

A Sub-Picosecond Resolution 0.5–1.5 GHz Digital-to-Phase Converter

Pavan Kumar Hanumolu, *Member, IEEE*, Volodymyr Kratyuk, *Member, IEEE*, Gu-Yeon Wei, *Member, IEEE*, and Un-Ku Moon, *Senior Member, IEEE*

Abstract—A digital-to-phase converter (DPC) is an essential building block in applications such as source-synchronous interfaces and digital phase modulators. The resolution of DPCs using analog phase interpolators is severely affected by the operating frequency and rise times of the interpolator inputs. In this paper, we present a new DPC architecture that achieves high resolution independent of both the operating frequency and the rise time. The 8 phases generated by a phase-locked loop are dithered using a delta-sigma modulator to shape the truncation error to high frequency and is subsequently filtered using a delay-locked loop phase filter. The test chip, fabricated in a 0.13 μm CMOS process, operates from 0.5–1.5 GHz and achieves a differential nonlinearity of less than ± 0.1 ps and an integral nonlinearity of ± 12 ps. The total power consumption while operating at 1 GHz is 15 mW.

Index Terms—Digital-to-phase converter, phase interpolation, noise shaping, delta-sigma modulation, delay-locked loop (DLL), phase-locked loop (PLL), phase filter, glitch-free phase switching.

I. INTRODUCTION

SOURCE-SYNCHRONOUS interfaces are a class of point-to-point links that are widely used in microprocessor-memory interfaces and communication switches. A simplified block diagram of a typical source-synchronous interface is shown in Fig. 1. In this system, a clock is transmitted along with the data on a separate dedicated channel to the receiver. The clock channel is typically shared among multiple data channels and clock edges are synchronized with the data transitions at the transmitter. If the data and clock transmission lines are perfectly matched, the time of flight of the data and the clock are equal, and as a result, clock and data remain synchronized at the receiver as well. However, as data rates increase to the multi-gigabit-per-second range, it is uneconomical to match the time of flight of clock and data paths to picosecond accuracy. This mismatch results in a skew between the clock and data at the receiver causing sub-optimal sampling of the incoming data. One can improve the timing margin by reducing the skew between the received clock and data by using a method to introduce a controlled phase shift on the clock. The focus of this paper is the implementation of




Fig. 1. Typical source-synchronous interface.

circuits that provide a means to introduce such a programmable phase shift. A digital-to-phase converter (DPC) is a circuit block that is often used to introduce a phase shift whose amount is controlled by an input digital word D_{IN} . It is important to note that the resolution of the DPC is of paramount importance as this determines the residual skew between the clock and data, which in turn directly affects the bit error rate (BER) of the link. This paper presents the design and experimental results of a DPC that utilizes delta-sigma modulation and phase filtering to achieve resolution much higher than what is achievable with traditional digital phase interpolators. Even though the design of the DPC is presented in the context of source-synchronous interfaces, it is worth mentioning that there are several other applications for DPCs in measurement instrumentation and the techniques developed here can directly be used in those applications.

Before we present the proposed DPC architecture, it is instructive to review the disadvantages of existing architectures. One of the earliest implementations of the DPC is shown in Fig. 2 [1]. It consists of a multi-phase clock generator, which provides N clock phases separated by a delay of ΔT . These multiple phases (Φ_1 to Φ_N) are typically generated through a chain of inverters whose delay is precisely adjusted to ΔT by a feedback loop. An N -to-1 multiplexer (MUX) is used to select one of the N phases based on the input digital word D_{IN} , thereby introducing a phase shift in steps of ΔT on the output. With a 16-phase multi-phase generator, this architecture achieves a phase resolution of 22.5° . There are several drawbacks with this approach. First, the resolution, ΔT , is limited by the minimum delay of the inverter in a given process. Second, since ΔT is

Manuscript received January 2, 2007; revised October 15, 2007. This work was supported by Intel Corporation. The test chip was fabricated by Samsung Electronics Ltd.

P. K. Hanumolu and U. Moon are with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331 USA (e-mail: hanumolu@eecs.oregonstate.edu).

V. Kratyuk is with Silicon Laboratories, Inc., Beaverton, OR 97006 USA.

G.-Y. Wei is with SEAS, Harvard University, Cambridge, MA 02138 USA.

Digital Object Identifier 10.1109/JSSC.2007.914287




Fig. 2. DPC using phase selection.




Fig. 3. DPC using phase selection and interpolation.




Fig. 4. Phase interpolator. (a) Operation. (b) Model.

equal to a fraction of the clock period, T_{period}/N , the resolution scales directly with frequency, thereby degrading resolution at lower operating frequencies. Finally, the phase selection process introduces unwanted discrete phase jumps in the output phase. Despite its simplicity, these performance limiting factors hamper the use of this DPC in multi-gigabit-per-second interfaces.

A more commonly used DPC architecture that overcomes some of these drawbacks is depicted by Fig. 3 [2]–[4]. This architecture combines the phase selecting multiplexer with a phase interpolator. The most significant bits (MSBs) of the input digital word are used to select two adjacent clock phases, Φ_j and Φ_{j+1} , from the N phases using an $N:2$ MUX. The interpolator, controlled by the least significant bits (LSBs), mixes these two phases to generate an intermediate phase Φ_{OUT} . As a result of phase interpolation, the resolution of this DPC is not limited to a minimum inverter delay. However, the effectiveness of the interpolation depends largely on the input rise time, phase separation ΔT , and the interpolator output time constant. Consider the conceptual phase interpolator block diagram and its model shown in Fig. 4. Ideally, the interpolator delay should depend only on interpolation weight, α , but in practice, the output phase also depends on the interpolator output time constant (RC), rise

time of the inputs (τ_r), and the time difference between the inputs. This dependence is illustrated in Fig. 5, in which the interpolator transfer function (α -to-output phase) is plotted for multiple values of ΔT and τ_r . All the time parameters, ΔT , τ_r , and output phase, are normalized to the output RC time constant. The output phase is referenced to the delay when the interpolation weight is zero as expressed by

$$\text{Normalized output phase at } \alpha 1 = \frac{T_D|_{\alpha=\alpha 1} - T_D|_{\alpha=0}}{RC} \quad (1)$$

where $T_D|_{\alpha=\alpha 1}$ and $T_D|_{\alpha=0}$ are the interpolator delays when the interpolation weights are equal to $\alpha 1$ and 0, respectively. When the rise time is very small compared to the phase spacing (Fig. 5(a)), the transfer function becomes grossly nonlinear as ΔT becomes larger than the output RC time constant. This non-linearity can be reduced by increasing the input rise time to 3 times the phase separation, as shown by Fig. 5(b). However, the slow rise times needed to achieve good linearity degrade the jitter immunity of the output clock [5]. Moreover, the resolution of this architecture also depends on operating frequency. The nonlinearity of the interpolator increases with increasing phase separation ΔT , thereby degrading the output phase resolution at lower operating frequencies. Finally, the output jitter of this architecture is severely affected by the discrete phase jumps introduced during the input phase switching of the interpolator. These drawbacks limit the phase resolution of this architecture to about 4° . A new DPC architecture is proposed to overcome these drawbacks and achieve better than 0.1° phase resolution.

II. PROPOSED ARCHITECTURE

The block diagram of the proposed DPC is shown in Fig. 6. Similar to the earlier implementations, the most significant bits (MSBs) of the input digital word D_{IN} are used to select three adjacent phases, Φ_{j-1} , Φ_j , and Φ_{j+1} out of the N phases generated by a multi-phase clock generator. However, as opposed to the previous implementations, the remaining least significant bits (LSBs) are quantized to 3 levels, -1 , 0 , and $+1$, by a second-order delta-sigma modulator (DSM). This 3-level DSM output is then used to select one of the three phases out of the N -to-3 MUX. As a result of the delta-sigma truncation of the LSBs, the resulting quantization error is shaped to high frequencies. By virtue of phase selection using the DSM output, this quantization error appears as shaped phase noise at the output of the 3-to-1 MUX, and filtering this high-frequency phase noise enables precise phase adjustment. Ideally, the phase resolution achieved by this architecture is equal to $1\text{UI}/2^B$, where B is the number of bits in the input digital word. At 1 GHz operation with a 14-bit input, this architecture ideally achieves about 60 fs of phase resolution.

Operational details of the proposed DPC are presented with respect to design parameters used in the prototype chip. In this implementation, the multi-phase generator provides 8 coarse phases Φ_1 to Φ_8 . The 3 MSBs of a 14-bit input digital word (D_{IN}) are used to select 3 out of 8 phases according to the mapping shown in Table I. For example, to generate an output phase between 67.5° and 112.5° , indicated by the shaded region in the phasor diagram of Fig. 7, phases Φ_2 , Φ_3 , and Φ_4 are selected. It is important to note that this mapping prevents overloading in




Fig. 5. Analysis of phase interpolator linearity. (a) Rise time τ_r much smaller than the phase spacing ΔT . (b) Rise time equal to three times the phase spacing. The solid line represents the transfer function with $\Delta T/RC = 0.5$ and dashed lines with $\Delta T/RC = 1, 1.5, 2$.

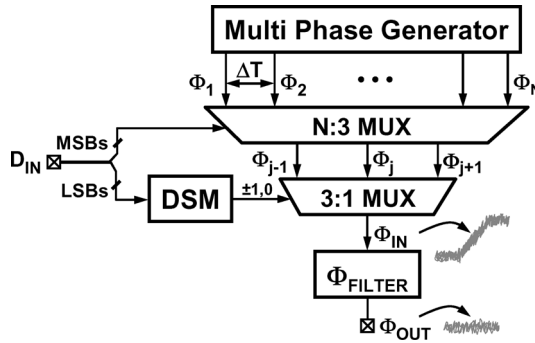


Fig. 6. Proposed DPC architecture.

the DSM because it guarantees that the input is only half of the full-scale of the DSM. In the prototype chip, the 3 levels of the DSM output ± 1 and 0 correspond to $\pm 45^\circ$ and 0° , respectively. The input to the DSM is then limited to an output phase corresponding to $\pm 22.5^\circ$. The selected coarse phases are dithered by the DSM according to the 11 LSBs of the input digital word. The power spectral density of the phase noise at the output of the 3:1 MUX, $S_{\Phi_q}(f)$ when a second-order DSM is used is given by [6]

$$S_{\Phi_q}(f) = \frac{1}{12F_s} \cdot \left(\frac{2\pi}{8}\right)^2 \cdot \left[2\sin\left(\frac{\pi f}{F_s}\right)\right]^4 \quad (2)$$

where F_s is the sampling frequency of the DSM. The low-pass response of a subsequent phase filter suppresses the shaped high-frequency noise. However, due to incomplete filtering, the shaped noise can leak to the output, resulting in residual phase noise at the output of the phase filter given by

$$S_{\Phi_{OUT}}(f) = S_{\Phi_q}(f) \cdot |\Phi_{FILTER}(f)|^2 \quad (3)$$

where $\Phi_{FILTER}(f)$ is the transfer function of the phase filter. Fig. 8 depicts the shaped phase noise at the output of the 3:1 MUX along with the residual noise denoted by the shaded region. For illustration purposes, a brick wall response is assumed for the phase filter. As expected, the figure shows the bandwidth of the phase filter should be low enough to not degrade

TABLE I
MAPPING BETWEEN OUTPUT PHASE Φ_{OUT}
AND COARSE PHASES $\Phi_{j-1}, \Phi_j, \Phi_{j+1}$

Φ_{OUT} range	Φ_{j-1}	Φ_j	Φ_{j+1}
$22.5^\circ \leq \Phi_{OUT} < 67.5^\circ$	1	2	3
$67.5^\circ \leq \Phi_{OUT} < 112.5^\circ$	2	3	4
$112.5^\circ \leq \Phi_{OUT} < 157.5^\circ$	3	4	5
$157.5^\circ \leq \Phi_{OUT} < 202.5^\circ$	4	5	6
$202.5^\circ \leq \Phi_{OUT} < 247.5^\circ$	5	6	7
$247.5^\circ \leq \Phi_{OUT} < 292.5^\circ$	6	7	8
$292.5^\circ \leq \Phi_{OUT} < 337.5^\circ$	7	8	1
$337.5^\circ \leq \Phi_{OUT} < 22.5^\circ$	8	1	2




Fig. 7. Phasor diagram to illustrate DPC operation.

the output phase noise. A more practical phase-filter response will be used in the next section to demonstrate the design considerations quantitatively.

This architecture offers several advantages. First, by virtue of noise shaping and phase filtering, this architecture is capable of achieving sub-picosecond phase resolution [7]. Second, since the digital-to-phase conversion is based on phase selection and filtering, as opposed to interpolation, this technique does not depend on the rise time of the clock phases. Consequently, the output clock is less sensitive to noise that causes jitter. Third, the




Fig. 8. Frequency domain view of the phase noise due to DSM noise shaping.

smoothing nature of the phase filter eliminates discrete phase jumps often present in conventional implementations. Finally, this technique relies mostly on digital circuitry and is, therefore, easily portable to different processes compared to more analog-centric implementations.

While the resolution of this architecture does depend on operating frequency, because of increased phase spacing ΔT at lower operating frequencies, this resolution dependence on operating frequency can be suppressed by designing the DPC to be limited by clock jitter. In other words, if the resolution of the DPC is much higher than the inherent jitter of the dithered phases, then the reduced resolution will be masked by the clock jitter. The phase quantization error of the DPC can be designed to be lower than the phase noise floor determined by intrinsic noise sources such as thermal and flicker noise.

III. PHASE FILTER IMPLEMENTATION

One of the most important building blocks of the DPC is the phase filter. A common implementation of a low-pass phase filter is a phase-locked loop (PLL). However, as is well known, the design of a high-performance PLL poses several challenges. Notably, jitter accumulation of the voltage-controlled oscillator (VCO) results in excessive output jitter and the suppression of this jitter requires large power dissipation. Also, the large gain of the VCO in deep submicron processes mandates a large loop filter capacitor that occupies considerable area to stabilize the loop. In addition to these drawbacks, PLLs suffer from an inherent noise–bandwidth tradeoff. The input phase noise is suppressed by a low-pass transfer function, while the VCO noise is shaped by a high-pass transfer function. In the context of using a PLL as a phase filter in the DPC, the low bandwidth required to suppress the delta-sigma noise exacerbates VCO noise. Because of these disadvantages, a PLL phase filter is not used in the prototype. It is worth mentioning that injection-locked oscillators can also serve as phase filters, but they require additional control to bring the oscillating frequency to within their pull-in range and also suffer from degraded noise performance due to jitter accumulation.

As opposed to a PLL, a delay-locked loop (DLL) offers superior jitter performance, is less sensitive to supply noise, occupies smaller area, and typically consumes lower power. Therefore, it is beneficial to consider using a DLL as the phase filter in the DPC. Since the noise from the voltage-controlled delay




Fig. 9. Modified DLL with low-pass transfer function.




Fig. 10. Small-signal model of the modified DLL.

line (VCDL) is not much of a concern, there is no noise–bandwidth tradeoff. However, the DLL suffers from a major disadvantage when considered for use in the DPC. The input-to-output transfer function $\Phi_{OUT}(s)/\Phi_{IN}(s)$ of the DLL is an all-pass, making it incapable of suppressing the shaped input noise. So instead, a modified DLL is used in the prototype and shown in Fig. 9. It achieves the needed low-pass transfer function while preserving all of the other advantages of a conventional DLL. In this architecture, the input phase Φ_{IN} is fed only to the phase detector and a separate reference phase Φ_{REF} is used as the input to the delay line. Consequently, the transfer function from the input Φ_{IN} is a low-pass while the transfer function from the reference Φ_{REF} is an all-pass. Using the small-signal model of the DLL shown in Fig. 10, the input transfer function can be derived as

$$LG(s) = \frac{I_{CP} \cdot K_{VCDL} \cdot F_{IN}}{Cs} \quad (4)$$

$$\frac{\Phi_{OUT}(s)}{\Phi_{IN}(s)} = \frac{LG(s)}{LG(s) + 1} \quad (5)$$

$$= \frac{I_{CP} \cdot K_{VCDL} \cdot F_{IN}}{s + I_{CP} \cdot K_{VCDL} \cdot F_{IN}} \quad (6)$$

where $LG(s)$ is the loop gain, I_{CP} is the charge pump current, K_{VCDL} is the gain of the VCDL, C is the loop filter capacitance, and F_{IN} is the input frequency.

When using this modified DLL as the phase filter, there are two important design parameters that determine the achievable resolution in the proposed architecture. First, the sampling rate of the DSM determines the effectiveness of noise shaping. For example, in a second-order DSM with a 3-level internal quantizer, the signal-to-quantization ratio improves by 15 dB with a doubling of the sampling frequency [8]. Second, as mentioned earlier, the bandwidth and the order of the phase filter determine the residual quantization error. These two parameters, the sampling frequency F_s and the filter bandwidth BW , are combined to define the effective over-sampling rate (OSR) as

$$OSR = \frac{F_s}{2BW}. \quad (7)$$

The effectiveness of a first-order DLL phase filter is illustrated by plotting the residual jitter (quantization error leakage)




Fig. 11. Residual jitter versus over-sampling ratio for a first-order DLL.




Fig. 12. Low-pass DLL with an active loop filter.

resulting from limited filtering versus OSR, shown in Fig. 11. This plot is obtained from behavioral simulations of the DPC using a DLL phase filter whose transfer function is given by (6). It shows that there is a considerable amount of residual jitter even at an OSR of 150. A high OSR translates to a larger sampling frequency, resulting in larger power dissipation in the DSM. The excessive residual jitter at lower OSR is mainly due to the fact that the delta-sigma modulator is second order while the DLL is only first order.

In addition to the ineffective filtering, this DLL also suffers from noise folding due to a nonlinearity in the charge pump that can result from current mismatch [9]. This mismatch is further exacerbated by a varying control voltage, V_C , needed to achieve the required large output phase range. To overcome both the charge pump nonlinearity and the incomplete filtering of the first-order DLL, an improved DLL that employs an active loop filter is used in the prototype. The block diagram of the modified DLL is shown in Fig. 12. The use of an active loop filter offers two main advantages. First, the feedback amplifier biases the output of the charge pump at a fixed reference voltage, V_{REF} , irrespective of the delay setting of the VCDL. As a result, current mismatch in the charge pump due to a varying control voltage is suppressed. Second, the higher order poles of the amplifier further suppress the shaped high-frequency noise, thus reducing the jitter due to quantization error leakage. The bandwidth of the amplifier is carefully optimized to achieve a second-order DLL transfer function without compromising the stability of the




Fig. 13. Residual jitter versus over-sampling ratio for a second-order DLL.




Fig. 14. Complete DPC architecture.

overall DLL feedback loop. Additionally, the amplifier bandwidth was chosen small enough so as to suppress the quantization error adequately even in the presence of PVT variations. A transfer function of the DLL that accounts for the limited amplifier bandwidth ω_{opamp} is given by

$$\frac{\Phi_{OUT}(s)}{\Phi_{IN}(s)} = \frac{K}{s^2 + s\omega_{opamp} + K\omega_{opamp}}$$

$$\text{where } K = \frac{I_{CP} \cdot K_{VCDL} \cdot F_{IN}}{C}. \quad (8)$$

The improvement in resolution of the DPC, due to the extra filtering offered by the finite amplifier bandwidth, is illustrated in Fig. 13. The resolution of the DPC is improved by more than 8X when compared to a first-order DLL at an OSR of 100 (see Fig. 11). This improved filtering allows for a lower sampling frequency in the DSM, which results in lower power. Incidentally, the bandwidth of the DLL is input frequency dependent, which makes the effective OSR of the DSM independent of the input frequency.

We now provide details of generating the reference phase Φ_{REF} used as the input to the VCDL in Fig. 12. In the DPC test chip, the reference input to the DLL is tapped off from one of the 8 phases out of a multi-phase clock generating PLL, as shown by the complete architecture illustrated in Fig. 14. As discussed in a later section, false locking in the DLL is avoided by maintaining an appropriate phase relation between Φ_{IN} and Φ_{REF} at start-up.




Fig. 15. Four-stage ring oscillator and the delay cell.

IV. CIRCUIT DESIGN

A. Phase-Locked Loop Design

The multi-phase clock generator is implemented by using a PLL. The PLL consists of a phase frequency detector (PFD), a charge pump (CP), a loop filter consisting of a series RC network, a 4-stage voltage-controlled ring oscillator (VCO), and a divider in the feedback. The PFD is implemented as a 3-state machine and generates a pair of digital pulses whose widths correspond to the frequency and phase error between the reference clock (REF) and the divided clock that is fed back [10]. The CP then converts the digital pulses into an analog current that is converted into a voltage via the passive loop filter. The resulting control voltage, V_C , drives the VCO toward phase lock. The VCO generates eight equally spaced phases Φ_1 to Φ_8 , one of which is buffered and fed back to the divider. Dummy inverters are used on the other unused phases to preserve equal loading and delay spacing between the adjacent phases.

The schematic of the VCO along with the delay cell is shown in Fig. 15. The delay cell is a simple pseudo-differential inverter in which a pMOS latch is used to couple the two single-ended current starved inverters to generate a differential output [11]. The output of the delay cell is buffered to nominally maintain a 50% duty cycle under process, voltage, and temperature variations. Transistor level simulations indicate that the operating range of the VCO is 0.3 GHz–2 GHz and the gain is 2 GHz/V. The simulated VCO phase noise is approximately -110 dBc/Hz at 3 MHz offset from the carrier frequency over the whole operating range. Using the design equations in [12], the charge pump current, loop filter resistor and capacitor values are determined to be 15 μ A, 8 k Ω , and 28 pF, respectively. These parameters result in a PLL bandwidth of about 5 MHz with a phase margin of 65°. The divider is implemented by a cascade of three TSPC divide-by-2 stages [13].

B. Delay-Locked Loop Design

A brief overview of the DLL with a low-pass transfer function was presented in Section III. The implementation details of the DLL are presented in this section. The schematic of the DLL used in the prototype is shown in Fig. 16. It consists of a phase-only detector (PD), a differential charge pump (CP), an active loop filter, and a VCDL. The phase-only detector generates digital pulses corresponding to the phase difference between the DLL input (IN) and delayed clock (DCK). The charge

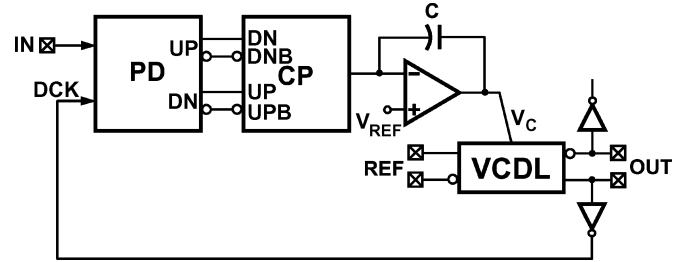


Fig. 16. Implemented delay-locked loop with active loop filter.




Fig. 17. Timing diagram illustrating a stuck at minimum delay fault.

pump converts these digital pulses into an output current, which is filtered by an active integrator. The integrator output drives the VCDL in a way that forces the phase error to zero. In the locked state, the delay of the VCDL is typically equal to the period of the input.

Despite the use of a phase-only detector, the DLL, if not properly designed, suffers from start-up problems that can result in a stuck at minimum delay fault or harmonic locking. Harmonic locking is avoided by resetting the VCDL to its minimum delay point on start-up [14]. This resetting of the VCDL does not, however, avoid the DLL from trying to acquire lock to a delay point that is below the minimum delay offered by the VCDL, resulting in a stuck at minimum delay fault. This problem, arising from two different start-up conditions, is illustrated in Fig. 17. In the first case, the minimum delay of the VCDL ($TD_{\min 1}$) is less than half of the clock period (TP_{IN}). The PD generates a down pulse (DN) indicating that the delay of the VCDL be further reduced, which results in the DLL getting stuck to this minimum delay point. Similarly, in the second case, if $TD_{\min 2}$ is greater than half of the clock period, the DLL also gets stuck

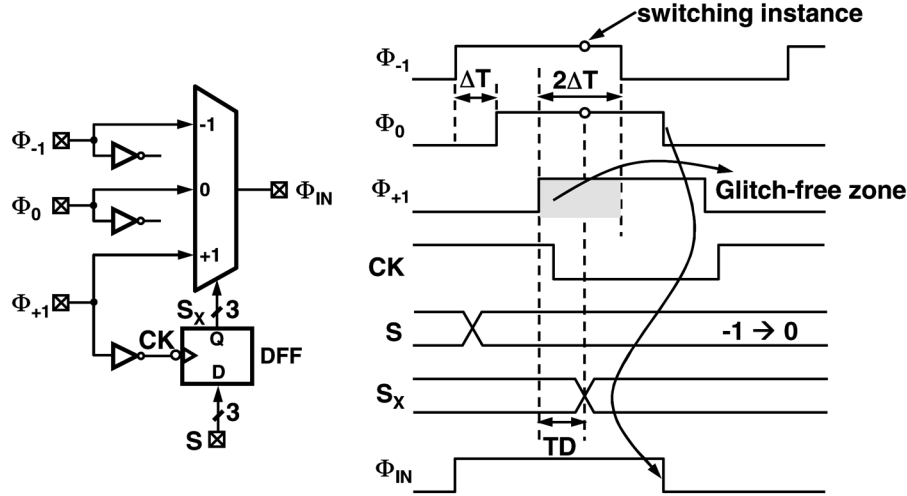


Fig. 22. Glitch-free switching scheme and the associated timing diagram.

diagram shown in Fig. 21. Consider the phase delay case where phase Φ_{-1} is switched to phase Φ_0 . If this phase switching occurs in the non-overlapping region indicated by the shaded region, a glitch occurs on the output phase Φ_{IN} as shown at the bottom of the figure. These glitches on the output phase can drive the DLL out of lock, resulting in complete operation failure of the overall DPC. Therefore, a method to prevent the glitches is needed.

It is useful to note that no glitches occur if the switching takes place during the overlap period in which both phases, Φ_{-1} and Φ_0 , are high (or low). The glitch-free switching scheme employed in the test chip is based on this observation and presented in Fig. 22. The control signal S_x is synchronized to the latest phase Φ_{+1} , so that phase switching occurs only in the glitch-free zone indicated by the shaded region in Fig. 22. This is achieved by synchronizing the MUX input control signal S to the phase Φ_{+1} . Dummy inverters are added on the other two phases, Φ_{-1} and Φ_0 , to preserve equal spacing. It is important that the sum of the delays of the inverter and clock-to-Q delay of the D-flip-flop (DFF) be less than $2\Delta T$ to ensure that a glitch-free zone exists. Mathematically, the following inequality should be satisfied for glitch-free switching:

$$TD = T_{INV} + T_{CK-Q} < 2\Delta T = \frac{2TP_{IN}}{8}. \quad (10)$$

V. EXPERIMENTAL RESULTS

A block diagram of the implemented prototype is shown in Fig. 23. In order to obviate complex circuitry for measuring sub-picosecond time differences, an exclusive OR (XOR) gate is used to convert the phase difference into a voltage. The filtered XOR output voltage is easily measured using a high-resolution sampling oscilloscope. A fully differential XOR gate is implemented with the symmetric architecture presented in [17] and its simulated transfer function is shown in Fig. 24. The simulated gain of this XOR gate is 2 mV/ps. In order to further simplify testing, an accumulator is used to generate the 14-bit input digital word with a serial input SD_{IN} . The complete DPC including the test blocks was fabricated in a 0.13 μm CMOS process and

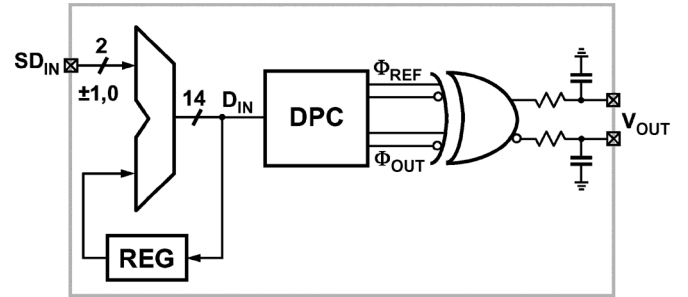


Fig. 23. Block diagram of the DPC prototype test chip.




Fig. 24. Simulated XOR transfer function.

the chip micrograph is shown in Fig. 25. The DPC occupies about 0.48 mm^2 of active die area. The die was packaged in a standard 48-pin LQFP plastic package. The packaged chip is attached to the four-layer test board through a clamp screw that is used to mechanically press down on the package and force its leads to contact solder pads on a printed-circuit board (PCB).

The measured transfer function of the DPC operating at 1 GHz is presented in Fig. 26. About 6% of the input codes on either end of the transfer curve are severely affected by the non-linearity of the XOR phase detector and are hence discarded.




Fig. 25. DPC chip micrograph.




Fig. 26. Measured transfer function of the DPC operating at 1 GHz.




Fig. 27. Measured DNL/INL of the DPC at 1 GHz operating frequency.

The linearity of the DPC is evaluated by plotting the differential and integral nonlinearities shown in Fig. 27. The maximum differential nonlinearity (DNL) is less than 0.1 ps while the maximum integral nonlinearity (INL) is about 12 ps. The low DNL validates the effectiveness of the proposed architecture, which relies on noise shaping via the delta-sigma modulator




Fig. 28. Effect of the multi-phase clock generator INL on DPC linearity.

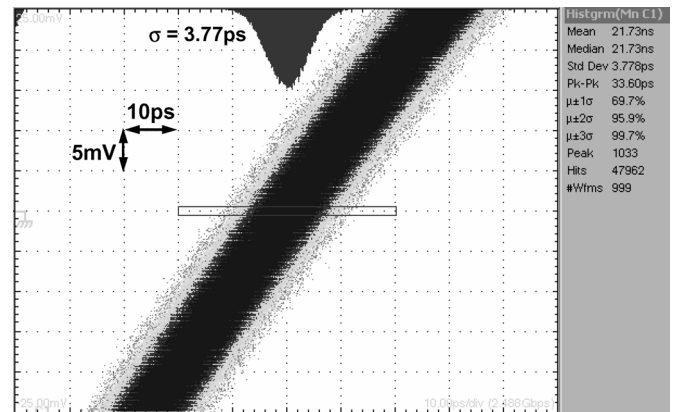


Fig. 29. PLL clock jitter at 1 GHz.

and subsequent filtering with a modified second-order DLL phase filter. Measured results also verify that the DNL and INL are less than ± 0.2 ps and ± 12 ps, respectively over the entire operating range of 0.5–1.5 GHz. This reinforces the earlier assertion that the resolution of the DPC is nearly independent of operating frequency. The measured output phase range of the DPC is greater than π radians over the entire operating range that was tested.

The symmetric nature of the INL reveals the cumulative effect of random phase mismatch and deterministic layout asymmetries in the multi-phase clock generator. In other words, the INL of the DPC is limited by the INL of the multi-phase clock generator. This is confirmed through behavioral simulations and the results are presented in Fig. 28. The output INL increases almost linearly with the input INL while the DNL is much less affected and remains nearly constant even for a large input INL.

The measured PLL clock jitter at 1 GHz [18] when the delta-sigma modulator and the DLL are reset is shown in Fig. 29. An rms jitter of 3.8 ps of the PLL sets the lower bound on the noise floor of the overall DPC. Fig. 30 shows the DPC output clock jitter when the input digital word is set to 100. The rms jitter of the phase-shifted output is 4.1 ps and this jitter increase translates to about 1.5 ps of jitter contribution from the delta-




Fig. 30. DPC clock jitter at 1 GHz.

TABLE II
DPC PERFORMANCE SUMMARY

Reference	This work	[4]
Technology	0.13μm CMOS	0.35μm CMOS
Supply voltage	1.2V	3.3V
Operating frequency	0.5GHz – 1.5GHz	50MHz – 250MHz
DNL/INL	±100fs/±12ps @ 1GHz	±31.25ps/±62.5ps @ 125MHz
Jitter	4.1ps rms	5.1ps rms
Phase span	> π radians	2π
Power consumption	15mW @ 1GHz	110mW @ 125MHz
Active die area	0.48mm ²	1.156mm ²

sigma modulator and the DLL, since it is uncorrelated to the noise floor.

The total power consumption of the DPC operating at 1 GHz with a supply voltage of 1.2 V is 15 mW of which 10 mW is consumed by the PLL. The DLL and all of the digital circuitry including the DSM and other test structures consume 3.5 mW and 1.5 mW, respectively. The performance of the DPC test chip is summarized in Table II.

VI. SUMMARY

A digital-to-phase converter architecture capable of achieving sub-picosecond resolution is presented in this paper. The use of a delta-sigma modulator to shape the phase noise to high frequencies and then filter it out with a low-pass filter presents an attractive alternative to the design of high-resolution digital-to-phase converters. The use of a DLL as a phase filter avoids the noise–bandwidth tradeoff of PLLs and facilitates the design of an area- and power-efficient low-pass filter. By relying on noise shaping and phase filtering, this architecture achieves high resolution that is independent of the operating frequency, rise time, and phase spacing of the input clock phases.

ACKNOWLEDGMENT

The authors thank Prof. Gabor Temes, Dr. José Ceballos, Dr. Younjae Kook, Dr. Gil-Cho Ahn, and Dr. Min Gyu Kim for useful discussion and critical feedback.

REFERENCES

- [1] J. Sonntag and R. Leonowich, “A monolithic CMOS 10 MHz DPLL for burst-mode data retiming,” in *IEEE ISSCC Dig. Tech. Papers*, 1990, pp. 194–195.
- [2] T. Lee, K. Donnelly, J. Ho, J. Zerbe, M. Johnson, and T. Ishikawa, “A 2.5 V CMOS delay-locked loop for 18 Mbit, 500 megabyte/s DRAM,” *IEEE J. Solid-State Circuits*, vol. 29, no. 12, pp. 1491–1496, Dec. 1994.
- [3] S. Sidiropoulos and M. Horowitz, “A semidigital dual delay-locked loop,” *IEEE J. Solid-State Circuits*, vol. 32, no. 11, pp. 1683–1692, Nov. 1997.
- [4] J. Chou, Y. Hsieh, and J. Wu, “A 125 MHz 8b digital-to-phase converter,” in *IEEE ISSCC Dig. Tech. Papers*, 2003, pp. 436–505.
- [5] A. Hajimiri, S. Limotyrakis, and T. Lee, “Jitter and phase noise in ring oscillators,” *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 790–804, Jun. 1999.
- [6] B. Miller and B. Conley, “A multiple modulator fractional divider,” in *Proc. Symp. Frequency Control*, May 1990, pp. 23–25.
- [7] P. Hanumolu, V. Kratyuk, G. Wei, and U. Moon, “A sub-picosecond resolution 0.5–1.5 GHz digital-to-phase converter,” in *Symp. VLSI Circuits Dig. Tech. Papers*, 2006, pp. 92–93.
- [8] R. Schreier and G. Temes, *Understanding Delta-Sigma Data Converters*. New York: Wiley-IEEE Press, 2005.
- [9] A. Ravi, R. Bishop, L. Carley, and K. Soumyanath, “8 GHz, 20 mW, fast locking, fractional-N frequency synthesizer with optimized 3rd order, 3/5-bit IIR and 3rd order 3-bit-FIR noise shapers in 90 nm CMOS,” in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, 2004, pp. 625–628.
- [10] M. Mansuri, D. Liu, and C. Yang, “Fast frequency acquisition phase-frequency detectors for Gsamples/s phase-locked loops,” *IEEE J. Solid-State Circuits*, vol. 37, no. 10, pp. 138–452, Oct. 2002.
- [11] J. Lee and B. Kim, “A low-noise fast-lock phase-locked loop with adaptive bandwidth control,” *IEEE J. Solid-State Circuits*, vol. 35, no. 8, pp. 1137–1145, Aug. 2000.
- [12] P. Hanumolu, M. Brownlee, K. Mayaram, and U. Moon, “Analysis of charge-pump phase-locked loops,” *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 51, no. 9, pp. 1665–1674, Sep. 2004.
- [13] J. Yuan and C. Svensson, “High speed CMOS circuit technique,” *IEEE J. Solid-State Circuits*, vol. 24, no. 1, pp. 62–70, Feb. 1989.
- [14] S. Sidiropoulos, C. Yang, and M. Horowitz, “A CMOS 500 Mbps synchronous point to point link,” in *Symp. VLSI Circuits Dig. Tech. Papers*, 1994, pp. 43–44.
- [15] S. Sidiropoulos, D. Liu, J. Kim, G. Wei, and M. Horowitz, “Adaptive bandwidth DLLs and PLLs using regulated supply CMOS buffers,” in *Symp. VLSI Circuits Dig. Tech. Papers*, 2000, pp. 124–127.
- [16] P. Raha, “A 0.6–4.2 V low-power configurable PLL architecture for 6 GHz–300 MHz applications in a 90 nm CMOS process,” in *Symp. VLSI Circuits Dig. Tech. Papers*, 2004, pp. 232–235.
- [17] B. Razavi, Y. Ota, and R. Swartz, “Design techniques for low-voltage high-speed digital bipolar circuits,” *IEEE J. Solid-State Circuits*, vol. 29, no. 3, pp. 332–339, Mar. 1994.
- [18] J. McNeill, “Jitter in ring oscillators,” *IEEE J. Solid-State Circuits*, vol. 32, no. 6, pp. 870–879, Jun. 1997.



Pavan Kumar Hanumolu (S’99–M’07) received the B.E. (Hons.) degree in electrical and electronics engineering and the M.Sc. (Hons.) degree in mathematics from the Birla Institute of Technology and Science, Pilani, India, in 1998, the M.S. degree in electrical and computer engineering from the Worcester Polytechnic Institute, Worcester, MA, in 2001, and the Ph.D. degree in electrical engineering from the Oregon State University, Corvallis, in 2006.

From 1998 to 1999, he was a Design Engineer at Cypress Semiconductors, Bangalore, India, working on phase-locked loops for low-voltage differential signaling (LVDS) interfaces. During the summers of 2002 and 2003, he was with Intel Circuits Research Labs, Hillsboro, OR, where he investigated clocking and equalization schemes for input/output (I/O) interfaces. Currently, he is an Assistant Professor in the school of Electrical Engineering and Computer Science, Oregon State University. His research interests include equalization, clock and data recovery for high-speed I/O interfaces, digital techniques to compensate for analog circuit imperfections, data converters, power-management circuits, and low-voltage mixed-signal circuit design.

Dr. Hanumolu received the Analog Devices Outstanding Student Designer Award in 2002, the Intel Ph.D. Fellowship in 2004, and was a co-recipient of

the Custom Integrated Circuits Conference (CICC) 2006 Best Student Paper Award. He currently serves as an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II, EXPRESS BRIEFS.



Volodymyr Kratyuk received the B.S. and M.S. degrees (Hons.) in physical electronics from the National Technical University of Ukraine, Kyiv, Ukraine, in 1997 and 1999, respectively, and the M.S. and Ph.D. degrees in electrical engineering from Oregon State University, Corvallis, in 2003 and 2006, respectively.

In 2004, he had a co-op position with Texas Instruments, Dallas, TX, working on pre-amplifiers for hard disk drives. Currently, he is with Silicon Laboratories, Inc., Beaverton, OR. His research interests involve frequency synthesizers, digital equivalent implementation of analog circuits, high-speed clocking and low-noise analog circuits.



Gu-Yeon Wei received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1994, 1997, and 2001, respectively.

He is currently an Associate Professor of electrical engineering in the School of Engineering and Applied Sciences at Harvard University, Cambridge, MA. After a brief stint as a Senior Design Engineer at Accelerant Networks, Inc. in Beaverton, OR, he joined the faculty at Harvard as an Assistant Professor in January 2002. His research interests span several areas: high-speed, low-power link design; mixed-signal circuits for communications; ultra-low-power hardware for wireless sensor networks;

and co-design of circuits and computer architecture for high-performance and embedded processors to address PVT variability and power consumption that plague nanoscale CMOS technologies.



Un-Ku Moon (S'92–M'94–SM'99) received the B.S. degree from the University of Washington, Seattle, in 1987, the M.Eng. degree from Cornell University, Ithaca, NY, in 1989, and the Ph.D. degree from the University of Illinois at Urbana-Champaign in 1994, all in electrical engineering.

He has been with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, since 1998, where he is currently a Professor. Before joining Oregon State University, he was with Bell Laboratories from 1988 to 1989, and from 1994 to 1998. His technical contributions have been in the area of analog and mixed-signal circuits including highly linear and tunable continuous-time filters, telecommunication circuits including timing recovery and data converters, and ultra-low-voltage analog circuits for CMOS.

Prof. Moon is a recipient of the National Science Foundation CAREER Award and the Oregon State University's Excellence in Graduate Mentoring Award. He has served as an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: ANALOG AND DIGITAL SIGNAL PROCESSING and on the Technical Program Committee of the IEEE Custom Integrated Circuits Conference. He currently serves as an Associate Editor of the IEEE JOURNAL OF SOLID-STATE CIRCUITS and the Editor-in-Chief of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II: EXPRESS BRIEFS, and on the Technical Program Committee of the IEEE VLSI Circuits Symposium and the Analog Signal Processing Technical Committee of the IEEE Circuits and Systems Society. He also serves on the IEEE Solid-State Circuits Society (SSCS) Administrative Committee (AdCom) and the IEEE Circuits and Systems Society (CASS) Board of Governors (BoG) as the SSCS representative to CASS.