An On-Chip All-Digital Measurement Circuit to Characterize Phase-Locked Loop Response in 45-nm SOI

Dennis Fischette, Richard DeSantis, and John Haeseon Lee

Advanced Micro Devices, Inc., Sunnyvale, CA 94085-3905 USA

Abstract—**An all-digital measurement circuit, built in 45-nm SOI-CMOS, enables on-chip characterization of phase-locked loop (PLL) response to a self-induced phase step. This technique allows estimation of PLL closed-loop bandwidth and jitter peaking. The circuit can be used to plot step-response vs. time, measure static phase error, and observe phase-lock status.**

I. INTRODUCTION

Many applications such as PCI Express™ require a PLL to produce a low-jitter clock at a given frequency while meeting stringent bandwidth and jitter peaking requirements. Process, voltage, and temperature (PVT) variations as well as random device mismatch make it difficult to guarantee a narrow range for PLL response. For example, loop parameters such as VCO gain could vary by more than 2X over PVT corners. In Fig. 1, we see the closed-loop jitter transfer functions of two PLLs with identical reference clock and output frequencies. One PLL exhibits large peaking and low bandwidth while the other shows little peaking but high bandwidth. Although differences in this example are more extreme than usual, similar but smaller differences often result from PVT variations.
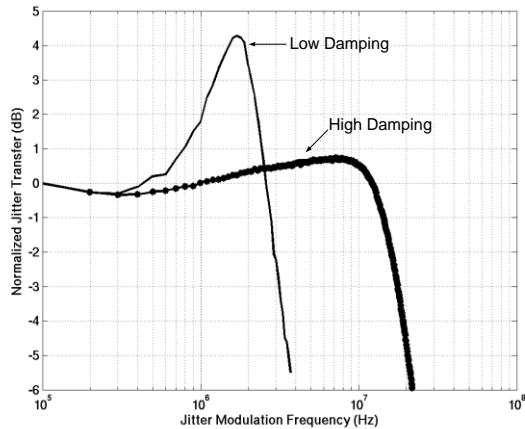

Fig. 1. PLL jitter transfer functions with different bandwidths and peaking.

PLL response is often measured on a test bench using signal generators, oscilloscopes, and/or spectrum analyzers. For example, the transfer functions in Fig. 1 were automatically generated by modulating the 100-MHz reference clock with various frequencies while observing the amplitudes of the resulting output spurs. Such methods, which may require many seconds to complete, motivate the need for faster, less expensive, and preferably on-chip techniques to characterize PLL response [1]–[3].

Fig. 2 shows the PLL output phase transient response to an induced phase step. Similar to other second-order feedback systems, the PLL tends to overcorrect (or overshoot) as it works to eliminate the induced phase error. If the PLL is underdamped, as in this example, the PLL may ring several times before settling to its final lock state. A key metric in the PLL step-response is $\tau_{crossover}$, defined here as the elapsed time from input step to onset of phase overshoot. Another key metric is *MaxOvershoot*. It measures the maximum overcorrection in the step response.
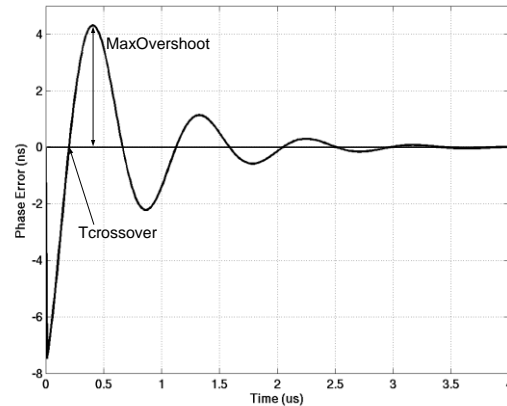

Fig. 2. Phase response of Type II second-order PLL to induced input phase step at time=0.

Transient simulations and closed-form loop equations [4] show that $\tau_{crossover}$ is inversely proportional to the PLL's −3dB closed-loop bandwidth; the smaller $\tau_{crossover}$ is, the higher the bandwidth (Fig. 3). Notice that $\tau_{crossover}$ is largely independent of the size of the phase step. Both simulations and loop equations also predict that *MaxOvershoot* is proportional to the maximum peaking in the closed-loop transfer function; the larger *MaxOvershoot* is, the greater the peaking (Fig. 4). Notice that the magnitude of the overshoot is also proportional to the input step size.
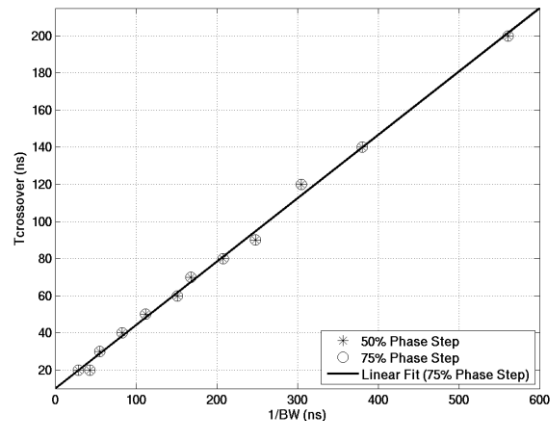

Fig. 3. Simulated $\tau_{crossover}$ vs. inverse of closed-loop −3dB bandwidth.

These relationships between time- and frequency-domain behaviors allow us to make fast time-domain measurements and then relate the results back to frequency-domain performance specifications. The circuit implementation presented in this paper shows that the PLL step response may be captured by an all-digital, on-chip finite state machine, allowing for fast PLL characterization. Silicon results indicate that this circuit could allow for power-on calibration of the PLL bandwidth and peaking for compensation of process variations.


Fig. 4. Simulated *MaxOvershoot* vs. closed-loop jitter peaking.

II. CIRCUIT DESIGN

The PLL under test (Fig. 5) is a standard integer-N charge-pump PLL. The only modification is the addition of loop measurement circuitry. The feedback divisor (N) is programmable from 5 to 63 although $N \geq 8$ during loop measurement tests. The charge-pump current, loop-filter resistance, and VCO gain are programmable to allow for bandwidth and peaking adjustments as well as jitter optimization. The PLL bandwidth may be configured from 3 to 25 MHz while the peaking may be varied from <1 to >4 dB. The VCO operates from 1.6 to 5 GHz. The expected reference clock frequency range is 100 to 200 MHz.
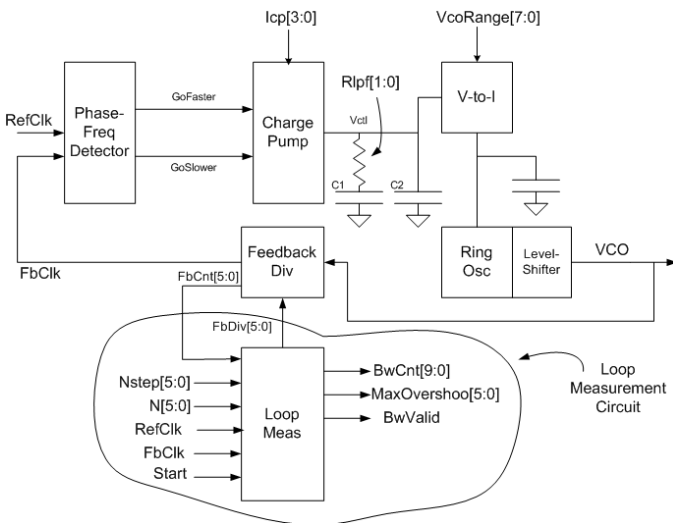

Fig. 5. PLL block-level diagram with loop measurement circuit.

A simple way to induce the required input phase step is to flip the polarity of the reference clock so its phase is advanced by half a clock cycle. A disadvantage to this approach is that the magnitude of the phase step is dependent on the reference clock duty cycle. This is undesirable because overshoot tests require a large and predictable input phase step. Instead, the circuit implementation presented here manipulates the feedback divisor to induce a known phase step. The circuit then automatically measures the resulting $\tau_{crossover}$ and *MaxOvershoot*. Fig. 6 shows a block diagram of the loop measurement test circuit. It includes three main units: control, $\tau_{crossover}$ detector, and *MaxOvershoot* detector. The control unit contains two synchronizers (to VCO clock), three edge detectors (rising and falling), and logic to enable the induced phase step. The $\tau_{crossover}$ detector includes a bang-bang phase detector, a phase-error change-of-sign detector, and a 10-bit counter. The *MaxOvershoot* detector contains a feedback count sampler, a comparator, and a maximum overshoot register.
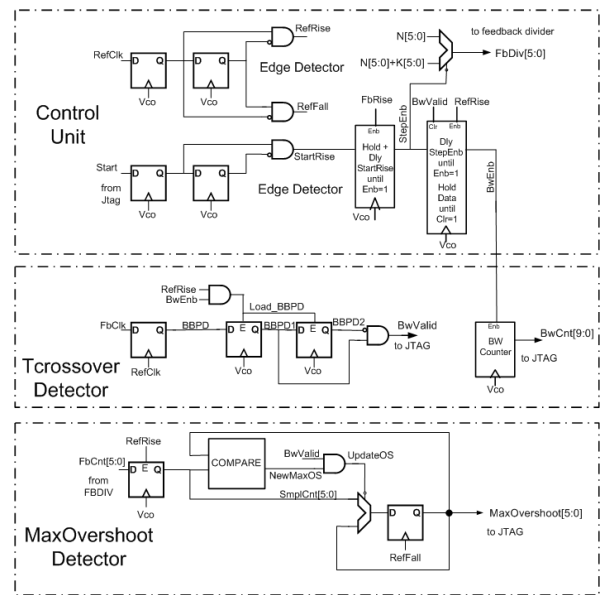

Fig. 6. Block diagram of loop measurement circuits in self-test mode.

Fig. 7 gives an example of how the input phase step is generated. The PLL is initially locked with $N=8$. When the *StepEnb* signal is asserted, $N=11$ is loaded into the incrementing feedback divider. One feedback clock cycle later, *StepEnb* is de-asserted and N is reset to the default value of 8. This process has the effect of delaying the feedback clock (*FbClk*) by 3 VCO clock cycles and instantaneously introducing a known phase error. At this point, the PLL begins to react to the induced phase error by increasing the VCO frequency.


Fig. 7. Feedback clock phase step and start of crossover-time measurement.

At the first rising *RefClk* after the phase step is applied, the *BwEnb* signal asserts to enable the crossover-time counter (*BwCnt[9:0]*) and begin the crossover-time measurement (Fig. 8). The bang-bang phase detector (*BBPD*) samples the *FbClk* level at every rising *RefClk*. Signals *BBPD1* and *BBPD2* are shifted versions of *BBPD*; they operate in the VCO clock domain and are updated once every reference clock cycle by the *Load_BBPD* signal. The state <*BBPD1*=1,*BBPD2*=0> indicates that the PLL has eliminated the induced phase error and that *FbClk* is now leading *RefClk*. The *BwValid* signal is set, halting the crossover-time counter and completing the crossover-time measurement. The 10-bit *BwCnt* value is converted to time with

$$\tau_{crossover} = \tau_{vco} \times (BwCnt - K) \qquad (1)$$

where τ_{vco} is the nominal VCO clock period and K is the induced step size in VCO clock cycles. K must be subtracted from the measurement results because the phase step causes the PLL to produce K additional VCO cycles during the re-lock process.


Fig. 8. Finding crossover-time.

When the *BwValid* signal is asserted, the maximum overshoot test begins (Fig. 9). In the overshoot test, the internal state of the feedback divider (*FbCnt[5:0]*) is sampled at every *RefRise* pulse where *RefRise* is a synchronized (to VCO) version of the rising edge of *RefClk*. *RefClk* is not used directly to sample the feedback divider count because it is not synchronous with the VCO clock. The sampled feedback divider count is placed in the *SmplCnt[5:0]* register.

During initial phase overshoot, the VCO speeds up, *FbClk* pulls ahead of *RefClk*, and the sampled feedback divider count increases in value. A circuit compares *SmplCnt[5:0]* to the previous maximum overshoot (*MaxOvershoot[5:0]*). If *SmplCnt[5:0]* is greater than *MaxOvershoot[5:0]*, then the UpdateOS signal is asserted and *SmplCnt[5:0]* replaces *MaxOvershoot[5:0]* at the next rising edge of *RefFall*, which is analogous to the aforementioned *RefRise* signal.


Fig. 9. Finding maximum overshoot.

Eventually, the VCO slows down and the sampled feedback divider count moves back toward zero. If the PLL exhibits ringing, then the sampled feedback divider count may continue past zero (undershoot) and begin recording values such as $N{-}1$, $N{-}2$, $N{-}3$, etc.. To filter these undershoot counts, the comparator ignores any *SmplCnt[5:0]* values greater than $N/2$.

The *RefClk* synchronizer latency must be subtracted from the measured *MaxOvershoot* count to calculate the actual maximum overshoot. The synchronizer latency in VCO clock cycles (N_{sync}) is measured in another test mode where *FbCnt[5:0]* is sampled by *RefRise* as previously described but no phase step is applied. The *MaxOvershoot* value measured in VCO cycles is converted to time using

$$\tau_{overshoot} = \tau_{vco} \times (MaxOvershoot - N_{sync}) . \qquad (2)$$

The precision of the overshoot detector is τ_{ref}/N, and so the measurement is less precise with small feedback divisors. In generating the input phase step, the feedback clock may be advanced instead of retarded. This allows for a larger phase step in cases in which the nominal feedback divisor is close to the maximum value of 63. A potential downside of advancing *FbClk* is that phase overshoots smaller than the *RefClk* synchronizer latency cannot be detected.

These step-response algorithms require that the PLL static phase error is less than the maximum overshoot. If not true, the required phase error sign-change does not occur, the bandwidth counter saturates at its maximum value, and the *BwValid* bit remains low. If the static phase error is large, then the *FbClk* phase can be advanced, forcing a phase error sign-change. The resolution of the bandwidth test is one reference clock period, and so the measurement becomes less precise as the PLL bandwidth approaches the reference clock frequency.

The loop measurement circuit can also be used to generate a time-trend of the PLL step-response, similar to a TIE plot. Instead of automatically detecting $T_{crossover}$ and *MaxOvershoot*, the feedback divider count is captured after exactly N reference clock cycles. By varying N from 1 to the maximum value of 63, the PLL step response may be plotted vs. time, as in Fig. 2. The loop measurement circuit may be used as a lock detector by repeatedly measuring N_{sync}. If it does not vary, then the PLL is locked. The static phase error may be estimated by comparing the measured N_{sync} to the expected synchronizer latency of 1–2 VCO cycles. By default, all loop measurement clocks are gated when not in use to minimize power. All flip-flops are of the sense-amplifier type for fast setup time and fast resolution of meta-stable signals.

III. EXPERIMENTAL RESULTS

While the loop measurement circuit has been successfully used with a wide range of PLL frequencies, the experimental results presented here focus on a PLL operating at 2.5 GHz with a 100-MHz reference clock. Programmable charge-pump currents and loop-filter resistances are used to vary the closed-loop bandwidth and jitter peaking. Table I shows the various loop-filter resistor and charge-pump settings as well as simulated and measured closed-loop bandwidth and jitter peaking. The measured results for cases 10–12 are nearly identical, probably due to premature saturation of the charge-pump current, although continuous-time loop equations are inaccurate at these low oversampling ratios. The unexpectedly similar measurement results for cases 1 and 2 could be due to second-order effects in the charge-pump at low

current settings. In general, the measured bandwidths are higher than the simulated values while the measured peaking is lower. Such differences between silicon and simulations motivated this work.

For each PLL setting in Table I, the loop measurement circuit captured $\tau_{crossover}$ and *MaxOvershoot* as described in Section II. Two phase step sizes were used: +13 and +19 VCO clock cycles. These phase steps are approximately 50% and 75% of the reference clock period, respectively. For each step size and PLL setting, the loop measurement test was conducted 25 times to verify repeatable results. These results confirm that run-to-run measurement variations are within the measurement precision. Three parts were tested, although — for reasons of clarity — the measurement results of a single part are plotted here.

<div align="center">TABLE I
COMPARISON OF SIMULATED AND MEASURED LOOP
PARAMETERS AT VARIOUS PLL SETTINGS</div>

Case	R_{lpf} (kΩ)	I_{cp} (μA)	Bandwidth (MHz)				Peaking (dB)			
			Sim-ulated	Measured			Sim-ulated	Measured		
				Part1	Part2	Part3		Part1	Part2	Part3
1	3.2	2.5	1.8	3.4	3.4	3.0	8.5	4.2	4.1	4.2
2	3.2	5	2.6	3.1	3.2	3.2	6.2	4.3	4.0	4.2
3	3.2	10	4.0	5.0	5.4	5.3	4.3	3.0	2.8	2.8
4	3.2	20	6.6	9.3	10.0	10.0	2.9	1.8	1.5	1.7
5	4.8	10	4.8	6.2	6.8	6.5	2.6	1.7	1.5	1.3
6	4.8	20	9.0	13.2	14.8	14.2	1.7	0.9	0.7	0.7
7	6.4	5	3.3	4.3	4.5	4.5	2.8	1.9	2.0	1.7
8	6.4	10	6.0	8.1	9.1	9.0	1.8	1.1	1.0	1.0
9	6.4	20	12.0	17.6	19.7	18.7	1.2	0.7	1.1	1.0
10	6.4	30	18.1	25.6	27.1	26.2	0.8	2.1	2.6	2.8
11	6.4	40	23.3	25.7	26.8	26.1	1.2	2.2	2.5	2.8
12	6.4	70	35.2	25.7	26.7	25.9	3.0	2.1	2.3	2.6

The measured $\tau_{crossover}$ (Fig. 10) shows the same inverse relationship to PLL bandwidth as simulated data. The slopes of the measured and simulated data are nearly identical, although the *y*-intercept of the measured data is approximately one reference clock period (10ns) higher than the simulated data. For a few PLL settings, the measured $\tau_{crossover}$ is slightly higher with the 75% step than with the 50% step, although the differences do not exceed the one reference clock precision of the test. This is true for all parts.
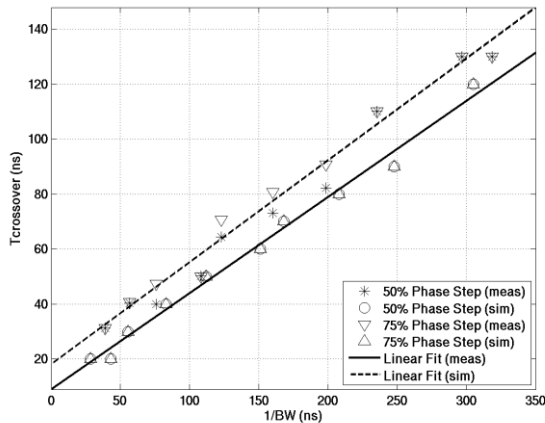

Fig. 10. $\tau_{crossover}$ vs. inverse of PLL bandwidth for feedback steps of +13 and +19 VCO clock cycles for Part 1.

In Fig. 11, the measured *MaxOvershoot* shows the same direct relationship to closed-loop jitter peaking as the simulated data. For both simulated and measured data, the slope of the overshoot vs. peaking curve is approximately 50% higher in the 75% phase step case compared to the 50% phase step case. All overshoot measurements except for one follow the expected curve within the measurement precision of one VCO clock. The exception: one measured overshoot (Part 2, case 2, 75% phase step) is one VCO clock (+14%) larger than predicted by the curve in Fig 11. While peaking in these plots is plotted in dB, *MaxOvershoot* is more accurately related to peaking plotted on a linear scale. However, for peaking values of 0.5−6.0dB, the curves in Fig. 11 remain nearly the same if peaking is plotted on a linear scale. This is because the relationship between $\log(x)$ and x is fairly linear in this peaking range.


Fig. 11. *MaxOvershoot* vs. jitter peaking for feedback steps of +13 and +19 VCO clock cycles for Part 1.

The simulated power for the loop measurements circuit is about 2.5 mW when operating at 2.5GHz and 1.2V power supply. The area is 2,750 μm², although it can easily be reduced by 40−50% by replacing some sense-amplifier flip-flops with smaller master-slave flip-flops and optimizing the overshoot comparator. Layout area is not a serious constraint in this design.

REFERENCES

[1] S. Kim and M. Soma, "An all-digital built-in self-test for high-speed phase-locked loops," *IEEE Trans. Circuits Systems II: Analog and Digital Signal Processing*, vol. 48, no. 2, Feb. 2001.
[2] B. Veillette and G. Roberts, "On-chip measurement of the jitter transfer function of charge-pump phase-locked loops," *IEEE J. Solid-State Circuits*, vol. 33, no. 3, pp. 483−491, Mar. 1998.
[3] S. Sunter and A. Roy, "BIST for phase-locked loops in digital applications," *Proc. International Test Conf.*, pp. 532−540, 1999.
[4] F. Gardner, *Phaselock Techniques*, 3rd ed., Wiley, p.102, 2005.

PCI Express is a registered trademark of the PCI-SIG Consortium.