Self-Biased High-Bandwidth Low-Jitter 1-to-4096 Multiplier Clock Generator PLL

John G. Maneatis, *Member, IEEE*, Jaeha Kim, *Student Member, IEEE*, Iain McClatchie, Jay Maxey, and Manjusha Shankaradas

***Abstract*—A self-biased phase-locked loop (PLL) uses a sampled feedforward filter network and a multistage inverse-linear programmable current mirror for constant loop dynamics that scale with reference frequency and are independent of multiplication factor, output frequency, process, voltage, and temperature. The PLL achieves a multiplication range of 1–4096 with less than 1.7% output jitter. Fabricated in 0.13-μm CMOS, the area is 0.182 mm^2 and the supply is 1.5 V.**

***Index Terms*—Adaptive bandwidth, analog circuits, clock generation, clock multiplication, frequency synthesis, phase-locked loop (PLL), self-biased.**

I. Introduction

ONE CHALLENGE in designing phase-locked loops (PLLs) for application-specific integrated circuits (ASICs) is providing ample flexibility for a wide variety of applications, including processors and video/chip interfaces. PLLs commonly are used to take low-frequency off-chip clocks, typically from crystals, and generate high-frequency on-chip clocks. The diversity of ASIC applications has also led to diversity in operating frequencies and multiplication factors required from PLLs.

For each PLL output frequency and multiplication factor, the loop parameters must be adjusted to minimize jitter and to guarantee stability. There are two jitter parameters of interest. One is long-term jitter, which is the deviation over time in the output clock edge time locations from those of an ideal clock output that is perfectly periodic. The other is period jitter, which is the variation over time in the period of the output clock. For a clock generator PLL, the output clock should track the input clocks as close as possible to minimize long-term jitter. It is also important to minimize the amount of period jitter.

These objectives pose a set of requirements on the loop parameters of the PLL. The loop bandwidth, which describes the response rate of the PLL, should be about 1/20 of the reference frequency. The damping factor, which describes the stability, should be about one. The third-order pole, which helps minimize period jitter, should be set at about 1/2 of the reference frequency. All of these loop parameters depend on specific circuit parameters, such as the charge pump current and the loop filter resistance. Thus, these parameters must vary with output frequency and multiplication factor.

The diverse values of output frequency and multiplication factor can be addressed by designing a different PLL for each ASIC. This strategy makes it easier to meet constrained target specifications with less challenging circuits, but verifying all the designs in silicon for the ASICs that a company plans to build would be time consuming and costly. A better strategy is to create a single PLL design that can be used for clock generation on a large set of ASICs. With only one design, verification in silicon is much easier, but the design becomes more difficult as loop parameters must adjust automatically to satisfy a wide range of output frequencies and multiplication factors.

Self-biased PLLs [2] can solve part of the problem by adjusting for different output frequencies. Specifically, they achieve a fixed loop-bandwidth-to-reference-frequency ratio and damping factor, which are largely independent of process, voltage, and temperature. This property allows the bandwidth to be set to a precise fraction of the reference frequency independent of the actual reference frequency, which will minimize long-term jitter over a wide reference frequency range. However, self-biased PLLs do not adjust for different multiplication factors. In particular, the bandwidth-to-reference-frequency ratio and the damping factor both vary with the multiplication factor. Also, with an additional third-order pole, the pole-frequency-to-reference-frequency ratio will also vary with the multiplication factor. To handle a large multiplication range, all of these ratios should be fixed and independent of the multiplication factor.

This paper describes a self-biased clock generator PLL capable of multiplying by 1 to 4096 with near-constant period jitter over the whole range [1]. The PLL extends the self-biased PLL architecture with a new loop filter structure that produces constant loop dynamics that scale with reference frequency and are virtually independent of the multiplication factor, output frequency, process, and environmental conditions.

This paper begins by reviewing the fundamentals of a self-biased PLL design and how it obtains tracking loop dynamics. Pattern jitter, a form of period jitter caused by multiplication, is discussed in Section III. Section IV presents a loop filter architecture that solves the scaled N problem while addressing pattern jitter. A number of key circuits used inside the PLL design are described in Section V. Finally, some experimental results demonstrating the effectiveness of the clock generator PLL architecture in minimizing output jitter are presented in Section VI.

Manuscript received April 15, 2003; revised June 22, 2003.

J. G. Maneatis and I. McClatchie are with True Circuits, Inc., Los Altos, CA 94022 USA (e-mail: maneatis@truecircuits.com).

J. Kim was with True Circuits, Inc., Los Altos, CA 94022 USA. He is now with Seoul National University, Seoul 151-742, Korea.

J. Maxey and M. Shankaradas are with Texas Instruments Incorporated, Dallas, TX 75243-0199 USA.

Digital Object Identifier 10.1109/JSSC.2003.818298

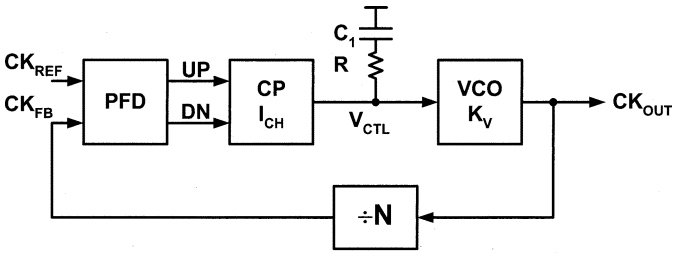

Fig. 1. Classic second-order PLL.

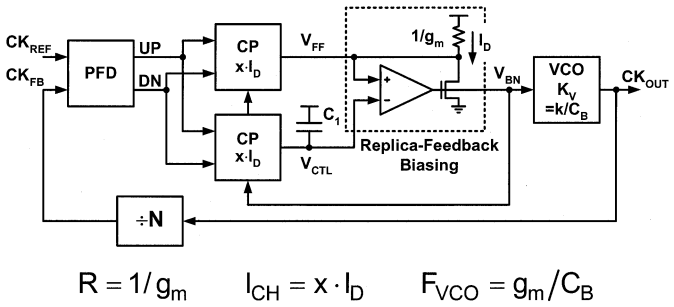

Fig. 2. Simple self-biased PLL.

II. Self-Biased PLL Fundamentals

Before considering a self-biased PLL, it is helpful to first review a classic second-order PLL, shown in Fig. 1. This PLL is composed of a phase-frequency detector (PFD), charge pump, loop filter, voltage-controlled oscillator (VCO), and a feedback divider. When in lock, the PLL generates an output frequency that is N times the reference frequency. It does this by adjusting the VCO frequency until it detects no phase or frequency difference between the reference and divided output clocks. Because of the integration of charge on the loop filter and the integration of phase in the VCO, the system has a second-order closed-loop response.

The frequency-domain phase response for the classic second-order PLL, given by the ratio of the output phase $P_{\rm O}(s)$ to the input phase $P_I(s)$, can be represented in standard form as

$$\frac{P_{\rm O}(s)}{P_{\rm I}(s)} = N \cdot \frac{1 + 2 \cdot \zeta \cdot \left(\frac{s}{\omega_N}\right)}{1 + 2 \cdot \zeta \cdot \left(\frac{s}{\omega_N}\right) + \left(\frac{s}{\omega_N}\right)^2}$$

where ω_N, defined as the loop bandwidth (rad/s), is given by

$$\omega_N = \sqrt{\frac{1}{N} \cdot I_{\rm CH} \cdot K_V \cdot \frac{1}{C_1}}$$

and ζ, defined as the damping factor, is given by

$$\zeta = \frac{1}{2} \cdot \omega_N \cdot R \cdot C_1.$$

The loop bandwidth characterizes the response rate of the system and the damping factor characterizes its stability. The system is underdamped with damping factors less than one and, thus, less stable. ω_N and ζ are functions of the various circuit parameters which are typically fixed for a particular design. These fixed parameters cause both ω_N and ζ to also be fixed. Ideally, ω_N should scale with $\omega_{\rm REF}$ to handle a wide frequency range.

A self-biased PLL, shown in Fig. 2, solves this problem with three important differences. First, rather than a fixed resistor in series with a capacitor, a self-biased PLL uses a $1/g_m$ resistance from the VCO bias generator that is proportional to the output period. Second, instead of a single charge pump to drive an *RC* network, two separate charge pumps are used to drive the capacitor and resistor separately, where the voltages are summed inside the VCO bias generator. Finally, the charge pump current is scaled from a current generated inside the VCO in order to make the open-loop gain related to the output frequency. By making these changes, the bandwidth-to-reference-frequency ratio will be proportional to the square root of $x \cdot N$, and the damping factor will be proportional to the square root of x/N, where x is the charge pump current scale factor. Thus, both results are constant with output frequency, which is desired, but not with N, which presents a problem.

However, before one can consider how to address this frequency multiplication scaling issue, one needs to first consider another problem related to frequency multiplication, called pattern jitter or spurious noise.

III. Pattern Jitter Issues

Pattern noise is caused by the phase corrections that occur on every rising edge of the reference clock. These phase corrections can briefly disrupt the control voltage $V_{\rm CTL}$ and, in turn, nearby output cycles, as illustrated in Fig. 3(a). The other output cycles will be unaffected. This noise pattern on $V_{\rm CTL}$ and the resultant period jitter will repeat on every reference cycle or N output cycles. Pattern jitter is typically caused by charge pump imbalances or leakage which gives rise to static phase offsets between the reference clock and the output clock, similar to that shown in the figure. A less periodic form of pattern jitter can result from jitter in the reference clock. The phase corrections resulting from this jitter can be concentrated in one of the N output cycles, giving rise to a substantial amount of period jitter relative to the much shorter output period.

This pattern jitter problem is typically solved by adding a shunt capacitor in the loop filter to create a third-order pole, which extends the disturbance on $V_{\rm CTL}$ with reduced amplitude over many output cycles, as illustrated in Fig. 3(b). The reduced noise amplitude on $V_{\rm CTL}$ makes the output cycles less distorted from one another, reducing the overall period jitter. Ideally, the number of output cycles for which the disturbances are extended should scale with N to maximize the filtering benefit over a wide range of N. Unfortunately, the number of cycles is fixed with a fixed capacitor. Thus, for large N, the number of cycles will be too small, leading to pattern jitter, while for small N, the number of cycles will be too large, which will lead to instability, as the increased loop bandwidth and this fixed third-order pole become too close.

Instead of a shunt capacitor, a switched capacitor or sampled filter network can be used to address the filter scaling issue, like that demonstrated in [3] and [4]. The basic idea is to scale down the amplitude of the error signal on $V_{\rm CTL}$ and spread it uniformly over exactly N output cycles with the same time integral, as illustrated in Fig. 3(c). The N output cycle duration is controlled directly by the switching network. The signal on $V_{\rm CTL}$

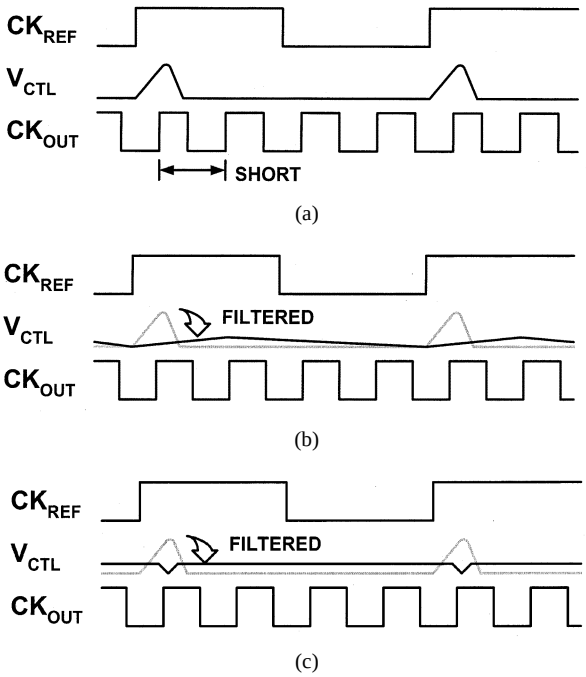

Fig. 3. Patten jitter on output clocks with (a) simple second-order PLL, (b) added shunt capacitor, and (c) sampled feedforward network.

due to periodic phase errors will be mostly constant, except for small disruptions as the network switches to the next sample. Thus, the output cycles will be virtually undisturbed with very little period jitter. The next section will describe a simple solution using this approach that is compatible with self-biased PLLs.

IV. Clock Generator PLL Architecture

To see how a sampled filter network can be constructed, consider the original loop filter network from the previously discussed self-biased PLL, shown in Fig. 4(a). A self-biased PLL uses two charge pumps to drive the capacitor and resistor separately. The path with the capacitor is the integral control path and does not need additional filtering. The path with the resistor is called the proportional or feedforward signal path. This path has no filtering and is the problem. Thus, a filter network should be added between the charge pump and the bias generator in the feedforward path.

A. Sampled Feedforward Network

The feedforward filtering can be performed by sampling the phase error and generating a proportional current that is held constant for N output cycles. The filter network shown in Fig. 4(b) accomplishes this filtering. It stores the output charge from the charge pump on capacitor C_2. This action generates a constant error voltage that drives the g_m stage to produce a feedforward current that is constant for N output cycles. This feedforward current develops a correction voltage across the $1/g_m$ resistor inside the bias generator, where it is summed with the control voltage $V_{\rm CTL}$.

The switch in Fig. 4(b) resets the C_2 capacitor voltage to a zero bias voltage level $V_{\rm RST}$ at the end of the reference cycle before the next phase comparison. This voltage is the point where

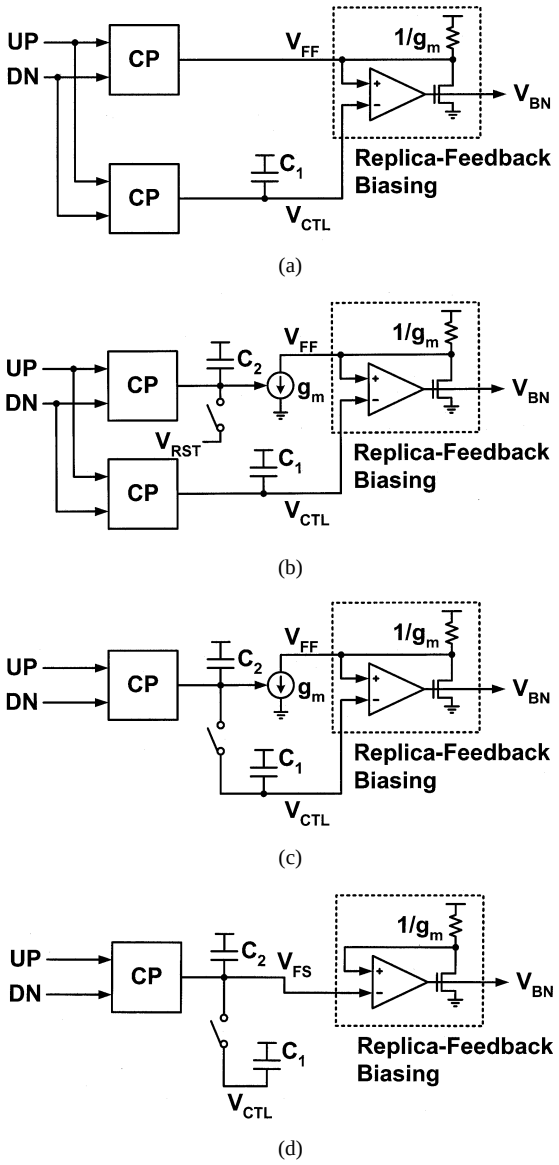

Fig. 4. PLL loop filter network (a) for simple self-biased PLL, (b) with sampled feedforward network, (c) with sampled feedforward network with C_2 reset to C_1, and (d) with sampled feedforward network with C_2 driven directly into amplifier.

the g_m stage will produce no current. Based on the g_m stage that is used in the implementation, this reset voltage must be equal to the control voltage. Rather than attempting to buffer the control voltage and possibly introducing error, C_2 can be reset directly to the control voltage, as illustrated in Fig. 4(c). This change eliminates the need for the integral charge pump, since the charge transferred from C_2 to C_1 in the reset process will be similar to the charge that the integral charge pump would have transferred to C_1.

A further optimization to the loop filter is shown in Fig. 4(d). Since the voltage $V_{\rm FS}$ across C_2 is reset to $V_{\rm CTL}$, and $V_{\rm FS} - V_{\rm CTL}$ is summed with $V_{\rm CTL}$ inside the bias generator resulting in $V_{\rm FS}$, $V_{\rm FS}$ can be simply driven to the $V_{\rm CTL}$ input of the bias generator instead of $V_{\rm CTL}$, eliminating the need for the g_m stage. This solution is much simpler, but does not allow any voltage scaling to be performed inside the transconductance stage, which can provide more flexibility on C_2 capacitor sizing.

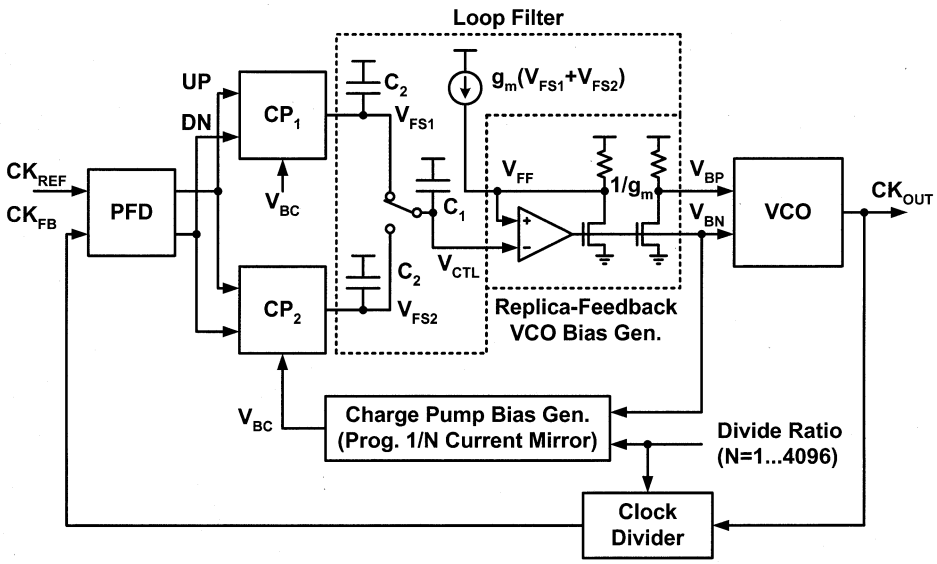

Fig. 5. Self-biased clock generator PLL architecture.

B. Loop Dynamics

In order to understand how this new loop filter affects the loop dynamics, consider how the feedforward gain has changed. Since the output current is generated for N output cycles, the output charge $Q_{\rm O}$ will be proportional to N times the input charge Q_I, so the added gain is proportional to N. Also, since the damping factor is proportional to the feedforward gain, it should be scaled by this added gain.

Thus, with the new loop filter, $\omega_N/\omega_{\rm REF}$ is unchanged from before, but the damping factor ζ, now given by

$$\zeta \sim \frac{Q_{\rm O}}{Q_I} \cdot \sqrt{\frac{x}{N}} \sim \sqrt{x \cdot N}$$

is multiplied by this added gain factor which is proportional to N. This change makes both ζ and $\omega_N/\omega_{\rm REF}$ proportional to the square root of $(x \cdot N)$, where x is the charge pump current scale factor. Thus, to keep both constant, x can be simply set to $1/N$, or equivalently, the charge pump current can be scaled with $1/N$.

A more careful analysis, similar to that in [2], performed in Section IV-D, will show that the bandwidth-to-reference-frequency ratio and the damping factor are given by

$$\frac{\omega_N}{\omega_{\rm REF}} = \frac{1}{2 \cdot \pi} \cdot \sqrt{\frac{C_B}{C_1'}}$$

$$\zeta = \frac{1}{4} \cdot \frac{\sqrt{C_B \cdot C_1'}}{C_2}$$

where C_B is the equivalent VCO capacitance and $C_1' = C_1 + C_2$ (if $m+1$ C_2 capacitors are used as discussed later, $C_1' = C_1 + (m+1) \cdot C_2$). Thus, they are both constants times the square root of ratios of capacitors. Both results are independent of $F_{\rm OUT}$ and N, as well as process, voltage, and temperature.

At this point, the frequency multiplication scaling problem has been completely solved: the loop dynamics are constant and independent of N, and the pattern jitter is minimized for all N.

C. Complete Self-Biased Clock Generator PLL

The complete self-biased clock generator PLL design, based on the loop filter structure in Fig. 4(c), is shown in Fig. 5. In order to allow ample time to reset the C_2 capacitor, the capacitor is duplicated, where one is in reset while the other samples charge and generates an error voltage for a complete reference cycle. The capacitors alternate function each reference cycle so that a continuous error signal is generated. Rather than trying to switch the output of one charge pump to the two capacitors, two charge pumps which are alternately enabled are used for simplicity. Also, the outputs of both C_2 capacitors are summed as currents to further minimize the need for switches. Finally, the charge pump current bias is generated by the programmable $1/N$ current mirror using the multiplication factor input.

Adding output switches to switch between the capacitor voltages as done in [4] has the apparent advantage of avoiding any extra pattern jitter caused by charge pump output current pulses. Such a scheme would ideally allow the output error signal to change monotonically between reference comparison cycles. However, the added switches will disturb the output error signal by an amount that does not depend on N due to charge injection. The proposed filter network will also cause a disturbance in the output error signal, represented as the sum of the C_2 capacitor voltages $(V_{\rm FS1} + V_{\rm FS2})$, due to a difference between the charge pump charging and reset switch discharging rates for the two capacitors. Even though the charge pump current is scaled inversely with N, which would otherwise reduce the disturbances in the error signal, the worst-case phase errors tend to increase linearly with N because of the reduced bandwidth, which tracks the reference frequency, and proportionally increased tracking jitter. Thus, both approaches can lead to disturbances in the output error signal.

Because the switch discharge rate is faster than the charge pump charging rate with the scaled-down charge pump currents in the proposed filter network, the disturbance in the error signal will be in the direction of less error. It is important to note that minimizing the disturbance in sum of the C_2 capacitor voltages

will reduce the maximum change in the output period from one cycle to the next (cycle-to-cycle jitter). However, the maximum deviation in the output period over all cycles (period jitter) is minimized as long as the error signal is spread out over N output cycles, independent of a disturbance caused by the sum of the C_2 capacitor voltages ramping up from a zero bias level each reference cycle. This result allows the proposed network to produce less period jitter than one with added output switches.

A further optimization can be made to avoid any disturbance in sum of the C_2 capacitor voltages without introducing extra output switches. Rather than immediately resetting one of the C_2 capacitors at the beginning of the phase comparison cycle by closing the corresponding reset switch, the polarity of the charge pump driving the capacitor could be reversed so that the charge pump discharges the capacitor at the same rate that the other charge pump charges the other capacitor. The reset switch would then be closed after the charge pump turns off to reset any residual error voltage, exposing the net change in the sum of the C_2 capacitor voltages without any added disturbance.

More filtering can be obtained by using $m+1$ C_2 capacitors where each is active for m cycles with a gain of $1/m$. Also, the filter network from Fig. 4(d) can be implemented with two or more C_2 capacitors by subdividing the differential pair and tail current source of the amplifier in the bias generator. The negative inputs can be connected to each C_2 capacitor while the positive inputs would be shorted together and connected to the $1/g_m$ resistor as before.

D. Detailed Analysis

The exact relationships for the bandwidth-to-reference-frequency ratio and the damping factor can be derived from the equivalent relationships for a simple self-biased PLL, similar to that shown in Fig. 2. For a simple self-biased PLL [2]

$$\frac{\omega_N}{\omega_{\mathrm{REF}}} = \frac{\sqrt{x \cdot N}}{2 \cdot \pi} \cdot \sqrt{\frac{C_B}{C_1}}$$
$$\zeta = \frac{y}{4} \cdot \sqrt{\frac{x}{N}} \cdot \sqrt{\frac{C_1}{C_B}}$$

where x is the ratio of the charge pump to the buffer bias current, y is the ratio of the $1/g_m$ resistance in the bias generator to that in the VCO buffer stages, and C_B is the equivalent VCO capacitance. For the self-biased clock generator PLL based on the sampled feedforward network, x is set to $1/N$ and y is equal to Q_{O}/Q_I for the feedforward network. Since C_1 is allowed to charge share with C_2 during the reset phase, C_1 becomes $C_1 + C_2$.

The exact relationship for Q_{O}/Q_I can be derived as follows. The charge pump dumps some charge Q_I onto capacitor C_2, which generates a voltage that drives the g_m stage, and produces a current for a duration of T_{REF}, so that

$$Q_{\mathrm{O}} = \frac{Q_I}{C_2} \cdot g_m \cdot T_{\mathrm{REF}}.$$

For a simple self-biased PLL, C_B is defined as $C_B = T_{\mathrm{VCO}} \cdot g_m$, so that

$$T_{\mathrm{REF}} = N \cdot \frac{C_B}{g_m}.$$

Thus

$$\frac{Q_{\mathrm{O}}}{Q_I} = N \cdot \frac{C_B}{C_2}$$

which is proportional to N as previously concluded. Substituting these relationships, the bandwidth-to-reference-frequency ratio and the damping factor are given by

$$\frac{\omega_N}{\omega_{\mathrm{REF}}} = \frac{\sqrt{\left(\frac{1}{N}\right) \cdot N}}{2 \cdot \pi} \cdot \sqrt{\frac{C_B}{C_1}} = \frac{1}{2 \cdot \pi} \cdot \sqrt{\frac{C_B}{C_1}}$$
$$\zeta = \frac{N \cdot \frac{C_B}{C_2}}{4} \cdot \sqrt{\frac{\left(\frac{1}{N}\right)}{N}} \cdot \sqrt{\frac{C_1'}{C_B}} = \frac{1}{4} \cdot \sqrt{\frac{C_B \cdot C_1'}{C_2}}$$

both of which are independent of F_{OUT} and N, as well as process, voltage, and temperature.

V. Key PLL Circuits

The next sections discuss some of the actual circuits used inside the PLL, including the filter network, the programmable $1/N$ current mirror, and the VCO. This self-biased clock generator technique can be applied to most VCO circuit families. In fact, we have implemented it with four different families. In this implementation, the VCO circuits were similar to those previously published [5].

A. Self-Biased Sampled Filter Network

The circuits for the filter network are shown in Fig. 6. The g_m stages are implemented with the same half-buffer replica stages used inside the VCO bias generator. The select block alternately enables the charge pumps on opposite reference cycles as the corresponding reset switches. The charge pumps are enabled by gating the UP and DN input signals just before the beginning of the comparison cycle in order to maximize the period that V_{FS1} (V_{FS2}) stays constant.

The reset switches used in the filter network pose a small problem because they must effectively switch bias voltages that can range from ground to V_{DD}. To solve this problem, a simple nMOS pass gate is used with a bootstrapped gate voltage that ranges from V_{CTL} to approaching $V_{\mathrm{DD}} + V_{\mathrm{CTL}}$, as shown in Fig. 7. This circuit provides a constant ON gate bias of V_{DD}. The gate bias is generated by a simple bootstrapped circuit. The V_{CTL} level is generated by using another VCO replica stage. The circuit works by alternately driving the lower voltage level to V_{CTL} and then allowing the capacitors to bootstrap the voltage up by V_{DD} ignoring charge sharing. This solution makes it possible to properly switch bias voltages independent of their common-mode level, eliminating a possible supply voltage headroom constraint.

B. Inverse-Linear Current Mirror

In order to scale down the charge pump currents by a factor of N, a programmable current mirror is needed that can implement

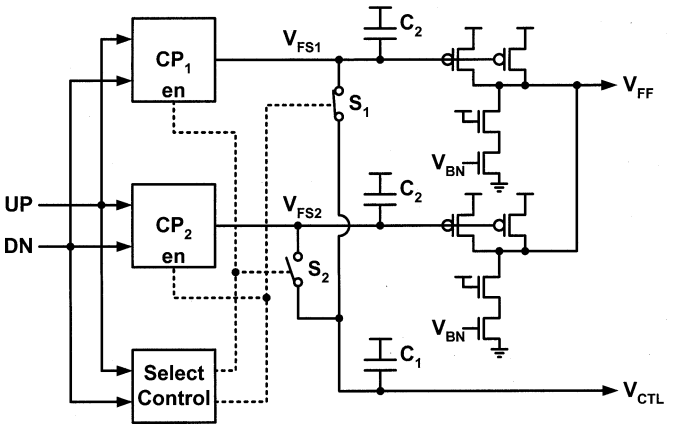

Fig. 6. Sampled feedforward circuits.

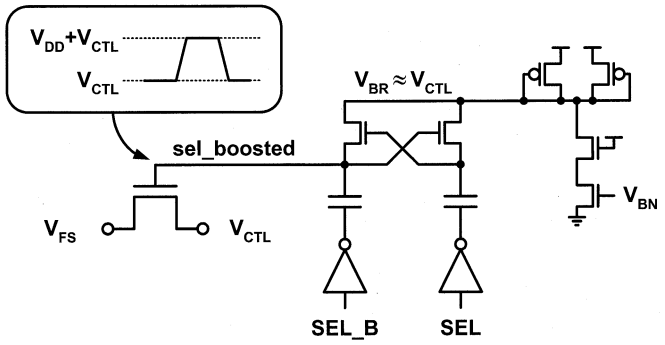

Fig. 7. Feedforward network bias switches.

the inverse-linear relationship. A simple way of accomplishing this task is to use switches to adjust the effective device size on the input side of a current mirror, as shown in Fig. 8(a). The individual input devices are binary weighted to allow any integer value of N to be programmed. However, to support a range of 1–4096, twelve binary-weighted legs are needed with a device size ratio of 2048 : 1, which will require too much area. A better solution would be to somehow segment the input side of the current mirror into different device groups operating from different gate biases so that only a small device size ratio is needed.

To see how this result can be accomplished, it is instructive to review how to make a multistage linear current mirror. In this case, multiple groups of binary-weighted devices operating from different but related current biases can be used so that they operate at different current densities. Fig. 8(b) shows an example of a multistage linear current mirror with two device groups of three devices and two current mirrors to establish a 1/8 current ratio between the two groups. This example can cover a range of 0–63 with a maximum device size ratio of 4:1. While this circuit example gives us a solution for a linear current mirror, the PLL actually needs an inverse-linear current mirror.

Consider the devices inside the box in Fig. 8(b). They implement a programmable-width device with some gate bias which is used as a current source. To create an inverse-linear programmable current mirror, this complex device can be diode connected and used as the input side of a current mirror, as shown in Fig. 8(c). It is important to note that the gate biases

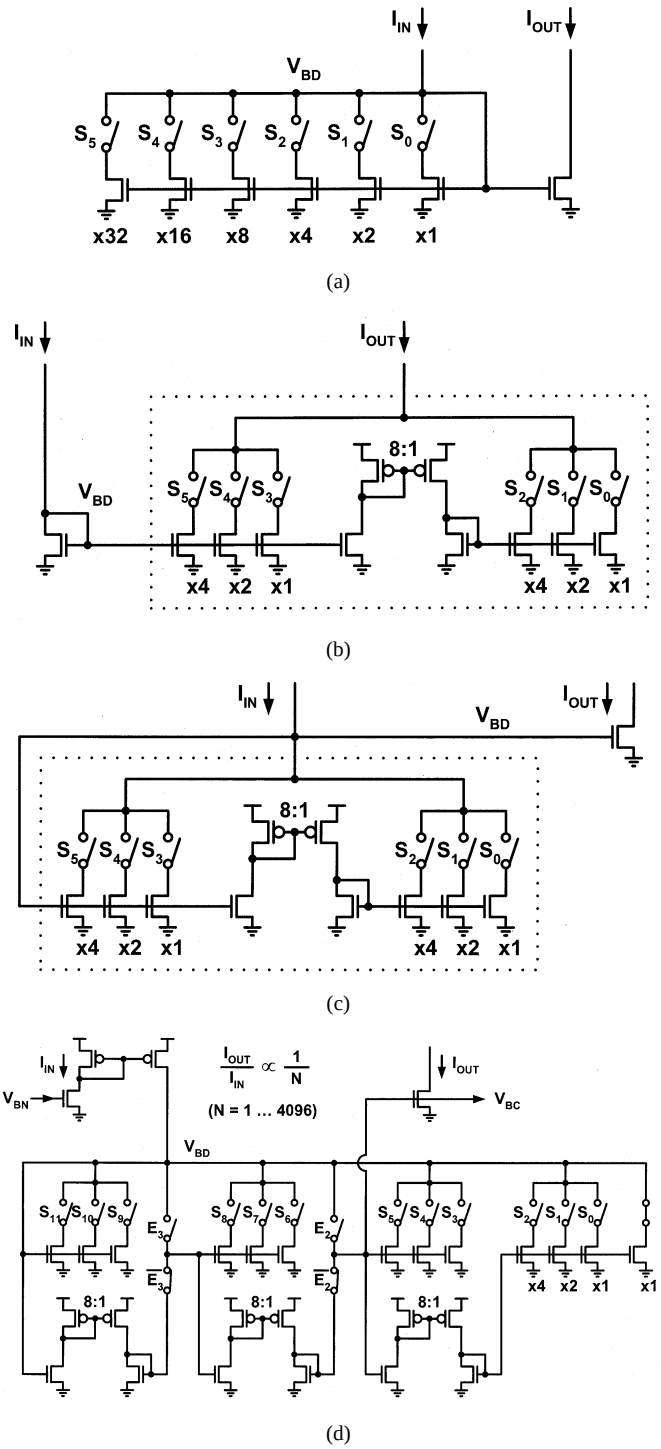

Fig. 8. Programmable current mirrors. (a) Simple inverse-linear current mirror. (b) Multistage linear current mirror. (c) Multistage inverse-linear current mirror. (d) Complete multistage inverse-linear current mirror.

from any device groups can be used to drive the current sources in order to obtain additional fixed scaling factors. Also, even though the feedback is more complicated than that for diode-connected devices, the network is completely stable as long as the gain stages used are gain reducing and not increasing.

Fig. 8(d) shows the complete programmable current mirror used in the PLL to scale the charge pump currents. It is similar to that in Fig. 8(c), but four groups instead of two are used.

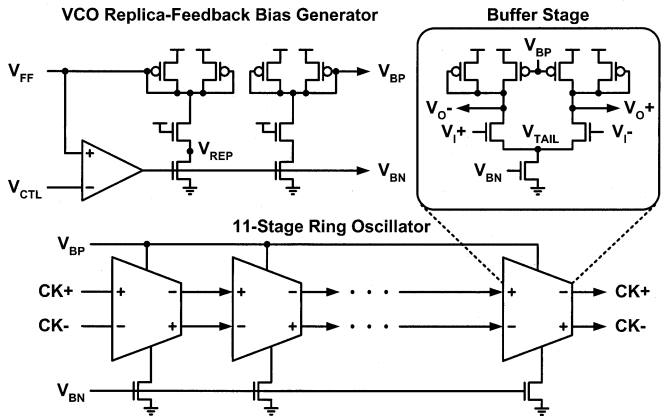

Fig. 9. VCO circuits.

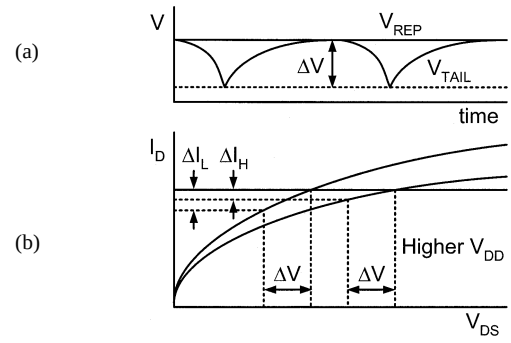

Fig. 10. Differential pair (a) tail node voltages and (b) tail currents as a function of time.

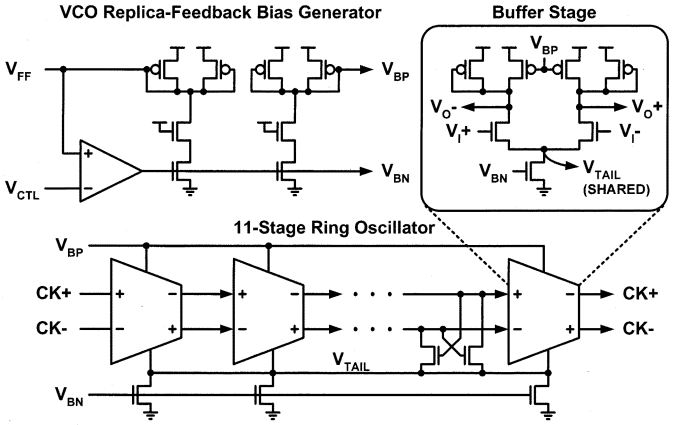

Fig. 11. Modified VCO circuits with shorted differential pair tail nodes.

As such, the mirror will divide the input current by N, where N ranges from 1–4096. Also, a fixed least significant bit (LSB) device is added in order to offset the N programming by one and so that the block does not attempt to output an infinite current. In addition, the switches E_2 and E_3 are added to bypass unused stages when programming small N. For small N, the first two gain stages would do nothing except reduce the gate voltage for the last two current-source groups. However, the voltage $V_{\rm BD}$ needed at the input of the gain-reducing stages could be excessive for moderate output current levels, which would drive the input pMOS current-source device into the linear region. By bypassing the unused gain stages, the voltage needed at $V_{\rm BD}$ is reduced for the same output current level. Since the output bias is tapped from the second to last group, the last gain-reducing stage cannot be bypassed.

C. Voltage-Controlled Oscillator

Another circuit with some challenges is the VCO. The VCO used in this design is a modified version of the differential ring oscillator with replica-feedback biasing from [5], shown in Fig. 9. The VCO bias generator uses an amplifier to establish a bias current such that the voltage across a replica load element equals $V_{\rm CTL}$. Since the I–V characteristics of the load element do not depend on $V_{\rm DD}$, the bias current will also be independent of $V_{\rm DD}$. By operating with a constant current, the VCO can generate a frequency that does not depend on supply voltage. This supply voltage independence helps to minimize the jitter produced by supply noise for the PLL.

However, for the VCO current to remain constant, the voltage at the differential pair tail nodes (in the VCO) must match the same point in the bias generator replica. Furthermore, these tail nodes tend to oscillate with the other VCO nodes, making the voltages match less well. Fig. 10(a) shows a plot of the $V_{\rm TAIL}$ node from the ring oscillator and the $V_{\rm REP}$ node from the bias generator as a function of time. Because of the oscillations, there is some range of deviation ΔV in these voltage levels. Fig. 10(b) shows these voltage levels mapped onto the current-source devices' $I_{\rm D}$ versus $V_{\rm DS}$ characteristics and the resultant drain current. The solid line is the current in the replica based on $V_{\rm CTL}$. The first group of dashed lines on the left show the voltages and resultant currents at a low supply voltage. Because of the finite output resistance of the current-source device, the buffer currents actually oscillate over some range ΔI_L. If the supply voltage increases, shown by the dashed lines at the right of the plot, the replica will adjust the gate bias to keep the current in the bias generator constant at the same level. The range of currents inside the VCO, ΔI_H, becomes smaller because of the decrease in slope of the I–V characteristics. This difference in average buffer current at different supply voltages gives rise to a supply voltage frequency sensitivity. This sensitivity makes the PLL more jitter sensitive to supply noise.

The root cause of the difference in average buffer current is the nonlinear output conductance of the current-source device, or, equivalently, the changing slope in its I–V characteristics. A simple solution to resolve this problem, shown in Fig. 11, is to short all of the tail nodes in the VCO so that their voltage is more or less constant and matching more closely to that in the bias generator. With all of the tail nodes tied together, the differential ring becomes two single-ended rings. To ensure differential operation, the input transistors of one prior stage pair are split and crossconnected to the stage outputs. Other cross-coupling points can be used to marginally increase or decrease the oscillation frequency [5].

Fig. 12 shows the simulated frequency as a function of supply voltage for three different control voltages at worst-case process and temperature. The solid curves are after shorting and the dashed curves are before shorting. As evident from the flatness of the solid curves, this change substantially improves the static supply-noise rejection of the VCO. Given a target minimum supply voltage of 1.2 V, the VCO can be operated at

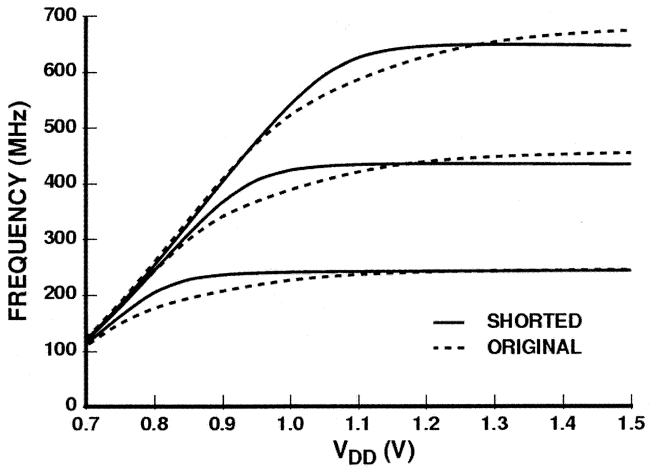

Fig. 12. VCO frequency as a function of supply voltage for three control voltages before shorting (dashed line) and after shorting (solid line) differential pair tail nodes.

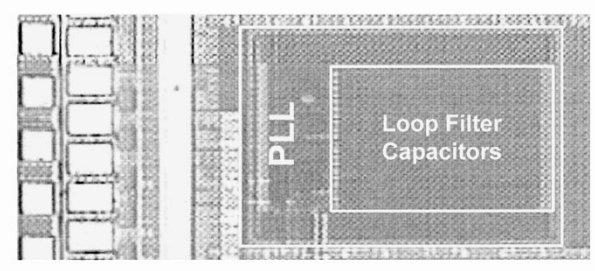

Fig. 13. Die micrograph of the self-biased clock generator PLL.

TABLE I
PLL SPECIFIED DESIGN TARGETS

Process Technology	0.13μm N-well CMOS
Nominal Supply Voltage	1.5V (designed for 1.2V)
Total Occupied Area	0.38 x 0.48mm^2
VCO Frequency Range	30 ~ 650 MHz
Multiplication Factor Range	N = 1 ~ 4096
Power Dissipation	7mW @ 240 MHz, 1.5V

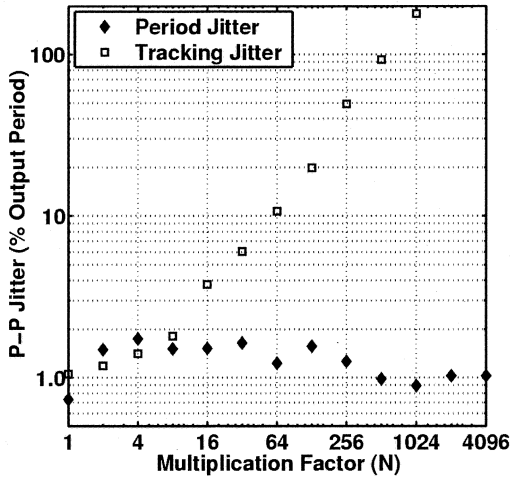

Fig. 14. Jitter versus multiplication factor at fixed 240-MHz output.

about three times the frequency of the original circuit with good supply-noise rejection.

VI. EXPERIMENTAL RESULTS

This PLL was implemented in a generic 1.5-V 0.13-μm n-well CMOS process. A micrograph of the fabricated PLL is shown in Fig. 13 and the specified design targets are summarized in Table I. While the nominal operating voltage for the PLL was 1.5 V, it was designed to operate down to 1.2 V to provide a 10% dc and 10% ac noise margin. The specified VCO frequency range is 30–650 MHz under worst-case conditions, which was divided by two before the output to improve duty cycle. The underlying VCO circuits could easily have been configured to run at much higher frequencies if the PLL application required them. Because the VCO supports a wide frequency range and the PLL has a tracking bandwidth, the PLL can be operated at frequencies well below 30 MHz, similar to that described previously [2]. The frequency multiplication range is 1–4096.

The focus of the measured results is on jitter in order to explore the effectiveness of the sampled feedforward network. Fig. 14 is a plot of the measured peak-to-peak tracking jitter and period jitter as a function of the multiplication factor N for a fixed output frequency of 240 MHz. Both the jitter and multiplication factor are plotted on a log-log scale to cover the large dynamic range of N. The period jitter is fairly constant over the complete range, mostly between 1% and 1.5% of the output period, which corresponds to 42 to 63 ps at 240 MHz. This period jitter data clearly demonstrates the effectiveness of the sampled feedforward network. Without it, one would expect the period jitter to scale with N or, equivalently, be a fixed faction of the reference frequency, because the magnitude of the steady-state proportional signal, due to various noise sources, would scale with the reference period, yet be dumped into only one of N output periods.

Fig. 14 also shows that the tracking jitter scales linearly with N, as expected, since the loop bandwidth scales inversely with N to be a constant fraction of the reference frequency. This relationship results because for low-frequency noise, the VCO will accumulate phase error for a duration that is inversely proportional to the bandwidth. The tracking jitter is less than the period jitter for low multiplication factors because it is measured between the edges of the reference clock and the output clock, and thus, does not include the effect of the period jitter on the output edges between reference edges.

Table II summarizes rms and peak-to-peak jitter levels under various operating conditions. The sensitivity to supply noise, like other implementations using the same underlying VCO circuits, is very good.

TABLE II
PLL JITTER MEASUREMENT SUMMARY

Period Jitter (quiescent)	30.5ps (p-p), 4.0ps (rms) @ N=1
@ 240MHz Output	37.2ps (p-p), 4.2ps (rms) @ N=1024
	42.7ps (p-p), 4.3ps (rms) @ N=4096
	72.6ps (p-p), 18.8ps (rms) @ worst-case
Period Jitter (noise*)	44.6ps (p-p), 4.9ps (rms) @ N=1024
Tracking Jitter (quiescent)	43.9ps (p-p), 4.6ps (rms) @ N=1
@ 240MHz Output	7.47ns (p-p), 939ps (rms) @ N=1024
Tracking Jitter (noise*)	7.62ns (p-p), 1.07ns (rms) @ N=1024
Reference Sidebands	-35 dBc @ N=1024, 240MHz Output

* At 240MHz output with 100mV of 100KHz square wave supply noise.

VII. CONCLUSION

The proposed PLL achieves a wide multiplication and output frequency range, satisfying the objective of using one clock generator PLL design without modification in a large set of ASICs. The PLL is self-biased with constant loop dynamics independent of multiplication factor, output frequency, process, voltage, and temperature. The sampled feedforward network suppresses pattern jitter with an effective third-order pole ω_C that tracks ω_{REF}. The PLL also achieves relatively constant period jitter of less than 1.7% of the output period as N is scaled from 1–4096.

REFERENCES

[1] J. Maneatis *et al.*, "Self-biased high-bandwidth low-jitter 1-to-4096 multiplier clock generator PLL," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2003, pp. 424–425.

[2] J. Maneatis, "Low-jitter process-independent DLL and PLL based on self-biased techniques," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1723–1732, Nov. 1996.

[3] A. Maxim *et al.*, "A low-jitter 125–1250 MHz process-independent 0.18 μm CMOS PLL based on a sample-reset loop filter," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, Feb. 2001, pp. 394–395.

[4] T. C. Lee and B. Razavi, "A stabilization technique for phase-locked frequency synthesizers," in *Symp. VLSI Circuits Dig. Tech. Papers*, June 2001, pp. 39–42.

[5] J. Maneatis and M. Horowitz, "Precise delay generation using coupled oscillators," *IEEE J. Solid-State Circuits*, vol. 28, pp. 1273–1282, Dec. 1993.

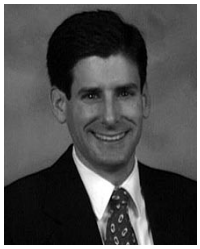

John G. Maneatis (S'90–M'92) received the B.S. degree in electrical engineering and computer science from the University of California, Berkeley, in 1988 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1989 and 1994, respectively.

While at Stanford University, from 1988 to 1994, his research interests included high-performance circuit design for phase-locked loops, microprocessors, data conversion, and clock recovery. From 1994 to 1998, he was a Lead Circuit Designer with Silicon Graphics, Inc., Mountain View, CA, working in the area of microprocessor design, clocking, and phase-locked loops. Currently, he is the President of True Circuits, Inc., Los Altos, CA, a company that designs and licenses a variety of circuit-level intellectual property (IP) and CAD software.

Jaeha Kim (S'94) received the B.S. degree in electrical engineering from Seoul National University, Seoul, Korea, in 1997 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1999 and 2003, respectively.

He was with True Circuits, Inc., Los Altos, CA, from 2001 to 2003, developing phase-locked loops and delay-locked loops for on-chip clock generation and off-chip interfaces. Currently, he is a Postdoctoral Researcher at Seoul National University, where his research interests include high-speed links, clock recovery PLL/DLLs, and low-power circuits and architectures.

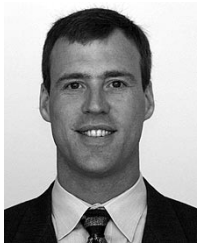

Iain McClatchie received the B.A. degree in computer science from the University of California, Berkeley, in 1992.

From 1991 to 1997, he was a Member of the Technical Staff with Silicon Graphics, Inc., Mountain View, CA, where he was responsible for floating-point verification and front-end CAD flows for two microprocessors. From 1997 to 2000, he pursued independent research on dynamic logic CAD flows. Since 2000, he has been with True Circuits, Inc., Los Altos, CA, working on phase-locked loops and high-speed logic.

Jay Maxey is a Distinguished Member of Technical Staff with Texas Instruments, Inc. (TI), Dallas, TX. He joined TI in 1978 and has been involved in IC design throughout his career, including low-power Schottky TTL, high-speed PLDs, FPGAs, and ASIC library development. He is currently responsible for TI's 130-nm ASIC PLLs in addition to driving circuit design methodologies for the ASIC Product Development group.

Manjusha Shankaradas received the B.S. degree in electrical engineering from S. R. M. Engineering College, Madras, India, in 1996 and the M.S. degree in electrical engineering from the University of South Florida, Tampa, in 1999.

She is currently with the ASIC Product Development Department, Texas Instruments Semiconductor Group, Dallas, TX. Her research includes design verification of phase-locked loops and test development on ATE using test chips.