An Integral Path Self-Calibration Scheme for a Dual-Loop PLL

Mark Ferriss, Jean-Olivier Plouchart, *Senior Member, IEEE*, Arun Natarajan, Alexander Rylyakov, Ben Parker, José A. Tierno, A. Babakhani, Soner Yaldiz, Alberto Valdes-Garcia, Bodhisatwa Sadhu, and Daniel J. Friedman, *Member, IEEE*

***Abstract*—An integral-path self-calibration scheme is introduced as part of a 20.1 GHz to 26.7 GHz low-noise PLL in 32 nm CMOS SOI. A dual-loop architecture in combination with an integral path measurement and correction scheme desensitizes the loop transfer function to the VCO's small signal gain variations. The spread of gain peaking is reduced by self-calibration from 2.4 dB to 1 dB, when measured at 70 sites on a 300 mm wafer. The PLL has a measured phase noise @10 MHz offset of −126.5 dBc/Hz at 20.1 GHz and −124.2 dBc/Hz at 24 GHz**

***Index Terms*—Bandwidth calibration, frequency synthesizers, phase locked loop, PLL.**

I. Introduction

THE next generation of wireless and wireline applications require high frequency (>20 GHz) clocks with stringent phase noise performance and well controlled jitter transfer characteristics. For example, a 20–26 GHz PLL followed by a frequency doubler has been proposed as a solution to generating the LO in a 60 GHz radio [1]. High data rate wireline I/O macros must support multiple standards while maintaining excellent jitter performance, where an example of such an I/O macro is described in [2].

In order to support these high performance applications while achieving high yield, several practical challenges must be addressed. For example, minimizing the integrated jitter requires selecting the PLL's phase transfer function as a compromise between meeting the goals of filtering the reference clock noise (which requires low bandwidth) and filtering the VCO phase noise (which requires high bandwidth). In a conventional analog PLL, as shown in Fig. 1(a) the phase transfer function can vary significantly due to changes in the small signal gain terms of components in the loop. Shown in Fig. 2 is an example of a simulated voltage to frequency transfer curve of a high frequency VCO. The transfer curve is highly non-linear. The analog tuning range represents a small fraction of the total tuning range of the VCO, while the remainder of the range is accessed by switching digitally-controlled fixed capacitors. Small manufacturing variations in the digitally switched or fixed capacitors must be compensated by moving the analog control to a different point on the VCO tuning curve, resulting in differences in small signal gain. In a conventional single path PLL (Fig. 1(a)), this gain variation affects both the proportional gain and integral gain in a similar fashion. In addition to the VCO non-linearity, the integrating capacitor and the ripple capacitor may themselves be non-linear, particularly if they are implemented with MOS-based devices. Consequently the PLL's small signal phase transfer function is frequency dependent, and sensitive to PVT variations. In a single path PLL, the VCO gain variations affect both the proportional gain and integral gain, and the non-linearity of the capacitors affects the position of the first zero in the open loop transfer function in addition to the frequency of the higher order pole.

Several techniques have been proposed to address these problems. In [3]–[6] an additional VCO control path is introduced which re-centers the original VCO control voltage to the middle of its range, as shown in Fig. 1(b). This additional path contains a low gain integrator in series with the loop filter, resulting in a type III control loop. As a result, the small signal gain of the primary path is desensitized to the non-linearity of the VCO and loop capacitor because the VCO's control voltage is always forced to the center of its range irrespective of the output frequency. In addition, as the primary control path is no longer required to support a broad frequency range, the VCO gain of this path can be reduced, and in order to maintain the bandwidth, the charge pump current proportionally increased. This approach has been shown in [4] to lead to an improvement in phase noise performance. However, as will be discussed in Section II, a drawback of the very low gain integrator approach is that PLLs built in this way are susceptible to capacitor leakage-driven effects.

A second method of reducing the sensitivity of the loop to uncertainty in the small signal gain terms is to use a foreground calibration scheme to calibrate the PLL's transfer function by measuring the PLL's transient response to injected phase steps, as described in [7]. This technique has the advantage of requiring no additional analog circuitry, allowing for a minimal area overhead. However, as will also be discussed in Section V, the scheme from [7] is difficult to implement in a PLL with a

Manuscript received August 30, 2012; revised November 26, 2012; accepted December 12, 2012. Date of publication January 30, 2013; date of current version March 22, 2013. This paper was approved by Guest Editor Vivek De. This work was supported in part by DARPA under AFRL contract FA8650-09-C-7924.

M. Ferriss, J.-O. Plouchart, A. Natarajan, A. Rylyakov, B. Parker, J. A. Tierno, A. Babakhani, S. Yaldiz, A. Valdes-Garcia, B. Sadhu, and D. J. Friedman are with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: mferriss@us.ibm.com).

S. Yaldiz is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA, and also with Carnegie Mellon University, Pittsburgh, PA 15213 USA.

B. Sadhu is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA, and also with the University of Minnesota, Minneapolis, MN 55455 USA.

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.2013.2239114

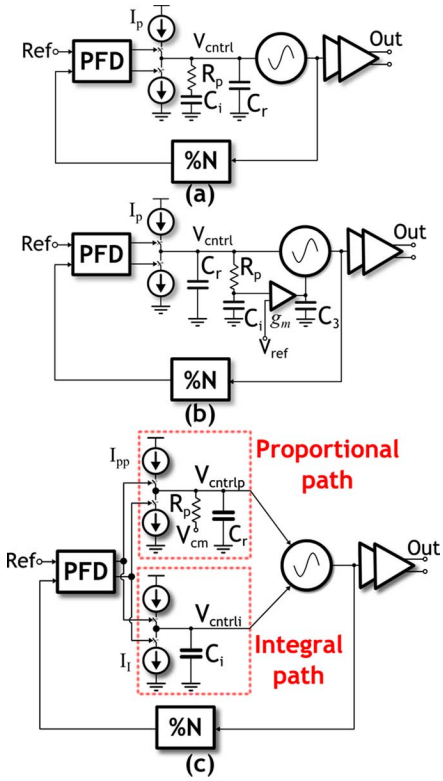

Fig. 1. Common PLL architectures. (a) Conventional single path PLL. (b) Dual loop type III PLL. (c) Dual loop type II PLL.

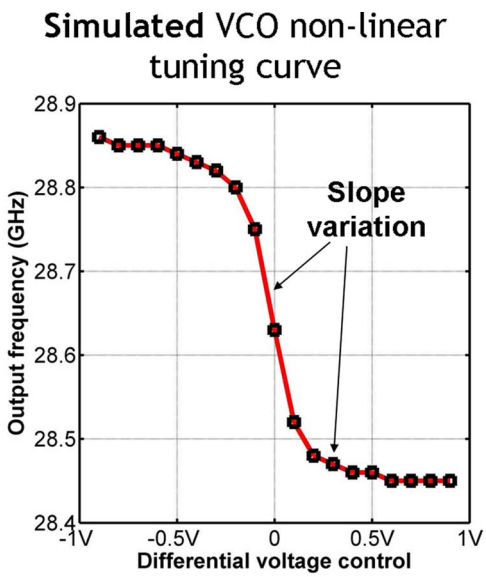

Fig. 2. A VCO's Voltage-to-frequency tuning curve extracted from a schematic simulation (The actual frequency of the VCO will be lower due to parasitic capacitance.)

narrow analog tuning range, where the calibration phase steps can move the PLL's control voltages sufficiently far from their small signal state so as to corrupt the measurement.

In this work we introduce a dual path analog PLL which, in conjunction with an integral path calibration scheme desensitizes the PLL's small signal transfer function to gain variations of the non-linear components of the loop. By implementing the PLL with separate integral and proportional paths, the gain of the proportional path is stabilized in a fashion similar to [4], without the need for a type III loop control. The integral path gain, which remains sensitive to VCO non-linearity is controlled with a novel calibration scheme. The scheme works by measuring the response of the PLL to opposite polarity phase steps while temporarily disabling the proportional path (i.e., setting $\zeta = 0$). This scheme avoids saturation problems on the proportional path, and works well even if the PLL has a limited linear range and/or has significant input-referred phase offsets. The dual path design and associated calibration scheme leads to a significant reduction in PLL phase transfer function variation.

This paper is organized as follows: In Section II we will review the benefits of the type II dual loop PLL. In Section III we will provide the block level details of the PLL's core components, and in Section IV we will present the details of a novel integral path PLL calibration system. In Section V we present measurement results

II. Dual-Path PLL Review

The PLL in this work was implemented in a dual-path configuration, as shown in Fig. 1(c). This architecture requires two charge pumps and a second set of varactors in the VCO, or, alternately, a method of summing the two sets of control voltages before the signal is applied to the VCO. This architecture has some distinct advantages compared with a conventional single path PLL which justify its relative increase in complexity. In this section we will review the benefits of the dual-path architecture, contrasting it with single path PLL and with dual path type III PLL architectures.

In a dual path PLL the control loop is split into separate integrating and proportional paths where the integrating path consists of a charge pump and loop capacitor $(\mathrm{C_i})$, and the proportional path consists of a charge pump, resistor $(\mathrm{R_p})$ and ripple capacitor $(\mathrm{C_r})$; a general discussion of an example of this architecture is provided in [8]. Variants of this architecture include a PLL with a resistor-less sample-reset proportional path, as presented in [9]. A significant advantage of this architecture is that the gain of the integrating and proportional paths can be set independently by changing the relative magnitudes of the charge pump currents. This feature allows for easier programmability of the transfer function of the PLL.

A. Sensitivity to VCO Non-Linearity

In a dual path PLL, the small signal gain of the proportional path is significantly desensitized to the non-linearity of the VCO. The proportional path has no long term memory; once the PLL has achieved lock the proportional path operates around its common-mode value, corresponding to a single point on the VCO transfer function. If the PLL locks to a new frequency, only the integral path control voltage must move to a new point on the VCO tuning curve. This separation also reduces the variation of the higher order pole as the voltage across the potentially non-linear ripple capacitor will also remain constant irrespective of the integral path voltage. In addition, the range of voltage over which the proportional path charge pump must be linear is reduced as the proportional path is decoupled from frequency tuning.

However, the integral path gain in the dual loop architecture still suffers from the potential non-linearity of the VCO and capacitor, resulting in uncertainty in the position of the open loop zero of the PLL. This uncertainty motivates the search for an integral path gain calibration scheme, as will be described in Section IV.

B. Noise in a Dual Path PLL

A further advantage of the dual path architecture over the single-path approach is that in dual-path designs it is easier to reduce the phase noise contribution from the charge pumps. In a high performance communication system, many of the PLL's parameters, such as reference frequency and division ratio, are fixed by the application, or, in the case of loop gain, are set to optimize the tradeoff between filtering reference noise and VCO noise. In a conventional single path PLL with a fixed VCO gain, $\mathrm{K_{vco}}$, this tradeoff fixes the ratio of charge pump current to loop filter impedance; in this case, only the absolute value of charge pump current and loop filter impedance can be modified. In particular, if the charge pump current is increased then the loop filter impedance should be decreased by the same amount to maintain the same transfer function. As shown in [10], the low frequency in-band phase noise contribution of the charge pump in a single path PLL is given by the following:

$$\overline{\Delta\phi^2_{\mathrm{out}}} = 4\pi^2 \frac{2\bar{I}^2_n}{I^2_P} \frac{2NT_{RST}}{T_{\mathrm{CLK}}} \tag{1}$$

In this equation, $\mathrm{I_n}$ is the noise from the charge pump, $\mathrm{I_p}$ is the charge pump current, N is the PLL's multiplication factor, $\mathrm{T_{RST}}$ is the reset time of the PFD, and $\mathrm{T_{CLK}}$ is the period of the reference clock. According to (1), there are few options available to reduce the charge pump noise contribution. $\mathrm{T_{RST}}$ generally cannot be reduced below a few gate delays to avoid dead-zones in the charge pump. N and $\mathrm{T_{CLK}}$ are fixed by the application. The sizes of the charge pump current source transistors can be optimized to minimize $\mathrm{I_n}$; further reductions in phase noise can only be achieved by increasing $\mathrm{I_p}$. Note that $\mathrm{I_n}$ is a function of $\mathrm{I_p}$; increasing $\mathrm{I_p}$ will also increase $\mathrm{I_n}$. If the charge pump current is doubled by switching in a second identical charge pump then $\mathrm{I_n}$ will increases by a factor of sqrt(2), leading to a 2x reduction of $\mathrm{I_n^2/I_p^2}$. In order to maintain the same loop transfer function, an increase in $\mathrm{I_p}$ must be compensated with a proportional reduction in the loop filter impedance. This leads to an unfortunate trade-off: improving the noise performance requires a reduction in loop filter impedance, which results in an expensive increase in the area of the loop capacitor. Note that the proportional path resistor must also be reduced, although this requirement is typically less problematic in terms of area.

Equation (1) provides the contribution to phase noise below the PLL bandwidth, and is derived assuming that charge pump current noise will be tracked by the PLL. In order to see the contribution to phase noise at frequencies at or above the bandwidth of the PLL, the derivation of (1) can be extended for single and dual paths PLL, to include the filtering effects of the PLL. The charge pump noise $\mathrm{I_n}$ flows into the loop filter for time $\mathrm{T_{RST}}$ of every reference period $\mathrm{T_{CK}}$. Consequently the net noise current flowing into the loop filter must include the harmonics mixed down from high frequency and scaled by $\mathrm{T_{RST}/T_{CK}}$. As shown in [10] the effective noise current from the charge pump is given by:

$$I^2_{cp\_noise} = 4\bar{I}^2_n \frac{T_{RST}}{T_{\mathrm{CLK}}} \tag{2}$$

This current noise is converted to a voltage in the loop filter which is then converted to phase noise by the PLL. It is straightforward to show that in a conventional type II second order PLL the resulting output phase noise is given by:

$$\overline{\phi^2_{\mathrm{out}}} = I^2_{cp\_noise} \left(\frac{\left(R_p + \frac{1}{sC_i}\right) 2\pi K_{vco}}{s + I_p \left(R_p + \frac{1}{sC_i}\right) K_{vco}/N} \right)^2 \tag{3}$$

where $\mathrm{K_{vco}}$ is the VCO gain in Hz/V, $\mathrm{R_p}$ is the proportional path resistance, $\mathrm{C_i}$ is the integrating capacitor, $\mathrm{I_p}$ is the charge pump current and N is the division ratio. As $\mathrm{s} \to 0$, (3) reduces to (1), in agreement with [10]. In a dual path PLL, there are two charge pumps with two noise currents: one from the proportional charge pump $\mathrm{I_{cpp\_noise}}$ and one in the integral path, $\mathrm{I_{cpi\_noise}}$. It is straightforward to show that the phase noise contribution in a type II dual path PLL from the charge pumps is given by

$$\overline{\phi^2_{\mathrm{out}}} = \bar{I}^2_{cpi\_noise} \left(\frac{\frac{1}{sC_i} 2\pi K_{vcoi}}{s + \left(I_{pp} R_p K_{vcop} + I_i \frac{1}{sC_i} K_{vcoi}\right) \frac{1}{N}} \right)^2 + \bar{I}^2_{cpp\_noise} \left(\frac{R_p 2\pi K_{vcop}}{s + \left(I_{pp} R_p K_{vcop} + I_i \frac{1}{sC_i} K_{vcoi}\right) \frac{1}{N}} \right)^2, \tag{4}$$

where $\mathrm{I_{cpi\_noise}}, \mathrm{I_{cpp\_noise}}$ are the contributions from the integral and proportional path charge pumps, $\mathrm{I_i}, I_{\mathrm{pp}}$ are the integral and proportional currents and $\mathrm{K_{vcop}}, \mathrm{K_{vcoi}}$ are the proportional path and integral path VCO gains. If we compare the noise contribution from the integral path charge pump current (the first term in (4)) to the charge pump noise in a single path PLL (3), we see that the contribution is the same up to the open loop zero of the PLL. However, for frequencies above the open loop zero, ω_{z}, the phase noise contribution is reduced by a factor of $|(\mathrm{s} + \omega_{\mathrm{z}})/\omega_{\mathrm{z}}|$, as illustrated in Fig. 3. This characteristic mitigates the problem of requiring an increase in capacitor area to reduce charge pump noise, at least above the first zero of the PLL. This advantage can be understood intuitively as follows: the noise from the integral path charge pump is converted to a voltage by the impedance at its output. In a dual path PLL, the integral path filter's impedance is just a capacitor, as shown in Fig. 3(b), in contrast to a resistor in series with a capacitor in a single path PLL, in Fig. 3(a). Therefore, the voltage on the integral path control voltage is less than that the same charge pump would produce in a single path PLL, resulting in lower noise at the PLL output.

The phase noise contribution from the proportional path charge-pump is similar to that of a single path PLL above the open loop zero. However, now that the proportional path is decoupled from the integral path, it is no longer required to support a wide frequency tuning range through the proportional path. By decreasing the VCO gain and increasing the charge pump current, the phase noise contribution of the charge pump

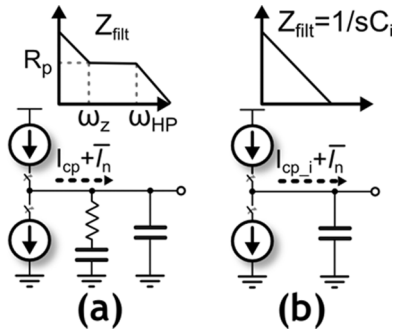

Fig. 3. (a) In a conventional PLL, the loop filter impedance includes a resistor and ripple capacitor. (b) In the dual path PLL, the integral path filter is just a capacitor.

is reduced without affecting the bandwidth, in a similar manner as described in [4]. While the approach offers significant advantages, there are practical limitations which should be kept in mind when employing this strategy in a multi-path PLL. In particular, leakage on the integral path capacitor or mismatch between the integral charge pump up/down currents will cause a phase offset at the input of the PLL, which will in turn cause the proportional path voltage to move away from its common-mode point. A static phase offset of $\Delta\theta$ will cause a DC shift in the proportional control voltage of $\Delta\theta/(2\mathrm{pi}) * \mathrm{I_{pp}} * \mathrm{R_p}$. The strategy outlined above of reducing $\mathrm{K_{vcop}}$, and increasing $\mathrm{I_{pp}}$ will increase the voltage offset caused by this phase offset. This consideration effectively puts a limit on the degree to which the proportional path VCO gain should be reduced. Nonetheless, provided that the phase offsets caused by the non-idealities in the integral path are modest, the proportional path VCO gain $\mathrm{K_{vcop}}$ can be reduced by at least an order of magnitude as compared with the integral path.

It is worth noting that the offset caused by capacitor leakage is potentially more problematic in the Type III PLL of Fig. 1(b). The gain, $\mathrm{g_m}/(\mathrm{s} * \mathrm{C_3})$, of the extra path must be made extremely small in order to make the frequency offset of the resulting zero small. Consequently, leakage in the capacitor, $\mathrm{I_{leak}}$ must be compensated with a change in the fine control voltage of $\mathrm{I_{leak}}/\mathrm{g_m}$. This offset is independent of the size of the capacitor; reducing the size of the capacitor reduces the amount of leakage, but $\mathrm{g_m}$ must also be proportionally reduced to maintain the same gain. This is a significant disadvantage of the type III PLL: it is potentially very sensitive to capacitor leakage.

In summary, the dual loop type II architecture is an attractive PLL topology as it leads to a stabilization of the proportional path gain, and makes it somewhat easier to deal with charge pump noise. In addition, it avoids the leakage sensitivity of type III PLLs. However the dual loop topology still suffers from gain variation in the integral path due to VCO and capacitor nonlinearity, which motivates the requirement for a new integral path calibration scheme, as described in Section IV.

III. Implementation Details

A block diagram of the PLL is shown in Fig. 4. The CMOS feedback divider consists of a divide by 16 prescaler, followed

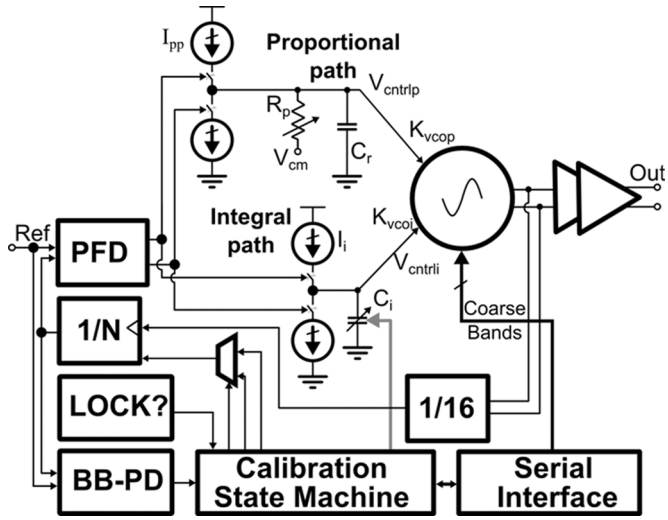

Fig. 4. Architecture of the PLL including circuitry required for integral path calibration.

by a digitally programmable divider. The dual-loop PLL includes separate fully differential proportional and integral control paths, the details of which are shown in Fig. 5. The differential integral path consists of 10 10 uA current slices each of which can be independently enabled. The loop capacitor includes a total of 217 pF implemented with a thick oxide capacitor over n-well, and segmented into a 5-bit binary weighted differential configuration. The output nodes also include capacitance to ground, as part of the stabilization for the common-mode feedback loop.

The proportional path is also implemented differentially. The charge pump consists of 15 50 uA independently controllable current slices and the loop resistor is segmented into a 4-bit binary weighted structure. The minimum required programmable range of the loop components must be sufficiently large so as to be capable of compensating for PVT variation in an application using a fixed reference clock; in this work the programmable range of charge pump currents and loop filter components is made very wide so as to enable the support of a wide range of bandwidths, applications and reference clock rates. The switches used to make the loop capacitor programmable are implemented with thick oxide transistors; the gate leakage of thin-oxide transistors was found to be sufficiently large to contribute to capacitor droop and reference spurs. The common mode of the proportional path is set by tying the mid-point of the proportional resistor to a fixed voltage, nominally set to half of the supply. This voltage, Vcm in Fig. 1(c), is generated with a simple resistor divider. In between UP/DWN pulses from the phase detector, a feedback loop on each side of the proportional path charge pump adjusts the up current to match the down current. Both paths are controlled with a classical tri-state phase detector.

The VCO architecture is based on the Colpitts noise shifting VCO described in [12]. This architecture provides the strong phase noise performance of the classical Colpitts configuration without degrading the oscillator's ability to start up robustly, even at K-band frequencies. A block diagram of the oscillator is

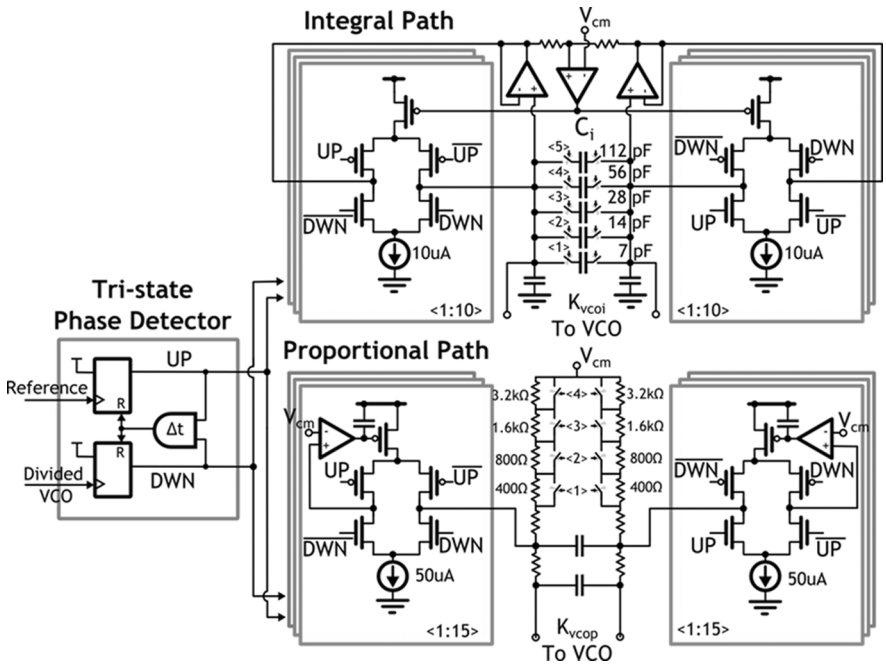

Fig. 5. Details of the fully differential charge pumps and loop filters.

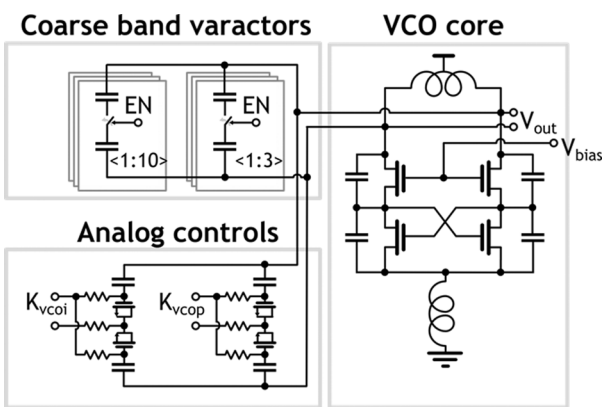

Fig. 6. Noise shifting Colpitts VCO including two sets of analog varactors for the proportional and integral paths. There are also two banks of digitally switched capacitors.

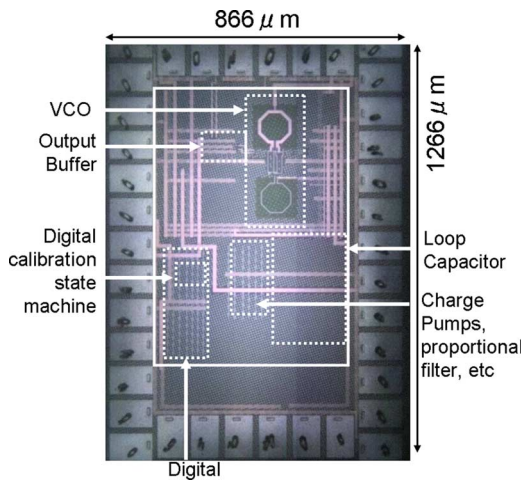

Fig. 7. Chip photo.

shown in Fig. 6; additional implementation details can be found in [13]. The proportional and integral analog controls required for the dual-path PLL utilize accumulation varactors which are isolated from the tank with back-end-of-the-line (BEOL) capacitors. The varactor structure is implemented in a symmetrical configuration so as to provide good common-mode rejection, even at the edges of the tuning band. The nominal integral and proportional path gains are approximately 400 MHz/V and 40 MHz/V, respectively. The majority of the VCO's tuning range is achieved using two sets of digitally controllable capacitor banks.

A chip photo is shown in Fig. 7. The state machine required for integral path healing (in Section IV) is embedded in the digital logic, and consumes a small fraction of the overall area (40 μm $\times$ 60 μm).

IV. Integral Path Healing Scheme

As discussed in Section II, in a dual-loop type II PLL, the integral path gain suffers from gain variation due to the non-linearity of both the VCO and the integral path capacitor. This integral path variation affects the position of the zero in the PLL's open loop response. The expected effects of changing the integral path gain on the closed loop PLL transfer function are shown in Fig. 8(a), calculated using a simple linear model. It is instructive to compare these results with measurements from hardware when the integral path voltage is at various locations on the transfer curve. Fig. 8(b) shows the closed loop response of the PLL, as measured from hardware; the methodology for measuring transfer function will be discussed in Section V. In this plot the closed loop phase transfer function of the PLL is measured at nine different points in the same VCO tuning band,

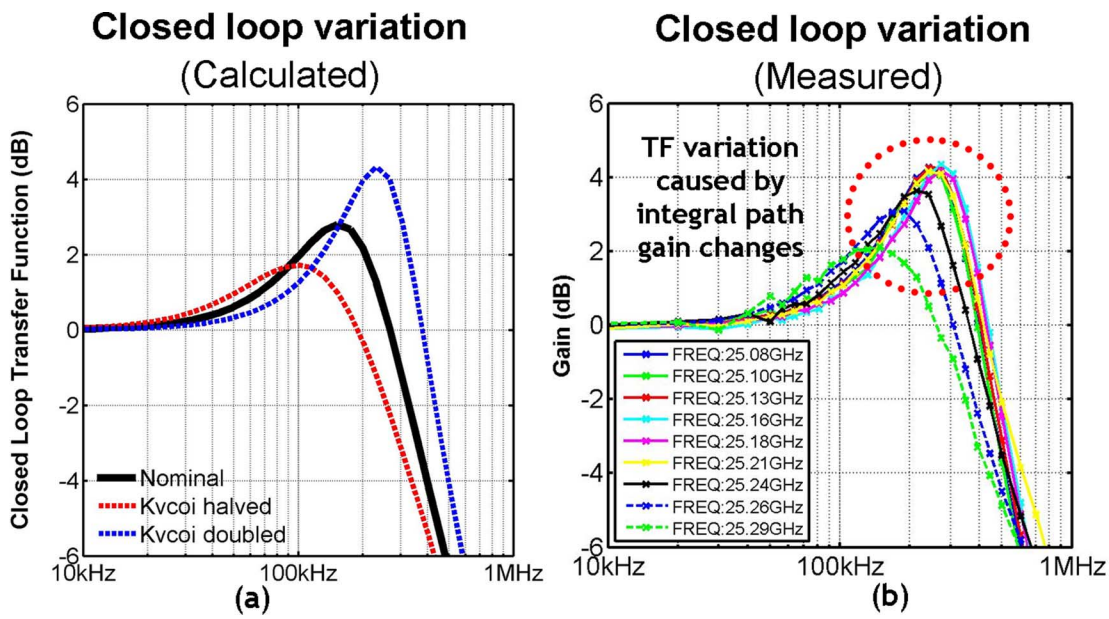

Fig. 8. Closed loop transfer function variation due to integral path gain changes. (a) Calculated using simple linear model. (b) Measured from hardware at 9 different points in a single coarse frequency band.

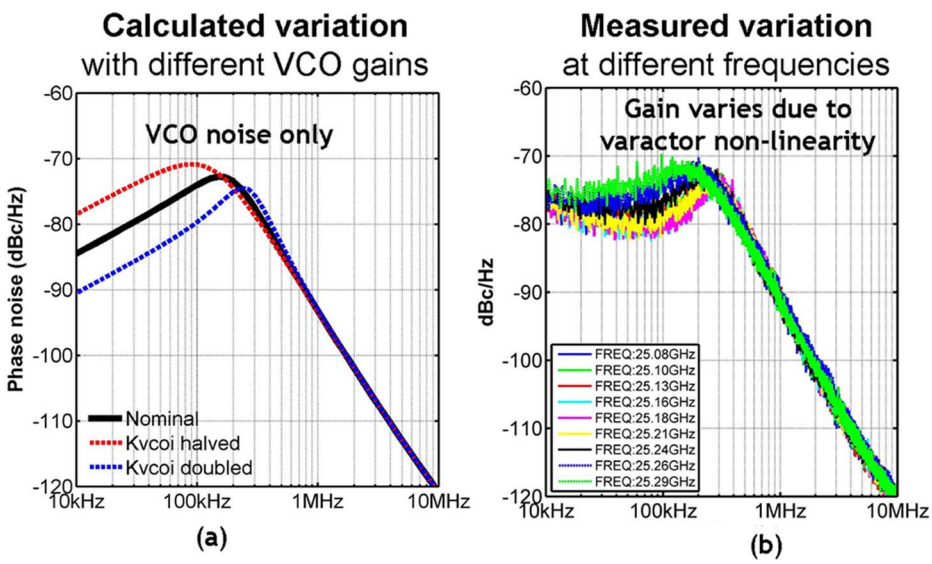

Fig. 9. Effect of changing the integral path gain on low frequency phase noise. (a) Calculated. (b) Measured from hardware.

equispaced from the bottom of the band to the top of the band. Therefore each of these measurements corresponds to a different point on the tuning curve of the VCO, leading to a change in the gain of the integral path. Significant variation in the amount of jitter peaking is observed, similar to the calculated case, despite the fact that in absolute terms the set of measured operating frequencies shown in the plot are relatively close together.

Jitter peaking can be reduced by reducing the integral path gain relative to the proportional path. However, the integral path gain also plays a significant role in suppressing the close to carrier phase noise of the VCO. In Fig. 9(a) the effect on phase noise of changing just the integral path gain is calculated. The integral path gain plays a significant role is suppressing low frequency VCO phase noise, in particular if the VCO noise is dominated by flicker noise. In Fig. 9(b), the PLL's phase noise is measured at the same frequencies as in Fig. 8(b). As can be seen in the figure, the measured phase noise is lowest at frequencies in the center of the tuning band, where the integral path gain is highest.

If the integral path gain is too high, the result is excessive jitter peaking in the PLL's transfer function. If the integral path gain is too low, then the low frequency phase noise of the VCO is insufficiently suppressed. The nonlinearity of the VCO results in uncertainty in the integral path gain. In this work we will present a method of precisely calibrating the integral path gain, enabling a more precise tradeoff between jitter peaking and VCO noise suppressions.

In [7], a foreground calibration technique has been proposed which has been demonstrated to be capable of correcting a PLL's transfer function in the presence of manufacturing variation.

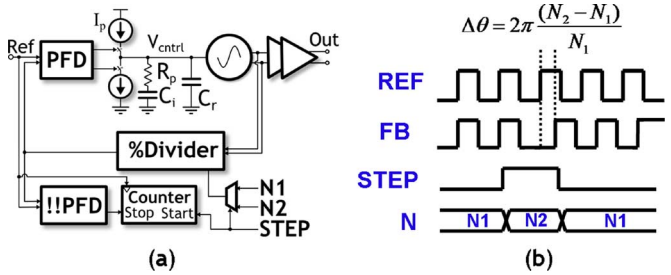

Fig. 10. Overview of time till crossover measurement. (a) Circuits used to add phase step and count time until crossover. (b) The phase step is added by temporarily changing the divider's count value from N1 to N2.

The additional circuitry required to implement this scheme is of a digital nature, resulting in a minimal area penalty. In addition, the method measures the PLL's small signal gain terms around its phase locked state, which means that it accounts for the small signal gain variation of non-linear components. The scheme works as follows: A phase step is added to the PLL by temporarily incrementing or decrementing the feedback divider's count value, as illustrated in Fig. 10. A bang-bang phase detector, placed in parallel with the PLL's phase detector, measures the time (in reference cycles) for the phase of the PLL to cross zero in response to the introduced phase step. This time-until-cross-over measurement, in conjunction with an over-shoot detector is used to calibrate the PLL's transfer function. As discussed in Section II, however, the PLL in the work described here features a limited frequency range on the proportional path, making the scheme proposed in [7] impractical.

Specifically, in contrast to the design described in [7], adding large phase steps to a PLL with a limited proportional path tuning range, causes the proportional control voltage to move from its nominal position significantly, resulting in a measurement that is corrupted by non-linear effects. If, instead, a very small phase step is used to avoid these non-linearities, the measurement is corrupted by phase offsets.

The transient step response of the PLL's phase, according to [11], is given by:

$$\phi(t) = \Delta\theta \left(\cosh(\sqrt{\zeta^2 - 1}\omega_n t) - \frac{\zeta}{\sqrt{\zeta^2 - 1}} \sinh(\sqrt{\zeta^2 - 1}\omega_n t) \right) \exp(-\omega_n \zeta t), \quad (5)$$

where the size of the phase step is $\Delta\theta$, and the natural frequency and damping ratio are given by ω_n and ζ. According to [7], this equation yields an expression for time to crossover for an over-damped PLL given by:

$$t_1 = \frac{1}{\omega_n \sqrt{\zeta^2 - 1}} \operatorname{arctanh}\left(\frac{\sqrt{\zeta^2 - 1}}{\zeta} \right). \quad (6)$$

In Fig. 11 the step response of an ideal second order type II PLL is shown, where the X-axis corresponds to time in reference cycles and the Y-axis corresponds to phase difference. (Practical PLLs are at least third order, but for now we will ignore the higher order pole). If the PLL is linear and offset-free, then the time until cross-over, t_1, is independent of the size of

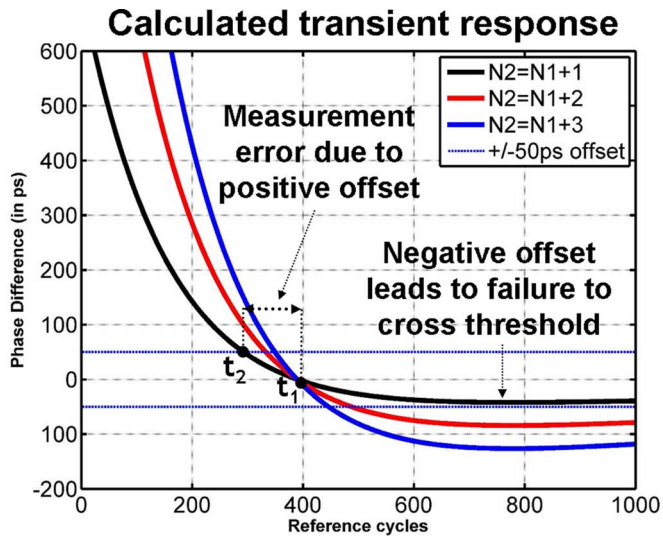

Fig. 11. Ideal response of a second order type II PLL to a phase step. The magnitude of the phase step does not change the time to crossover.

the phase step. The significance of the time t_1 is that it represents the time when the output of the bang-bang phase detector will switch polarity. This time can be easily measured with a counter, where the counter is clocked using the reference clock. However, analog PLLs may have input referred phase offsets for multiple reasons, including charge pump current mismatch and capacitor leakage; there may also be a static mismatch between the PLL's PFD and the BB-PFD that would yield an effective input referred offset. The time-until-cross-over measurement is corrupted by offsets from any of these sources; the actual time measured if the offset is positive offset is shown as t_2 in Fig. 11. In fact, if the offset is negative and larger than the overshoot, then the phase may actually never cross zero, resulting in an unbounded error.

The error caused by the PLL's offset can be reduced by increasing the size of the added phase step. As also can be seen in Fig. 11, the larger the phase step, the smaller the error due to the offset. However, the phase step in (5) causes the control voltage phase of the proportional path to move away from its common-mode point. As discussed in Section II, in order to minimize the phase noise contribution of the proportional path charge pump, $I_{pp} * R_p$ has been increased and K_{vcop} has been decreased. This strategy leads to an increase in the voltage excitation on the proportional path in response to phase steps. In order for the bandwidth calibration to be effective, any transient changes to the control voltages of the PLL must be sufficiently small so that the nature of the response is still small-signal. However, the reduction in proportional path gain (as described in Section II) makes this result difficult to achieve. In summary, if the added phase step is too small, then phase offsets of the PLL can corrupt the results. If the phase step is too large, then the time to crossover will not reflect small signal behavior of the system; instead, it will be a function of the non-linear elements in the PLL. The practical implications of the interaction between introduced phase step size and PLL characteristics can be seen in both simulation, Fig. 12(a), and hardware measurements from an early PLL prototype, Fig. 12(b). In both

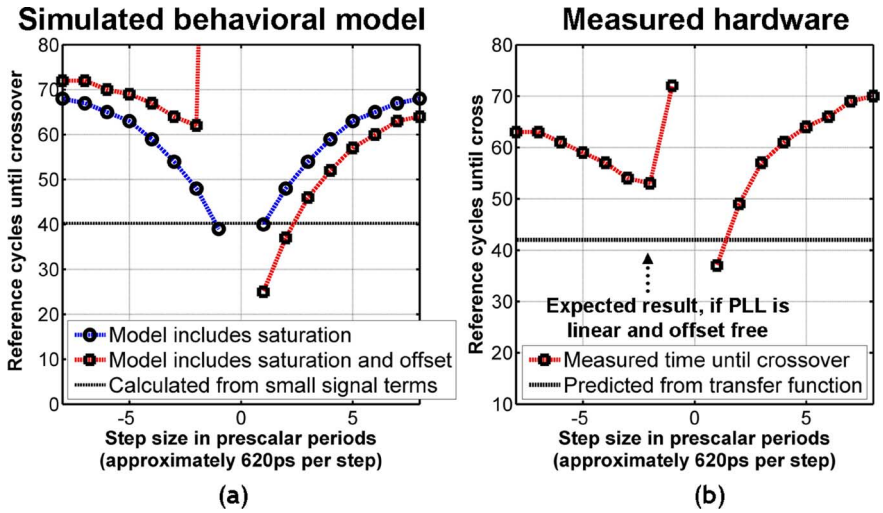

Fig. 12. Time-to-crossover measured as a function of the size of the added phase step. (a) Behavioral simulation results including non-linearity, with/without phase offsets. (b) Results measured from hardware, and results predicted from the PLL's small signal transfer function.

Fig. 12(a) and (b) the X-axis describes the magnitude of the added phase step in prescaler periods (one prescaler period = 16 output periods), and the Y-axis the time to crossover measured in reference cycles and recorded by the integrated state machine. Fig. 12(a) shows the time-to-crossover for three cases: 1) in the ideal case (i.e., calculated from the PLL's small signal parameters), 2) when the behavioral model includes non-linearity in the proportional path, and 3) when the model includes both non-linearity and phase offsets in the form of leakage in the loop capacitor. The simulated results agree with the ideal value only when offsets are not included and a minimum sized phase step is used, as can be seen in Fig. 12(a). When the size of the phase step is increased, or both non-linearites and phase offsets are modeled, then the simulated values never agree with the ideal value. The same experiment performed in hardware, as shown in Fig. 12(b), displays the same characteristics as the simulated results; the results are corrupted by non-linearity when large phase steps are used and by offsets when small phase steps are used.

As discussed in Section II, in the dual path PLL presented in this paper, the proportional path gain is quite insensitive to the nonlinearity of the VCO. The integral path gain, on the other hand, remains susceptible to variation due the VCO's non-linearity. Consequently, we introduce a scheme which calibrates just the integral path, and avoids the errors caused by offsets and proportional path non-linearity described above. The scheme works based on the principle that if a phase step is added to a PLL with the proportional path temporarily disabled ($\zeta = 0$) then the PLL's phase response will be

$$\phi(t) = \Delta\theta \cos(\omega_n t) \tag{7}$$

This periodic response, illustrated in Fig. 13, has some useful properties. Firstly, provided that the phase step is greater than the PLL phase offset, the phase will always cross the measurement threshold. Secondly, the cosine response is symmetrical around zero. As a result, if the polarity of the phase step is inverted, then the error caused by the offset will also be inverted

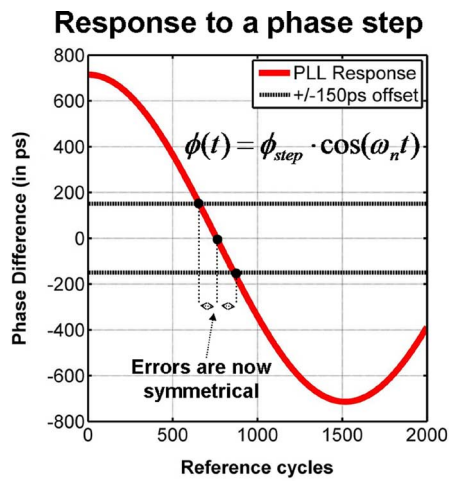

Fig. 13. Response of the PLL to a phase step when the proportional path is disabled.

but will maintain the same magnitude. Consequently, if the time to crossover measurement is performed twice but using opposite polarities of the introduced phase step, then the average of these two measurements will correspond to the offset-free time to crossover. It is straightforward to show that the average time till first cross-over with opposite polarity steps is simply $0.5 * \pi/\omega_{\rm n}$, with the proportional path disabled.

The measurement is performed as follows. First, the PLL is allowed to fully phase lock with an initial default set of charge pump and loop filter settings (i.e., with the proportional path enabled). Once the integrated lock detector detects lock, a phase step is injected into the loop while simultaneously disabling the proportional path. While the proportional path is required if the PLL is to be stable, the proportional path can be disabled for short durations without causing unlimited oscillations. In this scheme the proportional path is only disabled until the phase crosses zero, which is less than the natural period $(2 * \pi/\omega_{\rm n})$ of the PLL. Once the PLL's input phase crosses the measurement threshold, the number of reference cycles since the phase

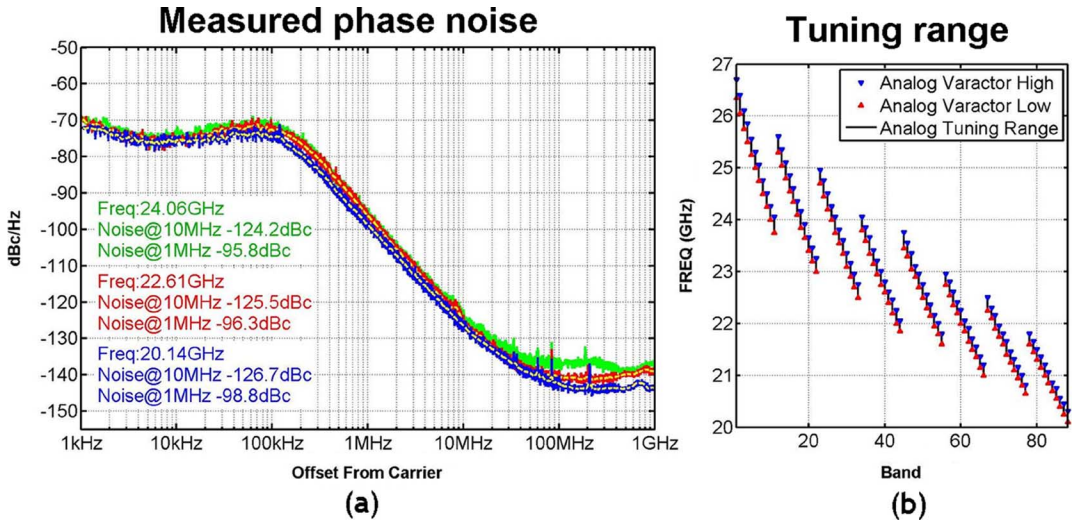

Fig. 14. The PLL's noise and tuning range performance. (a) Phase noise at multiple frequencies. (b) The PLL's tuning range is dominated by the digitally controlled capacitors.

step was introduced is stored, and the proportional path is re-enabled. Next, the PLL is allowed to relock and the measurement is repeated, but with the application of a phase step of opposite polarity. An integrated state machine averages the two time-to-crossover measurements, giving the offset-free time-to-crossover measurement, corresponding to the natural frequency of the PLL. As the effects of the phase offsets have been accounted for, a small step size can be used, which avoids the non-linearites of the VCO. Furthermore, it can be shown that magnitude of the voltage excitation in response to a phase step is much smaller on the integral path than it would be on the proportional path, which further eases the linearity requirements.

Note that through this approach, the effects of the higher order pole due to the ripple capacitor on the proportional path have been avoided, as the proportional path has been disabled. In contrast, a calibration scheme based on (6) does not include the effects of the ripple capacitor, which could lead to errors if the higher order pole is close to the crossover frequency of the PLL.

This method of measuring $\omega_{\rm n}$ was implemented as part of a full integrated integral path calibration state machine. In the prototype PLL both the integral path charge pump and loop capacitor are programmable, the charge pump is thermometer weighted, and the capacitor is binary weighted (details in Section III). Both are adjusted in the calibration scheme.

The algorithm operates as follows: Firstly, the PLL locks normally with the integral path charge pump current and the integral path capacitor set to their maximum value. Next, the average time until crossover is measured, using the technique described above. The measurement is compared with a value loaded through the serial interface, corresponding to a required integral path gain. If the result is less than the target, then the response was too fast, and the integral path charge pump current is decremented. This procedure is repeated until the measured result is greater than the target. Next, size of the capacitor is decremented, which increases the integral path gain and decreases the time until crossover measurement. Once the time to crossover measurement is equal to or less than the target, the algorithm is complete. In the current implementation the calibration is intended to occur on power up, and nominally takes approximately 1.5 ms to complete. The calibration time is dominated by the time it takes the lock detector to declare lock, which occurs when the input to the phase detector remains less than 40 gate delays (approximately 800 ps worst case) for 2^{11} reference clock cycles. In order to ensure the PLL is locked between phase measurements, the lock detector response time was made deliberately long. The calibration time could easily be reduced by reducing the lock detector counter, or by using a binary search algorithm. In this work the calibration is performed on power-up only. Subsequent temperature changes will result in changes to the center frequency of the VCO. The error caused by temperature changes depends on the temperature sensitivity of the VCO; in simulation a temperature change from 25 C to 85 C results in the VCO frequency moving by 17% of the range of a tuning band.

As both the integral path charge pump current and the loop capacitor are programmable, there are potentially multiple combinations that will produce the required ratio of charge-pump current to loop filter impedance and hence the same integral path gain. As discussed in Section II, for the same gain, however, using the largest current and the largest capacitor size will produce the best noise performance. Note that the algorithm described here will settle on that best noise performance solution.

V. Measurements

The realized PLL locks over a frequency range of 20.1 to 26.7 GHz and consumes approximately 33 mW (depending on VCO bias settings and frequency), excluding input and output buffers; Fig. 14 shows the PLL's measured phase noise performance at several output frequencies and the PLL's measured tuning range. At 20.14 GHz, close to the bottom of the tuning range, the measured phase noise was -126.7 dBc/Hz at a 10 MHz offset from the carrier. The phase noise gradually degrades as the output frequency increases.

In order to demonstrate the effectiveness of the integral path self-calibration system, a method of accurately measuring the PLL's phase transfer function is required. In this work, a method

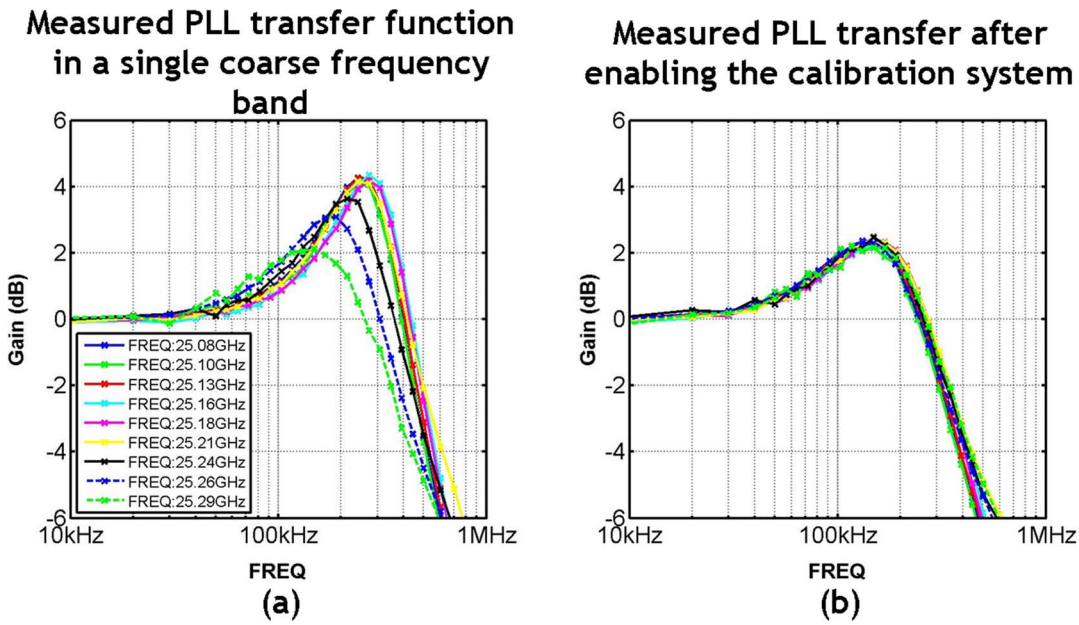

Fig. 15. Measured closed loop phase transfer function in a single frequency band. (a) Before enabling the calibration engine. (b) After enabling it.

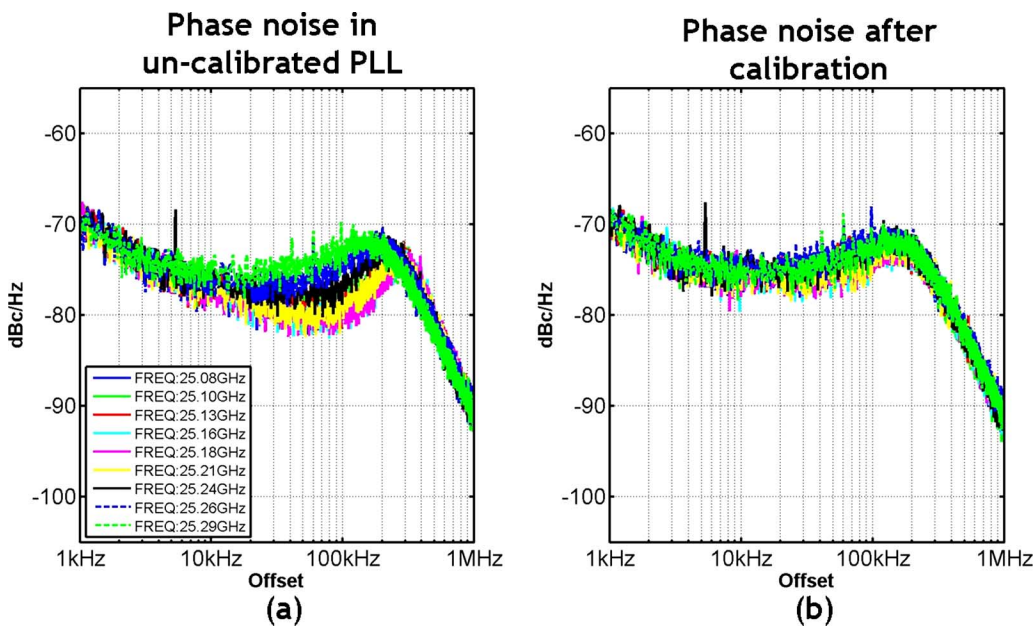

Fig. 16. Measured phase noise in a single frequency band. (a) Before enabling the calibration engine. (b) After enabling it.

similar to that described in [7] was used. First the reference clock is phase modulated at a fixed offset frequency. This introduces a tone in the output spectrum of the PLL. The magnitude of this tone is measured using a spectrum analyzer, and the ratio of the magnitude of the output tone to the reference clock's modulation index can be used to calculate the magnitude of PLL's phase transfer function at a given offset frequency. Next, the experiment is repeated at several modulation frequencies and the results are stitched together to estimate the PLL's small signal closed loop transfer function. This experiment requires careful calibration of the amplitude of the modulation tone; we have found that several commercially available clock generators with programmable phase modulation features filter the modulation tone with their own internal PLLs. Such non-idealities must be accounted for if the PLL's transfer function is to be accurately measured.

The PLL's loop transfer function, shown in Fig. 15, and in-band phase noise, shown in Fig. 16, were measured at 9 different frequencies, equally spaced from fastest to slowest in a single coarse tuning band. On the left side of these figures, the measurements without self-calibration are presented, showing significant variation in both phase noise and transfer function. On the right side of these figures are the transfer functions and phase noise measurements after self-calibration. As can be seen, a dramatic reduction in the variation of both loop transfer function and in-band phase noise is observed. In a second test

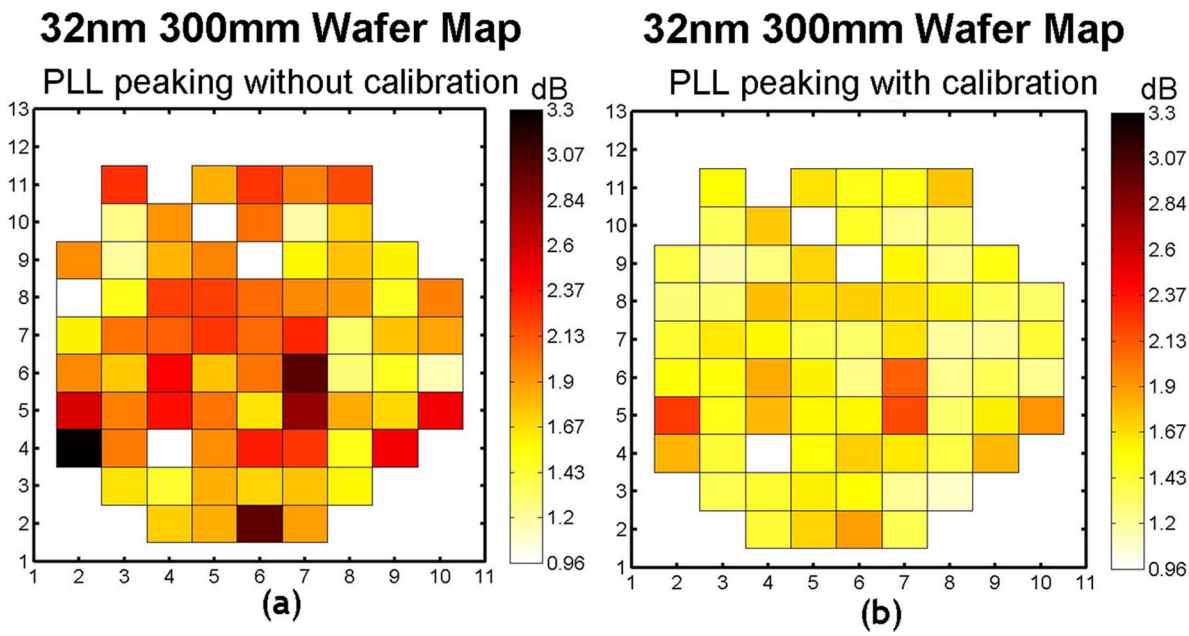

Fig. 17. Jitter peaking measured across a 300 mm wafer for every functional die at 25 GHz (a) before enabling the calibration engine, and (b) after enabling the calibration engine.

of the integral path self-calibrating PLL, the 3 dB bandwidth and jitter peaking were measured at an output frequency of 25 GHz, with a division ratio of 128 (reference frequency of 195.3125 MHz) across an entire 300 mm wafer. The results of this test, plotted in Fig. 17, show that the spread of jitter peaking is reduced from 2.4 dB to 1 dB through introduction of the calibration technique.

VI. Conclusion

In this work we have demonstrated that a dual path type II analog PLL is a suitable choice for implementing a fully integrated high frequency, low noise PLL. In addition, we have demonstrated a new foreground calibration scheme that can be used to calibrate the PLL's integral path gain in order to compensate for the non-linearity of the VCO and loop capacitor. This calibration scheme has been demonstrated to work well even in the presence of significant non-idealities within the loop, including a limited range proportional path, and phase offsets at the input of the PLL.

Acknowledgment

This work has been partially funded by DARPA under AFRL contract # FA8650-09-C-7924. The views, opinions, and/or findings contained in this presentation are those of the author/presenter and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.

References

[1] B. A. Floyd, "A 16–18.8 GHz sub-integer-N frequency synthesizer for 60 GHz transceivers," *IEEE J. Solid-State Circuits*, vol. 43, no. 5, pp. 1076–1086, May 2008.

[2] J. Bulzacchelli *et al.*, "A 28 Gb/s 4-tap FFE/15-tap DFE serial link transceiver in 32 nm SOI CMOS technology," in *IEEE International Solid State Circuits Conference*, 2012.

[3] S. Williams, H. Thompson, M. Hufford, and E. Naviasky, "An improved CMOS ring oscillator PLL with less than 4 ps RMS accumulated jitter," in *Proc. IEEE Custom Integrated Circuits Conf.*, Sept. 2004, pp. 151–154.

[4] R. Nonis, N. Dalt, P. Palestri, and L. Selmi, "Modeling, design and characterization of a new low-jitter analog dual tuning LC-VCO PLL architecture," *IEEE J. Solid-State Circuits*, vol. 40, pp. 1303–1308, June 2005.

[5] W. Rhee, H. Ainspan, D. Friedman, T. Rasmus, S. Garvin, and C. Cranford, "A uniform bandwidth PLL using a continuously tunable single-input dual-path LC VCO for 5 Gb/s PCI express Gen2 application," in *IEEE Asian Solid-State Circuits Conf.*, Nov. 2007, pp. 63–66.

[6] T. Wu, P. Hanumolu, K. Mayaram, and U. Moon, "Method for a constant loop bandwidth in LC-VCO PLL frequency synthesizers," *IEEE J. Solid-State Circuits*, vol. 44, no. 2, pp. 427–435, Feb. 2009.

[7] D. M. Fishette, A. L. S. Loke, R. J. DeSantis, and G. R. Talbot, "An Embedded All-Digital Circuit to Measure PLL Response," *IEEE J. Solid-State Circuits*, vol. 45, no. 8, pp. 1492–1503, Aug. 2010.

[8] J. Craninckx and M. Steyaert, "A fully integrated CMOS DCS-1800 frequency synthesizer," *IEEE J. Solid-State Circuits*, vol. 33, no. 12, pp. 2054–2065, Dec. 1998.

[9] A. Loke, R. Barnes, T. Wee, M. Oshima, C. Moore, R. Kennedy, and M. Gilsdorf, "A versatile 90-nm CMOS charge-pump PLL for SerDes transmitter clocking," *IEEE J. of Solid-State Circuits*, vol. 41, pp. 1894–1907, Aug. 2006.

[10] B. Razavi, "Phase-locking in wireline systems: Present and future," in *Proc. IEEE Custom Integr. Circuits Conf.*, Sept. 2008, pp. 615–622.

[11] F. M. Gardner, *Phaselock Techniques*, 3rd ed. Hoboken, NJ: Wiley, 2005, pp. 102–103.

[12] R. Aparicio and A. Hajimiri, "A noise-shifting differential colpitts VCO," *IEEE J. Solid-State Circuits*, vol. 37, Dec. 2002.

[13] J.-O. Plouchart, M. Ferriss, A. Natarajan, A. Valdes-Garcia, B. Sadhu, A. Rylyakov, B. Parker, M. Beakes, A. Babakani, S. Yaldiz, L. Pileggi, R. Harjani, S. Reynolds, J. A. Tierno, and D. Friedman, "A 23.5 GHz PLL with an adaptively biased VCO in 32 nm SOI-CMOS," in *IEEE Custom Integrated Circuits Conference*, 2012.

Mark Ferriss received a first class honors B.E. degree in electrical engineering from University College Cork, Ireland, in 1998 and a Ph.D. from the University of Michigan, Ann Arbor, in 2008. From 1998 to 2002 Mark worked for Analog Devices in Limerick, Ireland, during which time he worked on Digital-to-Analog converters, switches, controllers and phase-locked loops for fiber optic communications. In 2009 Mark joined the IBM T. J. Watson Research Center in Yorktown Heights, NY. His present research interests include self healing RF communication circuits, analog to digital interface circuits, and digital phase-locked loops.

Jean-Olivier Plouchart (M'96–SM'06) received in 1994 the Ph.D. degree in electronics from Paris VI University, France. From 1989 to 1996 he worked with Alcatel, France Telecom and the University of Michigan on HBT and MESFET MMICs for communication applications. In 1996, he joined as a research staff member the IBM T. J. Watson Research Center where his work involved the design of SiGe BiCMOS and CMOS RFIC circuits for wireless LAN applications, as well as RF product designs for Motorola. In 2000, he lead a team working on low-power high-performance SOI SoC technology and enablement, leading to the first demonstration of 130 and 90 nm SOI ASIC, and RF SOI circuits on high-resistivity substrate, as well as the enablement of the first 3.5 W 1 GHz Pentium class microprocessor. He also pioneered the design of millimeter wave SOI CMOS from 30 to 94 GHz in standard microprocessor technology. His research interests include solid-state technologies, the integration of RF transceivers, VCO, PLL and Analog to Digital converter with microprocessors for SoC applications, the RF to mmWave measurement automation and the Design For Yield in nanometer technologies. Currently, he leads the development of nanometer high-speed circuit design and high-yield nanometer design at the IBM T. J. Watson Research Center, Yorktown Heights, NY.

Dr. Plouchart is a Senior IEEE member, and coauthor of the best student paper award at the 2002 IEEE Radio Frequency Integrated Circuit Conference. He holds sixteen US patents, has published over 84 publications and one book chapter. He served as the Chairman of the IEEE CSICS CMOS committee in 2009, and currently serves as a member of the AMS ITRS roadmap as well as the SRC AMS Technical Advisory Board.

Arun Natarajan received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, in 2001 and the M.S. and Ph.D. degrees in electrical engineering from the California Institute of Technology (Caltech), Pasadena, in 2003 and 2007, respectively. From 2007 to 2012, he was a Research Staff Member at IBM T. J. Watson Research Center, NY and worked on mm-wave phased arrays for multi-Gb/s data links and airborne radar and on self-healing circuits for increased yield in sub-micron process technologies. In 2012 he joined Oregon State University as an assistant professor in the School of Electrical Engineering and Computer Science. His current research is focused on RF, mm-wave and sub-mmwave integrated circuits and systems for high-speed wireless communication and imaging. Dr. Natarajan received the National Talent Search Scholarship from the Government of India [1995–2000], the Caltech Atwood Fellowship in 2001, the Analog Devices Outstanding Student IC Designer Award in 2004, and the IBM Research Fellowship in 2005, and serves on the Technical Program Committee of the IEEE Radio-Frequency Integrated Circuits (RFIC) Conference, IEEE Compound Semi-conductor IC Symposium (CSICS), and the 2013 IEEE International Microwave Symposium (IMS).

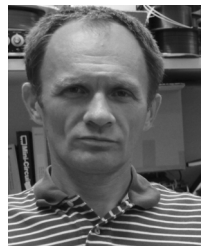

Alexander Rylyakov received the M.S. degree in physics from Moscow Institute of Physics and Technology in 1989 and the Ph.D. degree in physics from State University of New York at Stony Brook in 1997. From 1994 to 1999 he worked in the Department of Physics at SUNY Stony Brook on the design and testing of integrated circuits based on Josephson junctions. In 1999 he joined IBM T. J. Watson Research Center as a research staff member. Dr. Rylyakov's main current research interests are in the areas of digital phase-locked loops and integrated circuits for wireless and optical communication.

Ben Parker received a B.S. in physics from Bowdoin College, Brunswick, Maine, in 1979. He received a M.S. in physics from Brown University in Providence, Rhode Island in 1981. His graduate work dealt with the optical properties of adsorbed layers on metal surfaces. In 1986 he joined the GaAs group at IBM T. J. Watson Research Center, Yorktown Heights, New York, where he worked on characterization of III-V semiconductors. In 1991, he joined the Mixed-Signal Communications IC Design Group, working on design and verification of digital circuits in high speed serial communications.

José A. Tierno received the Engineering degree from the Universidad de la República, Montevideo, Uruguay, in 1988. He received the M.S. degree in electrical engineering in 1989 and the Ph.D. degree in computer science in 1995, both from the California Institute of Technology, Pasadena. From 1995 to 2012 he worked at the IBM T. J. Watson Research Center, Yorktown Heights, NY, on digital circuits for communications. Since 2012 he has worked at Apple Inc., Cupertino, California as a research scientist. His main areas of interest are self-timed digital circuits, and digital replacement of analog circuits.

A. Babakhani is an Assistant Professor of Electrical and Computer Engineering Department at Rice University and the Director of the Rice Integrated Systems and Circuits Laboratory. He won a prestigious DARPA Young Faculty Award (YFA) in 2012. He received the Caltech Electrical Engineering Department's Charles Wilts Best Ph.D. Thesis Prize for his work on Near-Field Direct Antenna Modulation (NFDAM). From 2006 to 2008 he was the Vice Chair of the IEEE Microwave Theory and Techniques Society Metro LA/SFV Joint Sections MTT-S Chapter 17.1. He was the recipient of the Microwave Graduate Fellowship in 2007, the Grand Prize in the Stanford-Berkeley-Caltech Innovators Challenge in 2006, Analog Devices Inc. Outstanding Student Designer Award in 2005, as well as Caltech Special Institute Fellowship and Atwood Fellowship in 2003. He was also the Gold Medal winner of the National Physics Competition in 1998, and the Gold Medal winner of the 30th International Physics Olympiad in 1999, in Padova, Italy.

Soner Yaldiz received the B.S. degree in microelectronics engineering from Sabanci University, Istanbul, Turkey in 2004, the M.S. and the Ph.D. degrees in electrical and computer engineering from Koc University, Istanbul, Turkey in 2006 and from Carnegie Mellon University, Pittsburgh, PA in 2012 respectively. His master's research was on energy optimization for multi-core systems using stochastic workload models. His doctoral research was focused on design methodologies for self-healing analog integrated circuits. He has interned at Nvidia, CA and IBM T. J. Watson Research Center, NY. He is working on analog circuit design methodologies at Intel Corporation, OR since January 2012.

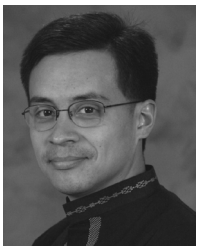

Alberto Valdes-Garcia received the Ph.D. degree in electrical engineering from Texas A&M University in 2006 and became a Research Staff Member at the IBM T. J. Watson Research Center. His present research work is on silicon-integrated millimeter-wave systems and carbon electronics.

From 2006 to 2009, Dr. Valdes-Garcia served in the IEEE 802.15.3c 60 GHz standardization committee and from 2010 to 2012 in the Technical Program Committee of the IEEE *Custom Integrated Circuits Conference* (CICC). Since 2009 he serves as Technical Advisory Board member with Semiconductor Research Corporation (SRC), where he served as Chair of the Integrated Circuits and Systems Sciences Coordinating Committee in 2011 and 2012. He holds 6 issued US patents with 20+ pending. As of 2012, his scholarly work (60+ authored or co-authored publications) has already received more than 900 independent citations in indexed journals. He is a co-editor of the book 60 GHz *technology for Gbps WLAN and WPAN: From Theory to Practice* (Wiley, 20011).

Dr. Valdes-Garcia is the winner of the 2005 Best Doctoral Thesis Award presented by the IEEE Test Technology Technical Council (TTTC), the recipient of the 2007 National Youth Award for Outstanding Academic Achievements, presented by the President of Mexico, and a co-recipient of the 2010 George Smith Award presented by the IEEE Electron Devices Society. Within IBM, he has been a co-recipient of two Research Division Awards (2012), an IBM Corporate Outstanding Innovation Award for the demonstration of wireless high definition video links with 60 GHz SiGe radios (2008), and the 2009 Pat Goldberg Memorial Award to the best paper in computer science, electrical engineering, and mathematics within IBM Research for the work *Operation of Graphene Transistors at GHz Frequencies* (Nano Letters, 2009).

Bodhisatwa Sadhu is currently a Post-doctoral Researcher at IBM T. J. Watson Research Center, NY. He received his B.E.(Hons.) degree in electrical and electronics engineering from Birla Institute of Technology and Science, Pilani, in 2007 and his Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis, in 2012. For his Ph.D., he worked on wideband circuits and architectures for software defined radio applications. In 2007, he was with Broadcom Corporation, Bangalore, where he worked on system integration and verification of ethernet switch SoCs. In Fall-2010 and Summer-2011, he was with the Mixed Signal Communications IC Design Group, IBM T.J. Watson Research Center where he worked on the analysis and design of low phase noise frequency synthesizers for 60 GHz and 94 GHz applications. Dr. Sadhu is the recipient of the University of Minnesota Graduate School Fellowship, 2007, 3M Science and Technology Fellowship, 2009 and the University of Minnesota Doctoral Dissertation Fellowship, 2011.

Daniel J. Friedman (S'91–M'92) received the Ph.D. degree in engineering science from Harvard University, Cambridge, MA, in 1992. After completing consulting work at MIT Lincoln Labs and postdoctoral work at Harvard in image sensor design, he joined the IBM Thomas J. Watson Research Center, Yorktown Heights, NY, in 1994. His initial work at IBM was the design of analog circuits and air interface protocols for field-powered RFID tags. In 1999, he joined the mixed-signal communications IC design group and turned his attention to analog circuit design for high-speed serializer/deserializer macros. He managed the mixed-signal team from 2000–2009, focusing efforts on serial data communication and clock synthesis applications. In 2009, he became manager of the communication circuits and systems group, adding responsibility for teams in millimeter-wave wireless and digital communications IC design. He has authored or co-authored more than 40 technical papers in circuit topics including serial links, PLLs, RFID, and imagers. He was a co-recipient of the Beatrice Winner Award for Editorial Excellence at the 2009 ISSCC and the 2009 JSSC Best Paper Award given in 2011; he holds more than 50 patents. He has been a member of the ISSCC international technical program committee since 2008; he has served as the Wireline sub-committee chair from ISSCC12 to the present. His current research interests include high-speed I/O design, PLL design, and circuit/system approaches to enabling new computing paradigms.