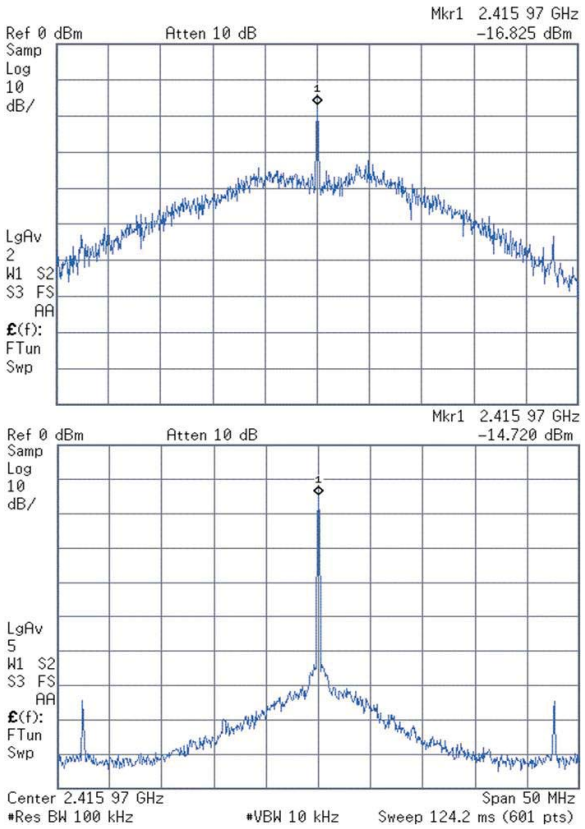

Fig. 19. Measured spectra in fractional-N mode with 2-tap FIR (top) and 35-tap FIR (bottom). (Horizontal scale: 5 MHz/div., vertical scale: 10 dB/div.)

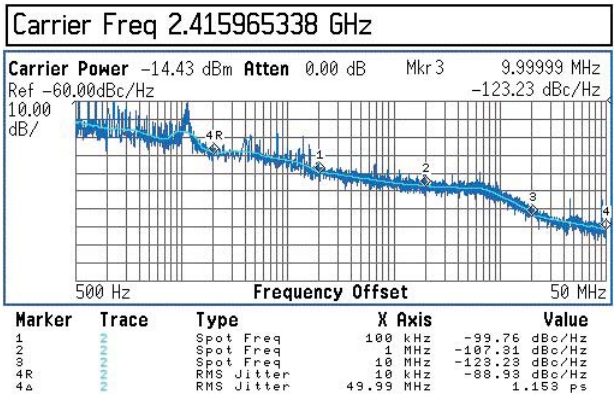

Fig. 20. Measured phase noise in integer-N mode.

up to 10-MHz offset. We also observe heavy noise folding due to the MSSF nonlinearity. In the latter, the phase noise at 10-MHz offset falls by 45 dB, reaching -121.4 dBc/Hz, which corresponds to 62 dB of reduction compared with the case without an FIR filter. The integrated rms jitter drops from 100 ps to 1.5 ps. Compared with the integer-N mode, the phase noise degrades by 3 dB at 1-MHz offset and 1.8 dB at 10-MHz offset.

Fig. 22 shows the measured fractional spur levels as a function of the frequency offset. The in-band level is around -41 dBc and the magnitude approximately follows the synthesizer response. These spurs are attributed to the supply line shared between the FIR filter and the XOR PDs.

TABLE I
PERFORMANCE SUMMARY

	[14]	[17]	[18]	[19]	[20]	[13]	This Work
Oscillator Topology	Ring	Ring	LC	LC	Ring	Ring	Ring
Reference Freq. (MHz)	32	26	48	49.2	26	22.6	22.6
Output Frequency (GHz)	0.86~1.26	1.87~1.98	2.2~2.4	2.6~3.9	2	2.3~2.6	2.31~3.05
f_{BW}/f_{REF}	0.1	0.077	0.01	0.014	0.058	0.5	0.25
Phase Noise @ 1−MHz offset (dBc/Hz)	−103	−98	−114	−118	−98	−109	−104
Phase Noise @ 10−MHz offset (dBc/Hz)	−114	−115	−135	−140.2	−115	−113.6	−121.4
RMS Jitter (ps) Integ. range (MHz)	N/A	3.4 (0.004~40)	0.36 (0.01~30)	0.23 (0.001~100)	2.4 (0.001~40)	1.7 (0.01~50)	1.5 (0.01~50)
Frac. Spur (dBc)	N/A	−50	−48	−62.3	−70	−52.5	−41
Ref. Spur (dBc)	−66	−67	−55	−60	−87	−70	−60
Freq. Resolution (kHz)	N/A	0.1	N/A	N/A	0.1	0.95	0.34
Power (mW)	16.8	10	17.3	11.5	9.1	6.4	10
Area (mm^2)	0.31	0.047	0.75	0.228	0.046	0.03	0.096
Tech. (nm)	130	40	180	65	40	45	45
FoM_1 (dB)	N/A	−219.4	−236.5	−241.8	−222.8	−227.3	−226.5
FoM_2 (dB)	150.7	153.4	168.7	175.9	154.4	168.4	161.7

$FoM_1 = 10\log_{10}[(\frac{\text{Jitter}}{1\text{ s}})^2(\frac{\text{Power}}{1\text{ mW}})]$ $FoM_2 = 10\log_{10}[(\frac{f_{osc}}{\Delta f})^2(\frac{1\text{ mW}}{\text{Power}})]$ − Phase Noise (dBc/Hz)

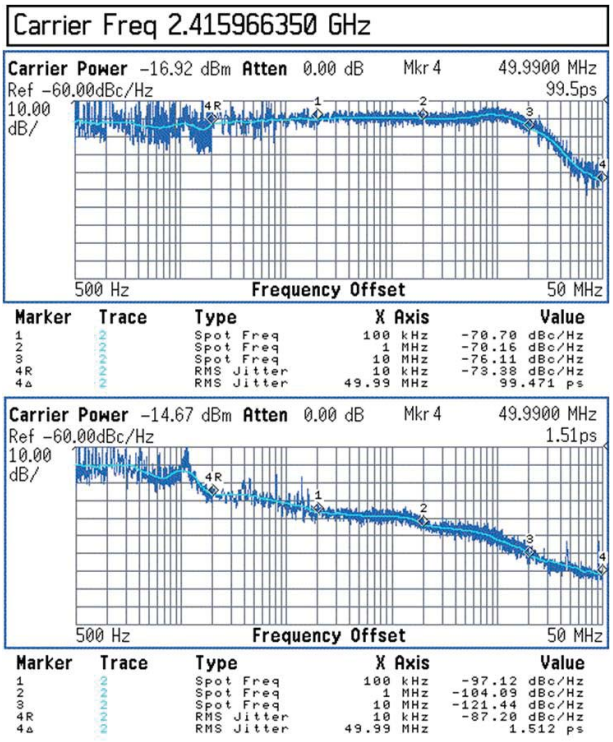

Fig. 21. Measured phase noise in fractional-N mode.

Simulations confirm that by separating these two supplies, the spur level will drop by 13 dB. We note that such in-band spurs negligibly degrade the signals in Wi-Fi applications as they simply cause a small fraction of one OFDM sub-channel to fall atop another. The spurs fall to −60 dBc at the edge of the RF channel.

Table I summarizes our performance and compares it with that of the state-of-the-art wideband fractional-N synthesizers. We note that our FoM_2 approaches that of the LC-VCO-based design in [18], if we account for their higher reference frequency.

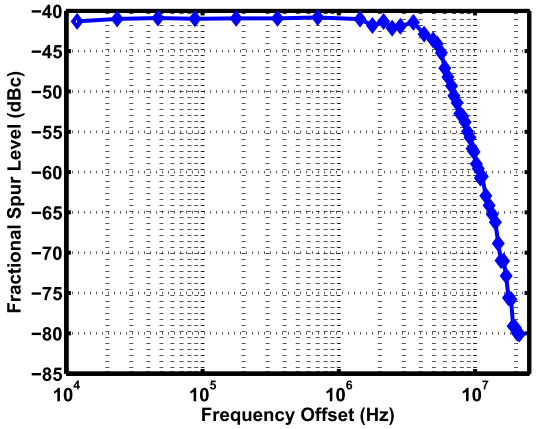

Fig. 22. Measured fractional spur level versus frequency offset.

VII. CONCLUSION

This paper has introduced a new fractional-N synthesizer architecture that incorporates a simple FIR filter and avoids analog circuits such as charge pumps and forward cancellation DACs. The heavy suppression of the $\Sigma\Delta$ noise afforded by the FIR filter allows a loop BW of about $f_{REF}/4$, thus providing the significant attenuation of the VCO phase noise as well.

APPENDIX A

In this appendix, we examine the effect of a constant shift in the resistor values that define the FIR coefficients in Fig. 4(b); for example, the XOR gates contribute a resistance of about 200 Ω.

The Kaiser–Bessel coefficients for an n-tap FIR filter are given by

$$\alpha_n = \frac{I_0\left[\pi a\sqrt{1-\left(\frac{2n}{N-1}-1\right)^2}\right]}{I_0(\pi a)} \tag{9}$$

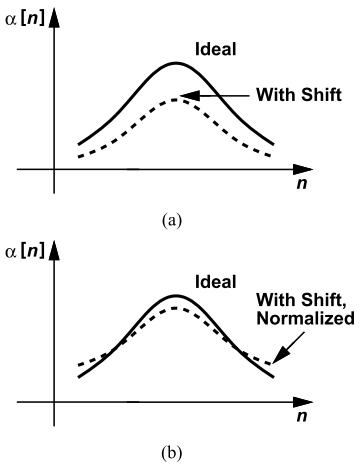

Fig. 23. (a) Coefficients plot in different scenarios (ideal and shifted cases). (b) Normalized plot.

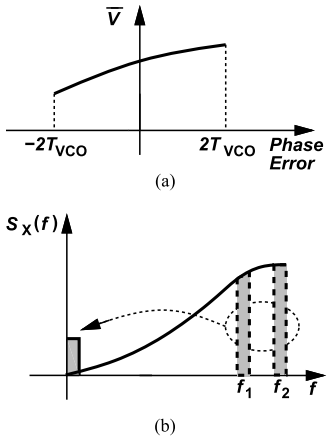

Fig. 24. (a) Nonlinear characteristic of PD. (b) Illustration of noise folding.

where I_0 denotes the modified Bessel function and parameter a provides an additional degree of freedom in creating the response. Fig. 23(a) conceptually sketches these values, whose total sum is unity. If R_0–R_{34} are shifted up by ΔR, then α_{17} experiences the greatest change and α_0 and α_{34} the least, yielding the dashed distribution. Normalizing the new coefficients so that their sum is unity gain, we arrive at the dashed plot in Fig. 23(b), which is, roughly, equivalent to adjusting a in (9). In this paper, a is adjusted by approximately 5.2%, negligibly affecting the response.

APPENDIX B

In this appendix, we derive an upper bound on noise folding for a second-order nonlinearity and apply the result to this paper. Let us approximate the nonlinear characteristic of a circuit by $y = \beta_1 x + \beta_2 x^2$, where x denotes the input phase and y the output voltage (as in a PD). Since the phase fluctuations produced by the feedback divider rarely exceed $\pm 2T_{VCO}$, we confine x to these two limits [Fig. 24(a)].

Suppose the input contains phase noise $\cos[\omega_0 t + x(t)]$ [Fig. 24(b)]. As a pessimistic approximation, we assume the entire noise is concentrated in two closely spaced narrow slivers at f_1 and f_2. The second-order intermodulation of these components causes folding to baseband. We approximate the two slivers by impulses and write $x(t) = A\cos\omega_1 t + A\cos\omega_2 t$. Note that, the peak phase fluctuation is equal to $2A$. We obtain the folded component as $(\beta_2 A^2/2)\cos(\omega_1 - \omega_2)t$. For example, we wish $\beta_2 A^2/2$ to be 80 dB below $\beta_1 A$.

For the PD characteristic in Fig. 12(b) with $C_0 = 40$ pF, we have $y \approx -2.74x + 0.69\ x^2$. Thus, $|(\beta_2 A^2/2)/(\beta_1 A)| = |(\beta_2/\beta_1)(A/2)|$ must be sufficiently small. Without the FIR filter, the peak phase fluctuation reaching the PD is about $2T_{VCO}$ and hence $2A = (2T_{VCO}/T_{REF}) \times (2\pi)$, yielding $|(\beta_2/\beta_1)(A/2)| = 0.0067 = -43.5$ dB. With the FIR filter, on the other hand, the peak fluctuation is around 3.5 ps, and the folded component at about −91 dB.

ACKNOWLEDGMENT

The authors are grateful to the TSMC University Shuttle Program for chip fabrication.

REFERENCES

[1] G. Gillette, "Digiphase synthesizer," in *Proc. Freq. Technol.*, Aug. 1969, pp. 25–29.
[2] J. Gibbs and R. Temple, "Frequency domain yield its data to phase-locked synthesizer," *Electronics*, vol. 27, pp. 107–111, Apr. 1978.
[3] L. Kong and B. Razavi, "A 2.4 GHz RF fractional-N synthesizer with 0.25 f_{REF} BW," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 330–331.
[4] L. Kong and B. Razavi, "A 2.4 GHz 4 mW integer-N inductorless RF synthesizer," *IEEE J. Solid-State Circuits*, vol. 51, no. 3, pp. 626–635, Mar. 2016.
[5] S. Pamarti, L. Jansson, and I. Galton, "A wideband 2.4-GHz delta-sigma fractional-NPLL with 1-Mb/s in-loop modulation," *IEEE J. Solid-State Circuits*, vol. 39, no. 1, pp. 49–61, Jan. 2004.
[6] M. Gupta and B.-S. Song, "A 1.8-GHz spur-cancelled fractional-N frequency synthesizer with LMS-based DAC gain calibration," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2842–2851, Dec. 2006.
[7] H. Hedayati, W. Khalil, and B. Bakkaloglu, "A 1 MHz bandwidth, 6 GHz 0.18 μm CMOS type-I $\Delta\Sigma$ fractional-N synthesizer for WiMAX applications," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3244–3252, Dec. 2009.
[8] A. Swaminathan, K. J. Wang, and I. Galton, "A wide-bandwidth 2.4 GHz ISM band fractional-N PLL with adaptive phase noise cancellation," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2639–2650, Dec. 2007.
[9] K. Raczkowski, N. Markulic, B. Hershberg, and J. Craninckx, "A 9.2–12.7 GHz wideband fractional-N subsampling PLL in 28 nm CMOS with 280 fs RMS Jitter," *IEEE J. Solid-State Circuits*, vol. 50, no. 5, pp. 1203–1213, May 2015.
[10] X. Gao *et al.*, "A 2.7-to-4.3 GHz, 0.16ps$_{RMS}$-Jitter, −246.8 dB-FOM, digital fractional-N sampling PLL in 28 nm CMOS," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2016, pp. 174–175.
[11] N. Markulic *et al.*, "A DTC-based subsampling PLL capable of self-calibrated fractional synthesis and two-point modulation," *IEEE J. Solid-State Circuits*, vol. 51, no. 12, pp. 3078–3092, Dec. 2016.
[12] D. Park and S. H. Cho, "A 14.2 mW 2.55-to-3 GHz cascaded PLL with reference injection and 800 MHz delta-sigma modulator in 0.13 μm CMOS," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, pp. 2989–2998, Dec. 2012.
[13] L. Kong and B. Razavi, "A 2.4-GHz 6.4-mW fractional-N inductorless RF synthesizer," *IEEE J. Solid-State Circuits*, vol. 52, no. 8, pp. 2117–2127, Aug. 2017.
[14] D.-W. Jee, Y. Suh, B. Kim, H.-J. Park, and J.-Y. Sim, "A FIR-embedded phase interpolator based noise filtering for wide-bandwidth fractional-N PLL," *IEEE J. Solid-State Circuits*, vol. 48, no. 11, pp. 2795–2804, Nov. 2013.
[15] J. F. Kaiser, "Nonrecursive digital filter design using the I_0-sinh window function," in *Proc. IEEE Int. Symp. Circuits Syst.*, Dec. 1974, pp. 20–23.
[16] J. Yuan and C. Svensson, "High-speed CMOS circuit technique," *IEEE J. Solid-State Circuits*, vol. 24, no. 1, pp. 62–70, Feb. 1989.

[17] T. K. Kao, C.-F. Liang, H.-H. Chiu, and M. Ashburn, "A wideband fractional-N ring PLL with fractional-spur suppression using spectrally shaped segmentation," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2013, pp. 416–417.
[18] W. S. Chang, P. C. Huang, and T. C. Lee, "A fractional-N divider-less phase-locked loop with a subsampling phase detector," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 2964–2975, Dec. 2014.
[19] Z.-Z. Chen *et al.*, "Sub-sampling all-digital fractional-N frequency synthesizer with −111 dBc/Hz in-band phase noise and an FOM of −242 dB," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, pp. 268–269, Feb. 2015.
[20] C. F. Liang and P. Y. Wang, "A wideband fractional-N ring PLL using a near-ground pre-distorted switched-capacitor loop filter," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2015, pp. 190–191.

Long Kong (M'16) received the B.E. degree in microelectronics from Shanghai Jiao Tong University, Shanghai, China, in 2011, and the M.S. and Ph.D. degrees in electrical engineering from the University of California at Los Angeles, Los Angeles, CA, USA, in 2013 and 2016, respectively.

In 2016, he joined Oracle, Santa Clara, CA, USA, as a Senior Hardware Engineer, where he was involved in high-speed SerDes transceivers. He is currently an RFIC Design Engineer with Apple, Cupertino, CA, USA. His current research interests include frequency synthesizers, clock and data recovery for data communication systems, and wireless transceivers.

Dr. Kong was a recipient of the Qualcomm Innovation Fellowship from 2013 to 2014, the Analog Devices Outstanding Student Designer Award in 2015, and the Broadcom Fellowship from 2015 to 2016.

Behzad Razavi (F'03) received the B.S.E.E. degree from the Sharif University of Technology, Tehran, Iran, in 1985, and the M.S.E.E. and Ph.D.E.E. degrees from Stanford University, Stanford, CA, USA, in 1988 and 1992, respectively.

He was an Adjunct Professor at Princeton University, Princeton, NJ, USA, from 1992 to 1994, and at Stanford University in 1995. He was with AT&T Bell Laboratories and Hewlett-Packard Laboratories in 1996. Since 1996, he has been an Associate Professor and subsequently a Professor of electrical engineering with the University of California at Los Angeles, Los Angeles, CA, USA. He has authored the books *Principles of Data Conversion System Design* (IEEE Press, 1995), *RF Microelectronics* (Prentice Hall, 1998, 2012) (translated to Chinese, Japanese, and Korean), the *Design of Analog CMOS Integrated Circuits* (McGraw-Hill, 2001, 2016) (translated to Chinese, Japanese, and Korean), the *Design of Integrated Circuits for Optical Communications* (McGraw-Hill, 2003, Wiley, 2012), and the *Fundamentals of Microelectronics* (Wiley, 2006) (translated to Korean and Portuguese). His current research interests include wireless transceivers, frequency synthesizers, phase-locking and clock recovery for high-speed data communications, and data converters.

Dr. Razavi is a member of the U.S. National Academy of Engineering. He has served as an IEEE Distinguished Lecturer. He served on the Technical Program Committees of the International Solid-State Circuits Conference (ISSCC) from 1993 to 2002 and VLSI Circuits Symposium from 1998 to 2002. He received the Beatrice Winner Award for Editorial Excellence at the 1994 ISSCC, the Best Paper Award at the 1994 European Solid-State Circuits Conference, the Best Panel Award at the 1995 and 1997 ISSCC, the TRW Innovative Teaching Award in 1997, the Best Paper Award at the IEEE Custom Integrated Circuits Conference in 1998, and the McGraw-Hill First Edition of the Year Award in 2001. He was a co-recipient of the Jack Kilby Outstanding Student Paper Award and the Beatrice Winner Award for Editorial Excellence at the 2001 ISSCC. He received the Lockheed Martin Excellence in Teaching Award in 2006, the UCLA Faculty Senate Teaching Award in 2007, and the CICC Best Invited Paper Award in 2009 and in 2012. He was a co-recipient of the 2012 and the 2015 VLSI Circuits Symposium Best Student Paper Awards and the 2013 CICC Best Paper Award. He was also recognized as one of the top ten authors in the 50-year history of ISSCC. He received the 2012 Donald Pederson Award in Solid-State Circuits, the American Society for Engineering Education PSW Teaching Award in 2014, and the 2017 IEEE CAS John Choma Education Award. He was an Editor of the *Monolithic Phase-Locked Loops and Clock Recovery Circuits* (IEEE Press, 1996) and *Phase-Locking in High-Performance Systems* (IEEE Press, 2003). He has also served as a Guest Editor and an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS, the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, and the *International Journal of High Speed Electronics*.