# Low-Power Dividerless Frequency Synthesis Using Aperture Phase Detection

Arvin R. Shahani, Derek K. Shaeffer, *Student Member, IEEE*, S. S. Mohan, *Student Member, IEEE*, Hirad Samavati, *Student Member, IEEE*, Hamid R. Rategh, *Student Member, IEEE*, Maria del Mar Hershenson, *Student Member, IEEE*, Min Xu, *Student Member, IEEE*, C. Patrick Yue, *Student Member, IEEE*, Daniel J. Eddleman, *Student Member, IEEE*, Mark A. Horowitz, *Senior Member, IEEE*, and Thomas H. Lee, *Member, IEEE*

***Abstract*—A phase-locked-loop (PLL)-based frequency synthesizer incorporating a phase detector that operates on a windowing technique eliminates the need for a frequency divider. This new loop architecture is applied to generate the 1.573-GHz local oscillator (LO) for a Global Positioning System receiver. The LO circuits in the locked mode consume only 36 mW of the total 115-mW receiver power, as a result of the power saved by eliminating the divider. The PLL's loop bandwidth is measured to be 6 MHz, with a reference spurious level of −47 dBc. The front-end receiver, including the synthesizer, is fabricated in a 0.5-$\mu$m, triple-metal, single-poly CMOS process and operates on a 2.5-V supply.**

***Index Terms*—Frequency synthesizers, Global Positioning System, phase detection, phase-locked loops, radio-frequency integrated circuits, radio receivers.**

## I. Introduction

THE growing demand for portable, low-cost wireless-communication devices has spurred interest in radio-frequency integrated circuits. Part of offering a completely integrated solution involves identifying a low-power, monolithic gigahertz local oscillator (LO) implementation. A quartz-crystal-based oscillator cannot be used directly for the LO, since the fundamental modes of inexpensive quartz crystals are limited to approximately 30 MHz [1], and overtone orders of 50 are impractical. However, a crystal oscillator can be used as the reference in a static-modulus phase-locked-loop (PLL) frequency synthesizer. As is well known, the stability of the frequency-multiplied reference is retained by a wideband loop. This ability to synthesize a stable high-frequency source is beneficial, but it comes at the expense of significant power consumption. This paper addresses the power issue by introducing a new type of phase detector capable of phaselocking the synthesizer's frequency-multiplied output to its reference input, without the use of a divider. Eliminating the need for the divider allows the synthesis of a 1.573-GHz output on only 36 mW of power in this technology.

Section II examines the PLL-based LO used for the Global Positioning System (GPS) receiver architecture shown in Fig. 1 [2] and introduces the element that eliminates the need for a divider: the aperture phase detector (APD). Treatment begins at the architectural level and descends into the APD's detailed nature. Both the theory and the implementation of an APD are covered. Section III presents experimental results on the APD PLL.

Manuscript received May 7, 1998; revised August 4, 1998.

The authors are with the Center for Integrated Systems, Stanford University, Stanford, CA 94305 USA.

Publisher Item Identifier S 0018-9200(98)09432-3.

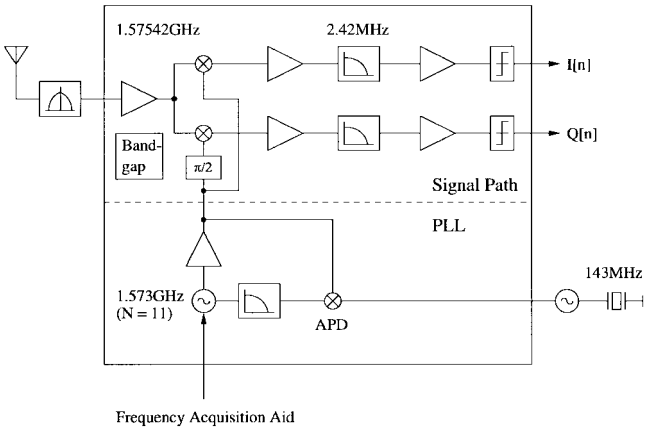


Fig. 1. GPS receiver architecture.

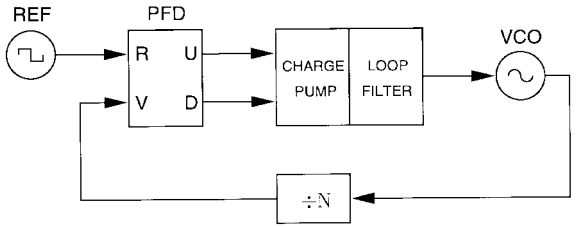


Fig. 2. Integer-$N$ synthesizer block diagram.

## II. PLL

### A. Architecture

The conventional and widely used implementation of the PLL frequency synthesizer with static modulus is the integer-$N$ synthesizer [3]. The traditional divide-by-$N$ block shown in Fig. 2 can be realized with a single counter. However, there are two drawbacks associated with the divider: power consumption and switching noise. Power consumption is large, particularly at high frequencies, because of the well-known $CV^2f$ relationship. For example, a recently published 1.6-

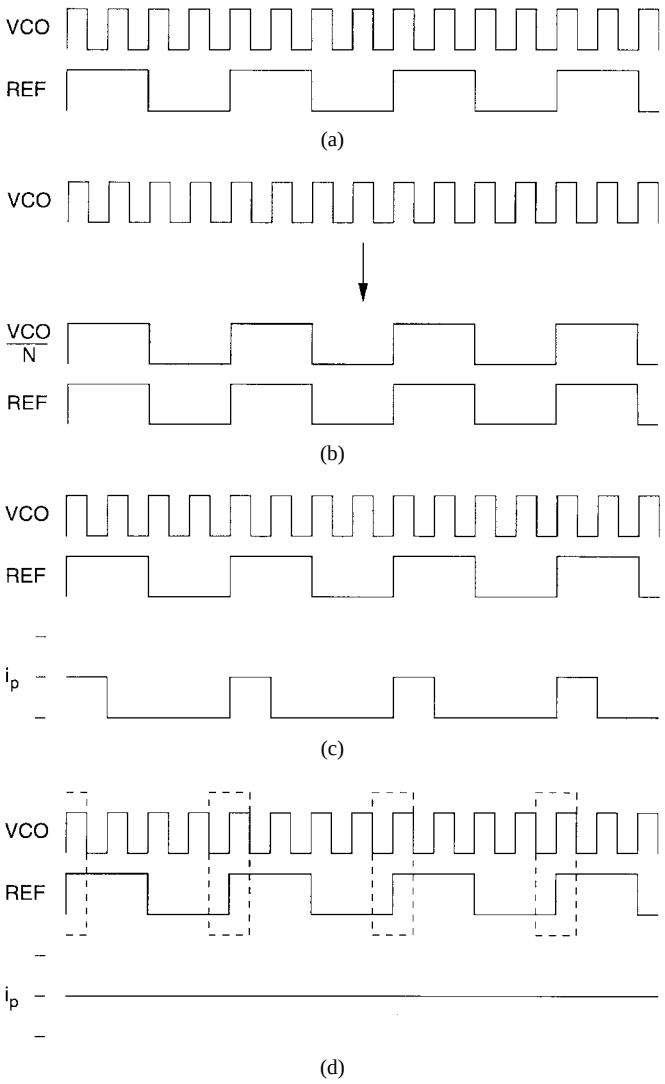


Fig. 3. Phaselock techniques. (a) Phaselocked signals. (b) Phaselock with a divider and PFD. (c) PFD along; negative charge pump current commands the VCO to decrease its frequency, breaking phaselock. (d) Phaselock with an APD.

GHz integer-$N$ synthesizer built in a 0.6-$\mu$m CMOS technology reported a total power consumption of 90 mW, of which 22.5 mW were used by the divider [4]. A further disadvantage of the divider is the on-chip interference generated by its high-speed digital transitions. This is particularly worrisome if the synthesizer is to be integrated with the front end's sensitive low-noise amplifier.

To reduce power consumption and high-frequency noise, a windowing technique that eliminates the divide-by-$N$ block for phase comparisons is investigated here. To appreciate how windowing may be of benefit, it is worthwhile to revisit the phenomenon of locking in a conventional PLL. To retain phaselock, it is necessary to align every $N$th rising or falling edge of the voltage controlled oscillator (VCO) with a corresponding reference edge. Phaselock is demonstrated in Fig. 3(a) for $N = 4$, where every fourth rising VCO edge lines up with a rising reference edge. A divider with a phase-frequency detector (PFD) accomplishes edge alignment by first dividing down the VCO by the right multiple so that edge

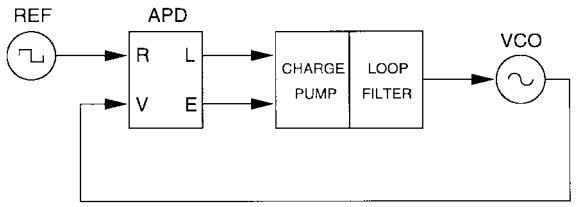


Fig. 4. APD synthesizer block diagram.

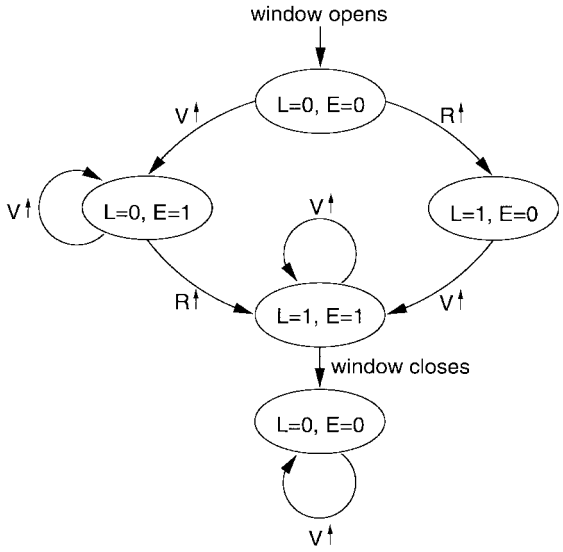


Fig. 5. Idealized APD state diagram.

alignment is unambiguous, as pictured in Fig. 3(b). Because the PFD compares phase over the entire reference cycle, a PFD cannot phaselock two inputs at different frequencies. In fact, it is precisely this property that makes the PFD popular.

Now consider using a PFD without a divider. Clearly, there would be an edge ambiguity problem, rendering the PFD quite ineffective, as seen in Fig. 3(c). The reason is that the PFD responds to every edge of the VCO, evidenced by the charge pump current's net negative value. This erroneously commands the VCO to decrease its frequency. However, by restricting the time interval during which phase is examined, one may eliminate the edge ambiguity, and hence the frequency divider. The dashed boxes in Fig. 3(d) define the window during which phase may be compared, even if the two inputs are of different frequency. The window can be controlled by the reference time base, since it periodically opens at that rate. Furthermore, the window need only be wide enough so that a VCO edge falls within it, which is equivalent to requiring that the window be active for a time longer than the instantaneous VCO period. No dividers are thus necessary to maintain phaselock, and this phase detector, called an APD, can operate with two inputs that are at different frequencies, as shown in Fig. 4.

A more substantive description of the APD's operation is provided in Fig. 5, which illustrates the state diagram for an idealized APD. When the window opens, the phase detector becomes active. The $R$-input rising edge sets the $L$ (denoting "late") terminal true, and the $V$-input rising edge sets the

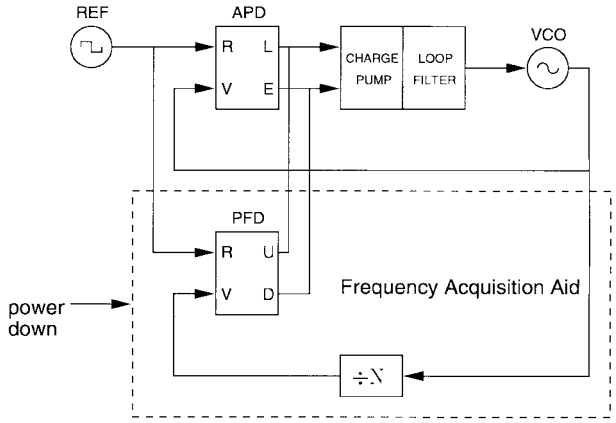


Fig. 6. APD synthesizer block diagram with integer-$N$ FAA.

$E$ (denoting "early") terminal true. Subsequent edges of the $V$-input are ignored until the next window opens. The time difference between the rising edges of the $L$ and $E$ signals is proportional to the phase error between the reference phase and the VCO phase. If $L$ is set first, the VCO phase is late; and conversely, if $E$ is set first, the VCO phase is early. When the window closes, the $L$ and $E$ terminals are reset (to false).

Fig. 1 shows that some type of frequency acquisition aid (FAA) is required to bring an APD-based loop initially into lock. This necessity is a consequence of restricting phase comparisons to a window, which eliminates the phase detector's ability to perform frequency detection. This issue is discussed in further detail in Section II-D. For this work, an external acquisition aid was used for experimental purposes. An integrated implementation of the acquisition aid, Fig. 6, uses the traditional divider with PFD to lock the loop and then powers down the acquisition aid, transferring control to the low-power APD. An APD can be used once in lock because the reference is derived from a stable crystal oscillator.

### *B. Loop Theory*

Having provided an overview of APD operation, we now develop a linearized APD PLL model relating input and output phase. This model is important for quantitative loop design and ensures that the synthesized output has the desired stability and noise performance.

From the description of the late and early APD signals given in the previous subsection, the average charge pump current over one reference cycle is given by

$$i_d = I_p(t_v - t_r)\frac{\omega_r}{2\pi} \tag{1}$$

where $I_p$ is the magnitude of the charge pump current, $t_v$ is the time of the first VCO rising edge in the window, $t_r$ is the time of the reference rising edge in the window, and $\omega_r$ is the angular reference frequency. The current $i_d$ can be expressed as a function of the reference and VCO phases by relating these phases to $t_r$ and $t_v$, assuming small phase errors. Expressions relating edge time to signal phase are

$$t_r = -\frac{\theta_r}{\omega_r} \tag{2}$$

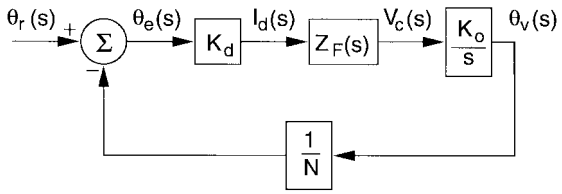


Fig. 7. APD PLL block diagram in lock.

where $\theta_r$ is the reference phase and

$$t_v = -\frac{\theta_v}{\omega_v} \tag{3}$$

where $\theta_v$ is the VCO phase and $\omega_v$ is the angular VCO frequency. The average charge pump current over one reference cycle can thus be written as

$$i_d = I_p(t_v - t_r)\frac{\omega_r}{2\pi} = \frac{I_p}{2\pi}\left(\theta_r - \theta_v\frac{\omega_r}{\omega_v}\right). \tag{4}$$

When the loop is in lock, $\omega_v = N\omega_r$, giving

$$i_d = \frac{I_p}{2\pi}\left(\theta_r - \frac{\theta_v}{N}\right) = \frac{I_p}{2\pi}\theta_e = K_d\theta_e \tag{5}$$

where $K_d$ is the phase-detector gain constant. Note that even though there is no explicit divider in the loop, the VCO phase is divided by $N$ in (5), just as in a conventional loop.

This model can be used in place of the APD in Fig. 4, and the other blocks in the same figure can be replaced by their corresponding linear time-invariant (LTI) models, yielding the overall system model shown in Fig. 7. Fig. 7 is an LTI representation of the APD PLL in lock, from which the phase transfer function is readily found to be

$$H(s) = \frac{\theta_v}{\theta_r} = \frac{NK_dK_oZ_F(s)}{Ns + K_dK_oZ_F(s)} \tag{6}$$

where $K_o$ is the VCO gain constant and $Z_F(s)$ is the loop filter's impedance, expressed in the $s$-domain.

### *C. APD Characteristic ($i_d$ Versus $\theta_e$)*

The derivation in the previous subsection treats the APD for small phase errors. For completeness, it is instructive to examine the response of the APD to arbitrary phase errors. Now, the delay between the time the window opens and the time at which the reference edge occurs becomes important. This delay is designated by $d$, which is a positive quantity whose least restrictive range is limited to $[0, T_r)$, where $T_r$ is the reference period. However, the loop can lock if and only if $d$ is in the interval $[0, T_v]$, where $T_v$ is the VCO period. Otherwise, the first VCO edge within the window will always precede the reference edge.

From Fig. 8, it is apparent that the characteristic will be periodic in VCO phase, because when the VCO waveform has moved one VCO period to the right, the situation is identical to the start. As the VCO waveform moves to the right, the time difference $t_v - t_r$ varies proportionally with phase error $\theta_e$. Therefore, to find the APD's characteristic, $\theta_e$ and $i_d$ need to be calculated at only two points, and the remainder of the

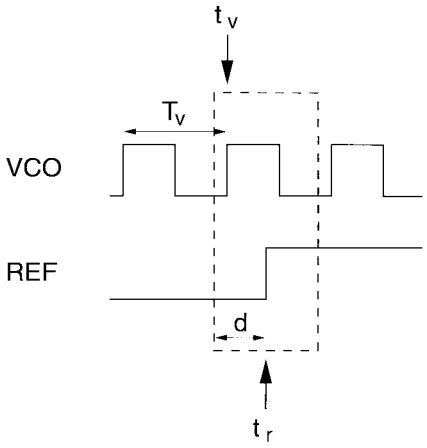


Fig. 8. Position of VCO and reference edge in window.

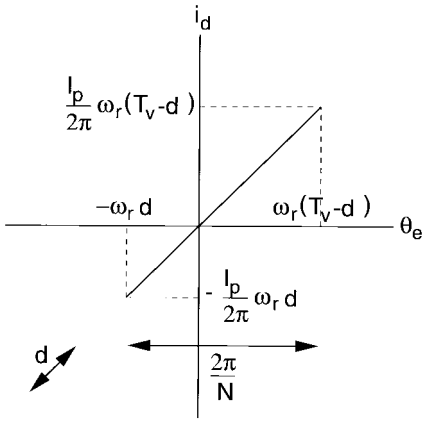


Fig. 9. APD characteristic over $(2\pi)/N$ interval.

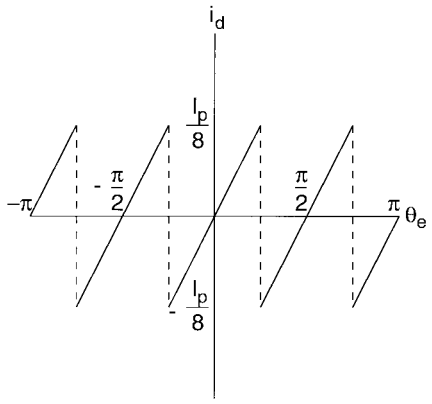


Fig. 10. APD characteristic for $d = (T_v)/2$ and $N = 4$.

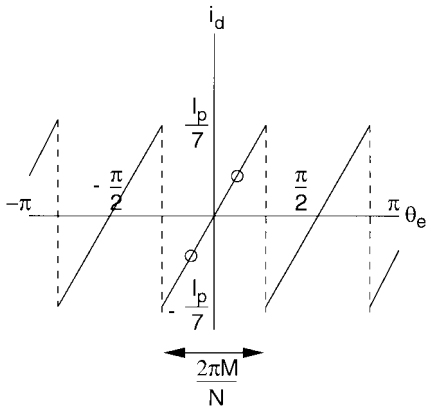


Fig. 11. $M = 2$, $N = 7$ subharmonic-lock mode.

characteristic is generated by connecting these endpoints. $\theta_e$ and $i_d$ are first calculated for $t_v - t_r = -d$

$$\theta_e = -\omega_r d \tag{7}$$

$$i_d = -\frac{I_p}{2\pi}\omega_r d. \tag{8}$$

Next, $\theta_e$ and $i_d$ are calculated at the other extreme where $t_v - t_r = T_v - d$

$$\theta_e = \omega_r(T_v - d) \tag{9}$$

$$i_d = \frac{I_p}{2\pi}\omega_r(T_v - d). \tag{10}$$

From this information, the portion of the APD characteristic shown in Fig. 9 can be constructed.

The influence of two parameters—$d$, the delay between the time the window opens and when the reference edge occurs, and $N$, the ratio between the VCO and reference frequencies—warrants special attention. Decreasing $d$ shifts the characteristic diagonally up (along the line of the characteristic), and increasing $d$ shifts the characteristic diagonally down. It is desirable to have $d$ equal to half the synthesized frequency's period. By designing for this condition, the APD characteristic will be centered about $i_d = 0$ to provide a symmetrical correction range. The parameter $N$ affects the phase error's periodicity, with larger values increasing the periodicity. Fig. 10 shows the complete APD characteristic (a $2\pi$ variation in $\theta_e$) for the specific case where $d = \frac{T_v}{2}$ and $N = 4$.

### *D. Subharmonic-Lock Modes*

The existence of *subharmonic*-lock modes explains the need for an acquisition aid. During each window, which opens periodically at the reference rate, the APD makes a single phase comparison. It is this property that allows an APD to phaselock the VCO's output to an integer multiple of the reference input. But the ability to examine the phase of two signals at different frequencies introduces more modes than just the desired integer-lock modes. Additional subharmonic-lock modes occur if the net current delivered over multiple cycles of the reference is zero, allowing the loop to stay locked at an undesired frequency [5].

If we designate by $M$ the number of reference cycles over which the net charge delivered to the loop filter is zero, then an expression relating the reference frequency to the VCO frequency when phaselock occurs is $M\omega_v = N\omega_r$. Fig. 11 displays the points on the APD characteristic between which the loop ping-pongs for the specific case where $M = 2$, and $N = 7$. Because $M = 2$, the charge pump alternates between pumping up on one cycle and pumping down on the next cycle, balancing the charge to the loop filter over two cycles.

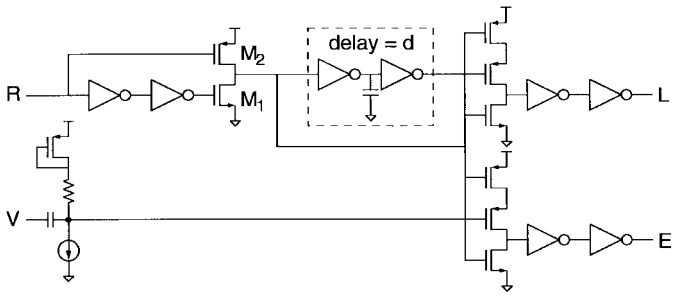


Fig. 12. APD circuit diagram.

These subharmonic-lock modes are problematic because they are spaced, in frequency, closer than the neighboring integer-lock modes. However, the APD favors integer over subharmonic modes for two reasons. First, the loop's bandwidth imposes a limit on $M$. If the number of reference cycles over which the charge pump current averages to zero grows too large, the loop will act on partial information because the loop responds to signals averaged over a loop period. The loop period is the reciprocal of the closed-loop bandwidth. Another reason the APD favors integer over subharmonic modes is that the subharmonic modes have a lower detector gain $K_d$ because the VCO edge arrives at a different time in each of the $M$ cycles. If the APD characteristic is nonlinear, then the overall detector gain is the average of the $M$ individual linearized detector gains.

Using an FAA to ensure frequency lock eliminates the concern of locking in a subharmonic mode. Once lock has been achieved, and control transferred to the APD, the APD is capable of maintaining lock at the desired frequency.

### E. APD Circuit Implementation

Fig. 12 shows a circuit implementation of an APD. The reference clock (which has about a 50% duty cycle) is shaped by the structure preceding the delay to have fast falling edges, since these are the edges that enable the precharged gates. When the reference input is low, $M_1$ is off and $M_2$ is on, causing the output to be high. After the reference rises, $M_2$ shuts off before $M_1$ turns on due to the two inverter delays. Therefore, $M_1$ does not fight $M_2$ to pull the output low, and creates a fast falling edge. The window opens on this fast falling edge. The delay $d$ between the opening of the window and the reference edge is determined by two inverters with a capacitor in the middle. The APD uses two precharged gates to evaluate the reference and VCO phases. An advantage of precharged gates is that they only respond once while active. In this case, the precharged gates are precharged low, and rise on detection of low levels. Also, the precharge action does not affect the loop to first order, because the state $L = 1$, $E = 1$ has the same action as the state $L = 0$, $E = 0$.

The behavior of this APD circuit differs somewhat from the ideal APD discussed in Section II-A. In particular, the circuit implementation responds to falling edges instead of rising edges, and more precisely, the precharged gates act as level detectors of a low voltage level instead of as edge detectors.

A simulation of the APD characteristic over a $\frac{2\pi}{N}$ interval is shown in Fig. 13, where $N = 11$. The phase error $\theta_e$ is plotted against the average charge pump current over one reference cycle $i_d$, and includes the nonidealities of the charge pump as well. The flat section near $-0.2$ rad is where the $E$ signal driving the charge pump is compressing due to the level detection nature of the precharged gates. Another imperfection in this circuit's APD characteristic is the section with finite negative slope instead of a discontinuity. From the characteristic, the phase detector's gain constant $K_d$ in a state of zero static phase error $(\theta_e = 0)$ is evaluated to be 7.4 $\mu$A/rad.

### F. PLL Circuit Implementation

In Section II-B, a general model for a locked APD PLL was developed, expressing the closed-loop phase transfer function in terms of the loop filter's $s$-domain impedance and an idealized VCO. We now provide specific expressions for the loop as actually implemented.

The loop filter used is the conventional network shown in Fig. 14, whose $s$-domain impedance is

$$Z_F(s) = \frac{1 + sRC_1}{(C_1 + C_2)s\left(1 + sR\frac{C_1C_2}{C_1+C_2}\right)}. \tag{11}$$

A single-pole amplifier was used to interface with the VCO's varactor, thus the ideal VCO transfer function $(K_o)/s$ must be modified to

$$\frac{K_o}{s\left(1 + \frac{s}{2\pi f_{3\,\mathrm{dB}}}\right)} \tag{12}$$

where $f_{3\,\mathrm{dB}}$ is the 3-dB bandwidth of the VCO's preamplifier. Using (11) and (12) in (6) enables us to write the complete phase transfer function of the implemented APD PLL as shown in (13) at the bottom of the page. In the next section, we compare measured data to (13).

## III. Experimental Results

A test chip (see Fig. 15) containing a copy of the APD PLL used in the complete GPS receiver is used to evaluate the APD PLL. Two separate tests are performed; one to verify the derived closed-loop transfer function of the APD PLL and the other to observe the synthesized LO spectrum for the GPS receiver. In the second test, the synthesized LO is also checked

$$H(s) = \frac{NK_dK_o(1 + sRC_1)}{N(C_1 + C_2)s^2\left(1 + \frac{s}{2\pi f_{3\,\mathrm{dB}}}\right)\left(1 + sR\frac{C_1C_2}{C_1+C_2}\right) + K_dK_o(1 + sRC_1)}. \tag{13}$$

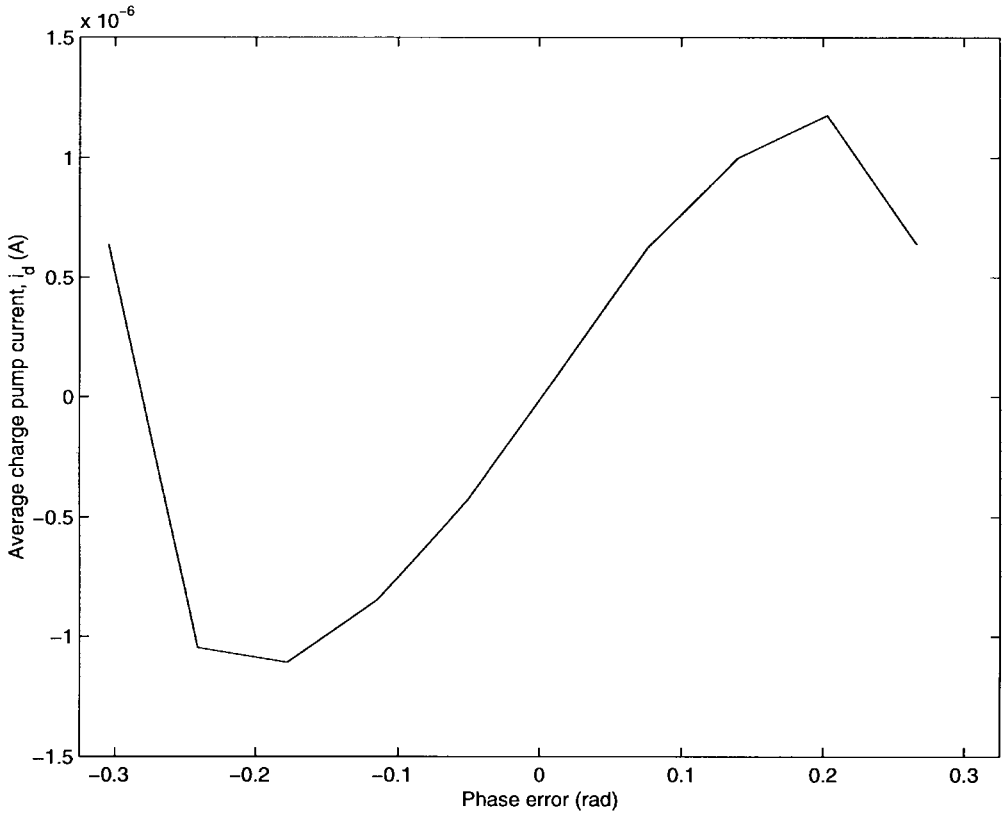


Fig. 13. Simulated APD circuit characteristic.

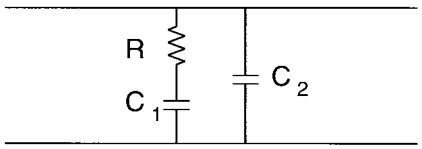


Fig. 14. Loop filter used in APD PLL.

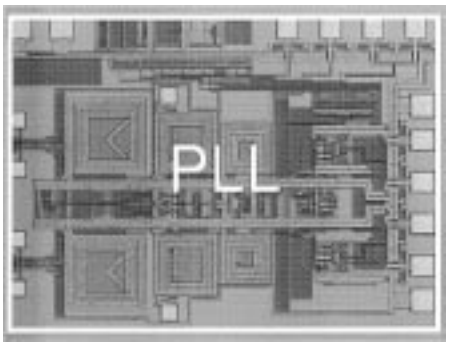


Fig. 15. Die photograph of PLL on GPS receiver test chip.

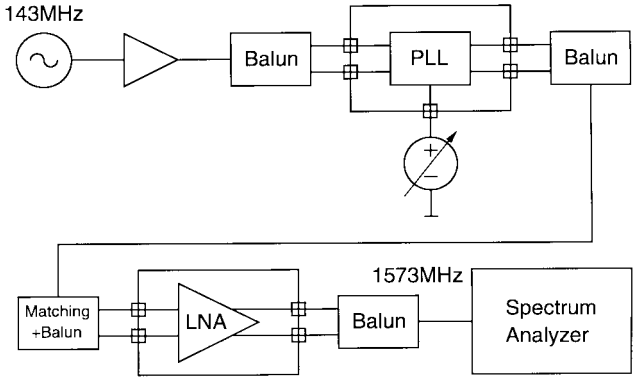


Fig. 16. PLL test setup number 1.

with a microwave frequency counter to verify its long-term stability.

Fig. 16 shows the experimental setup for the first test. Phase noise is measured for offsets from 1 kHz to 10 MHz with the HP8563E spectrum analyzer, which has special phase-noise-measurement software. Ten MHz is used as the upper limit since the loop is designed to have a bandwidth less than 10 MHz. Beyond the loop's bandwidth, the PLL's phase noise is determined by the VCO's phase noise, making measurement of the PLL's transfer function difficult. One of the largest factors affecting measurement accuracy is the noise floor of the instrument. To minimize this error source, measurements of the floor with a clean source are performed first. These results are later used to calibrate the data. Reference phase noise and PLL phase noise are also both measured. After some data processing, the PLL's closed-loop phase transfer function $H(s)$ is determined.

Fig. 17 shows the measured $|H(f)|$ and the predicted $|H(f)|$, from (13), for the case where the reference frequency is 143 MHz and the VCO frequency is 1.573 GHz $(N = 11)$. The seven loop parameters in (13) are set as follows: $N$ is known; $K_o$ is taken from measured VCO data; $R$, $C_1$, and

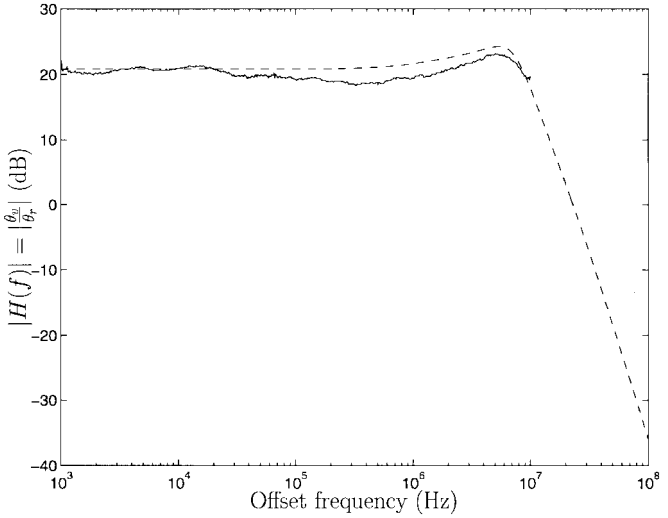


Fig. 17. Measured and predicted $|H(f)|$.

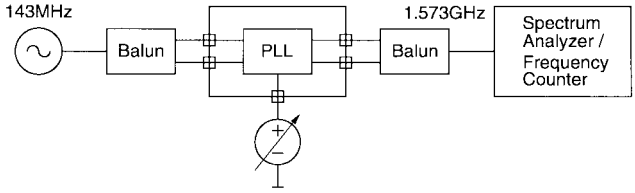


Fig. 18. PLL test setup number 2.

$C_2$ are taken to be their designed loop filter values; $f_{3\,\mathrm{dB}}$ is calculated from the technology data; and $K_d$ is fit. The fit value of $K_d$, 6.6 $\mu$A/rad, is a little less than the simulated value noted in Section II-E, 7.4 $\mu$A/rad. Since $K_d = (I_p)/(2\pi)$ for an ideal APD, one could argue that the discrepancy in $K_d$ is due to an actual pump current that is lower than the pump current used in simulations. But when $I_p$ is measured, it is found to be correct. Still, the discrepancy in $K_d$ is readily explained. The simulation in Fig. 13 establishes an upper bound on $K_d$ because it is measured in a state of zero static phase error. The simulated APD circuit characteristic illustrates that the detector gain (i.e., the slope) decreases the farther that one departs from zero radians. The charge pump is known to have some offset; thus, the loop has some static phase error in lock to overcome the offset, resulting in a slightly lower $K_d$.

Fig. 18 shows the experimental setup for the second test. The LO spectrum is measured with the HP8563E spectrum analyzer, and the frequency is checked with an HP5350B microwave frequency counter. Fig. 19 displays the synthesized output spectrum, in which the PLL's ability to track the low close-in phase noise of the reference can be seen. The visible skirts are due to the VCO's phase noise outside the 6-MHz bandwidth of the PLL. Spurious tones at $-47$ dBc are primarily due to control-line ripple resulting from charge pump leakage. In GPS applications, the measured spurious level is acceptable because of the absence of blockers at the corresponding offset frequencies. In more demanding applications, one may reduce ripple through improved charge pump design and the use of analog phase interpolation [3].

Table I provides a summary of the APD PLL's performance.

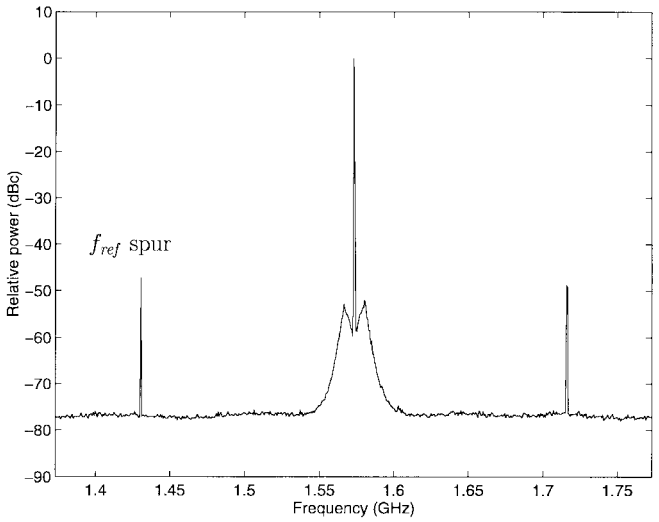


Fig. 19. LO spectrum.

TABLE I
MEASURED APD PLL PERFORMANCE

| | |
|---|---|
| Synthesized frequency | 1.573GHz |
| Reference frequency | 143MHz |
| Loop bandwidth | 6MHz |
| $f_{ref}$ spur | $\leq$ -45dBc |
| $2 * f_{ref}$ spur | $\leq$ -55dBc |
| VCO power consumption | 26mW |
| Total power consumption | 36mW (2.5V supply) |
| Die area | 3.1mm$^2$ |
| Technology | 0.5$\mu$m CMOS |

The PLL has a wide bandwidth of 6 MHz, and the APD circuit consumes only one-quarter of the total synthesizer power. With the elimination of the divider, the main power consumer in the synthesizer is now the VCO.

## IV. CONCLUSION

A new method for performing phase detection that eliminates the divide-by-$N$ function within a PLL has been presented. A frequency acquisition aid circuit, which can be powered down once lock is established, is required. By using an aperture phase detector, a 1.573-GHz local oscillator can be synthesized on roughly half the power of a loop containing a conventional divider. Additionally, elimination of the divider also reduces the frequency of transitions that might cause substrate and supply bounce. The power savings and noise reduction make the APD PLL an attractive design for low-power, integrated frequency synthesizers.

## ACKNOWLEDGMENT

The authors gratefully acknowledge Rockwell International for fabricating the receiver and Dr. C. Hull and Dr. P. Singh for their valuable assistance. In addition, the authors acknowledge Tektronix, Inc., for supplying simulation tools and E. McReynolds for his invaluable support of, and assistance with, CMOS modeling issues. Last, the authors thank IBM for generous student support through IBM fellowships.

## REFERENCES

[1] *M-tron Engineering Notes*, Dec. 1997.
[2] D. K. Shaeffer, A. R. Shahani, S. S. Mohan, H. Samavati, H. R. Rategh, M. Hershenson, M. Xu, C. P. Yue, D. J. Eddleman, and T. H. Lee, "A 115-mW, 0.5-$\mu$m CMOS GPS receiver with wide dynamic-range active filters," *IEEE J. Solid-State Circuits*, vol. 33, pp. 2219–2231, Dec. 1998.
[3] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge: Cambridge Univ. Press, 1998.
[4] J. F. Parker and D. Ray, "A 1.6 GHz CMOS PLL with on-chip loop filter," *IEEE J. Solid-State Circuits*, vol. 33, pp. 337–343, Mar. 1988.
[5] F. M. Gardner, *Phaselock Techniques*, 2nd ed. New York: Wiley, 1979.

**Arvin R. Shahani,** for a photograph and biography, see this issue, p. 2041.

**Derek K. Shaeffer** (S'98), for a photograph and biography, see this issue, p. 2230.

**S. S. Mohan** (S'98), for a photograph and biography, see this issue, p. 2231.

**Hirad Samavati** (S'98), for a photograph and biography, see this issue, p. 2041.

**Hamid R. Rategh** (S'98), for a photograph and biography, see this issue, p. 2231.

**Maria del Mar Hershenson** (S'98), for a photograph and biography, see this issue, p. 2231.

**Min Xu** (S'97), for a photograph and biography, see this issue, p. 2231.

**C. Patrick Yue** (S'93), for a photograph and biography, see this issue, p. 2231.

**Daniel J. Eddleman** (S'98), for a photograph and biography, see this issue, p. 2231.

**Mark A. Horowitz** (S'77–M'78–SM'95) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1978 and the Ph.D. degree from Stanford University, Stanford, CA, in 1984.

He is the Yahoo Founders Professor of Electrical Engineering and Computer Science at Stanford. His research area is in digital system design. He has led a number of processor designs including MIPS-X, one of the first processors to include an on-chip instruction cache; TORCH, a statistically scheduled, superscalar processor; and FLASH, a flexible DSM machine. He has also worked on a number of other chip design areas, including high-speed memory design, high-bandwidth interfaces, and fast floating point. In 1990, he took a leave from Stanford to help start Rambus, Inc., a company designing high-bandwidth memory interface technology. His current research includes multiprocessor design, low-power circuits, memory design, and high-speed links.

Dr. Horowitz received a 1985 Presidential Young Investigator Award and an IBM Faculty Development Award, as well as the 1993 Best Paper Award from the International Solid-State Circuits Conference.

**Thomas H. Lee** (S'87–M'87), for a photograph and biography, see this issue, p. 2041.