AN-1061
APPLICATION NOTE

One Technology Way • P.O. Box 9106 • Norwood, MA 02062-9106, U.S.A. • Tel: 781.329.4700 • Fax: 781.461.3113 • www.analog.com

Behavior of the AD9548 Phase and Frequency Lock Detectors in the Presence of Random Jitter

by Ken Gentile

INTRODUCTION

The AD9548 is a digital PLL with a direct digital synthesizer (DDS) serving the role of the VCO appearing in an analog PLL. The digital nature of the AD9548 enabled the designers to implement digital phase lock and frequency lock detectors, as well (see Figure 1).

The sole purpose of the detectors is to indicate to the user whether the PLL control loop has reached a state that signifies a lock condition. As such, the detectors play no role in the phase and frequency acquisition process of the PLL, but serve only as status indicators.

The AD9548 digital phase/frequency detector logic (DPFDL) has two output signals. One comprises time error samples, which constitute the time difference between the reference and feedback edges. The other comprises period error samples, which constitute the difference between the periods of the reference and feedback signals.

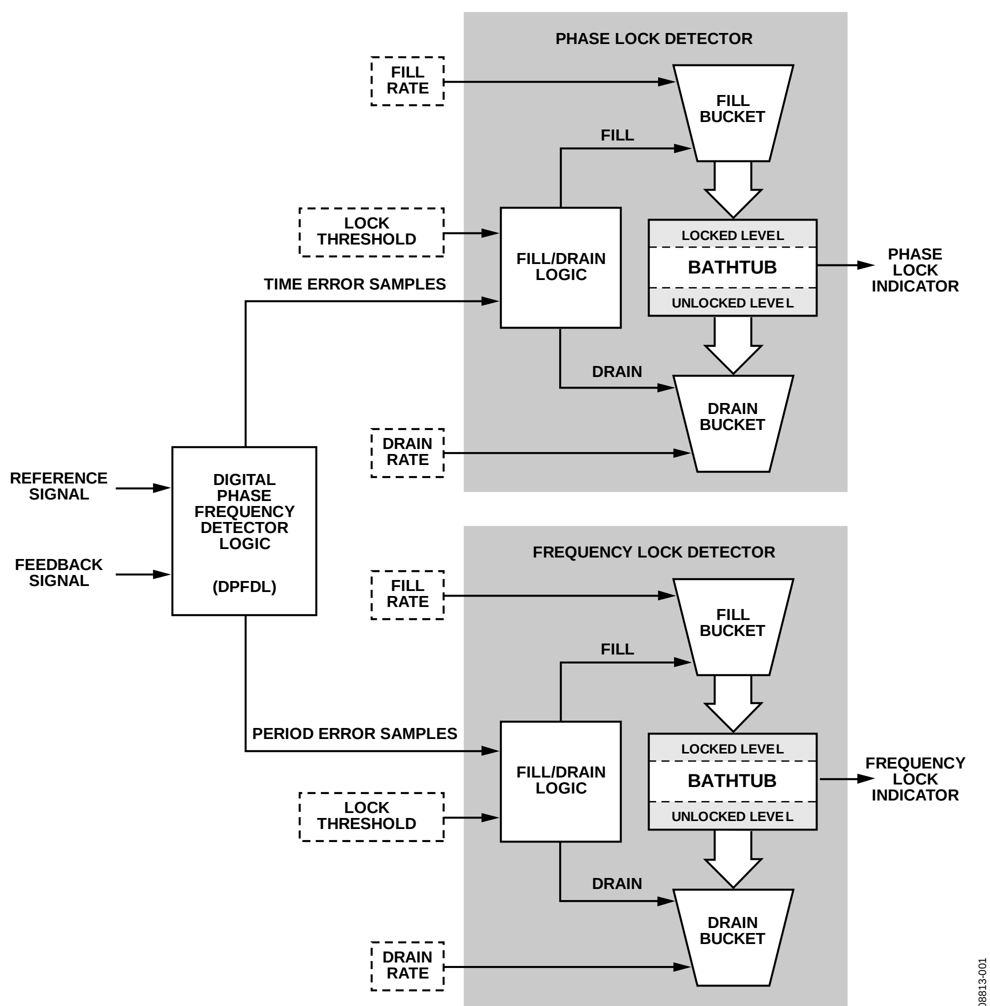

Figure 1. The AD9548 Phase and Frequency Lock Detectors

TABLE OF CONTENTS

PURPOSE OF THE APPLICATION NOTE

The fact that the AD9548 detectors are digital provides the option of programmability, thereby making them extremely flexible compared to their analog counterparts.

Taking advantage of this flexibility, however, requires an understanding of the relationship between the programming parameters (lock threshold, fill rate, and drain rate) and the jitter characteristics of the input signal, which is the purpose of this application note.

THE LOCK DETECTOR BATHTUB ANALOGY

As described in the AD9548 data sheet, the detectors behave like a bathtub that has water either added with a fill bucket or removed with a drain bucket (see Figure 1). The tub can hold 4096 gallons of water and has a level mark at the ¼ and ¾ levels (−1024 and +1024 gallons, respectively, with 0 denoting ½ full).

Whenever the water level in the tub reaches or exceeds ¾ full, the lock detector indicates locked. Conversely, whenever the water level in the tub drops to or below ¼ full, the lock detector indicates unlocked.

During those times that the water level in the tub is between the ¼ and ¾ marks, the lock detector retains its previous indication (either locked or unlocked, as the case may be). The ¼ and ¾ marks provide the detectors with hysteresis so that their lock/unlock output signal is less likely to chatter when the level in the tub is near the ¼ or ¾ mark.

At the start of the PLL acquisition process, the tub starts out half full and the lock detector indicates unlocked. The fill bucket adds water whenever a sample (time error or period error) is within the lock threshold value. The drain bucket removes water whenever a sample is outside the lock threshold value.

The lock threshold value is programmable, as is the size of the fill and drain bucket, via the fill rate and drain rate parameter, respectively. The bucket size is programmable in 1-gallon increments from 1 to 255 gallons. Note that the phase lock detector and frequency lock detector are identical, but have independent lock threshold, fill rate, and drain rate parameters.

LOCK THRESHOLD DETAIL

The lock threshold value is the key to controlling the lock detectors as it sets the decision point for adding or removing water from the tub. The lock detector continuously tests every sample generated by the DPFDL to see if it is within or outside of the lock threshold value. If the sample is within the threshold, one fill bucket of water adds to the tub. If the sample is outside the threshold, one drain bucket of water subtracts from the tub.

The phase lock detector uses a 16-bit number to establish the phase lock threshold in units of picoseconds (ps). Suppose, for example, the input frequency to the DPFDL is 50 kHz and the expectation is that the phase lock detector should indicate lock when the DPFDL inputs are within 1° of alignment.

This works out to a time difference of

$$(1°/360°)/(50 \times 10^3) = 55{,}556 \text{ ps}$$

Therefore,

$$\text{phase lock threshold} = 55{,}556 \text{ ps}$$

The frequency lock detector uses a 24-bit number to establish the frequency lock threshold in units of picoseconds (ps).

Again, suppose the input frequency to the DPFDL is 50 kHz and the expectation is that the frequency lock detector should indicate lock when the DPFDL input frequencies are within 10 Hz of each other. This works out to a time difference of

$$1/50{,}000 - 1/(50{,}000 + 10) = 3999 \text{ ps}$$

Therefore,

$$\text{frequency lock threshold} = 3999 \text{ ps}$$

FILL AND DRAIN RATE DETAIL

The fill rate and drain rate parameters control the responsiveness of the lock detectors, which affects how quickly the detector swings between locked and unlocked indications. Recall that the fill and drain buckets can be programmed to any value from 1 to 255 gallons. Because of the separation between the locked (¾) and unlocked (¼) levels, there is a 2048-gallon difference between a locked indication and an unlocked indication. Thus, the programmed bucket size sets the minimum number of buckets required to traverse the lock/unlock span. For example, a 255-gallon bucket takes at least 9 buckets to cover the 2048-gallon span, whereas a 1-gallon bucket takes at least 2048 buckets to cover the span (assuming that every DPFDL sample results in the same decision with regard to the lock threshold, either always filling or always draining).

At the beginning of the lock acquisition process, however, the tub always starts out ½ full, so it only takes 1024 gallons to reach the locked or unlocked level from a cold start. Therefore, from a cold start, a 255-gallon bucket takes at least 5 buckets to cover the span, whereas a 1-gallon bucket takes at least 1024 buckets to cover the span (again, assuming every DPFDL sample results in the same decision).

On the other hand, if the detector manages to saturate (either a fully locked condition at +2048 gallons or fully unlocked condition at −2048 gallons), then there is a 3072 gallon difference between saturation and the alternate indication. Thus, a 255-gallon bucket takes at least 13 buckets to cover the 3072-gallon span, whereas a 1-gallon bucket takes at least 3072 buckets to cover the span (again, assuming every DPFDL sample results in the same decision).

Clearly, the programmed bucket size has a significant impact on the number of DPFDL samples required to reach the locked or unlocked levels. A small bucket causes the detector to be sluggish, while a large bucket causes the detector to be quite responsive. The AD9548 detectors are unique in the fact that they allow the user independent control of the detector's responsiveness to indicating both lock and unlock.

Keep in mind that the preceding paragraphs assume that every DPFDL sample yields the same decision result. In reality, the input signal will exhibit jitter due to noise in the system, which means that when the DPFDL output signal is near the lock threshold the fill/drain decisions could be noisy. This is why the jitter characteristics of the input signal can have a significant effect on the choice of the lock threshold, fill rate, and drain rate, as explained in this application note.

GENERIC PLL LOCK ACQUISITION PROCESS

Figure 2 shows an example of a generic phase error vs. time plot for a typical PLL acquiring phase lock (the blue trace in the top part of the figure). The beginning of the trace (left to right) shows cycle slips. These can occur when the difference between the reference and feedback frequencies is relatively large.

In this case, as the PLL gradually drives the feedback frequency toward the reference frequency, the phase difference between the two crosses the ±π limits of the phase detector causing it to jump from +π to −π (or vice versa). When the two frequencies are close enough, however, the PLL is able to hold the phase difference within the ±π limits of the phase detector. This begins the linear response region of the PLL where it gradually forces the feedback and reference signals into phase alignment.

Depending on the closed loop dynamics of the PLL, the locking process may exhibit ringing as shown by the multiple excursions of the blue trace through zero in the upper portion of Figure 2. Eventually, the PLL loop is able to drive the phase error to near zero, constituting complete phase lock. The area of interest with regard to this application note is in the region of the zoom box.

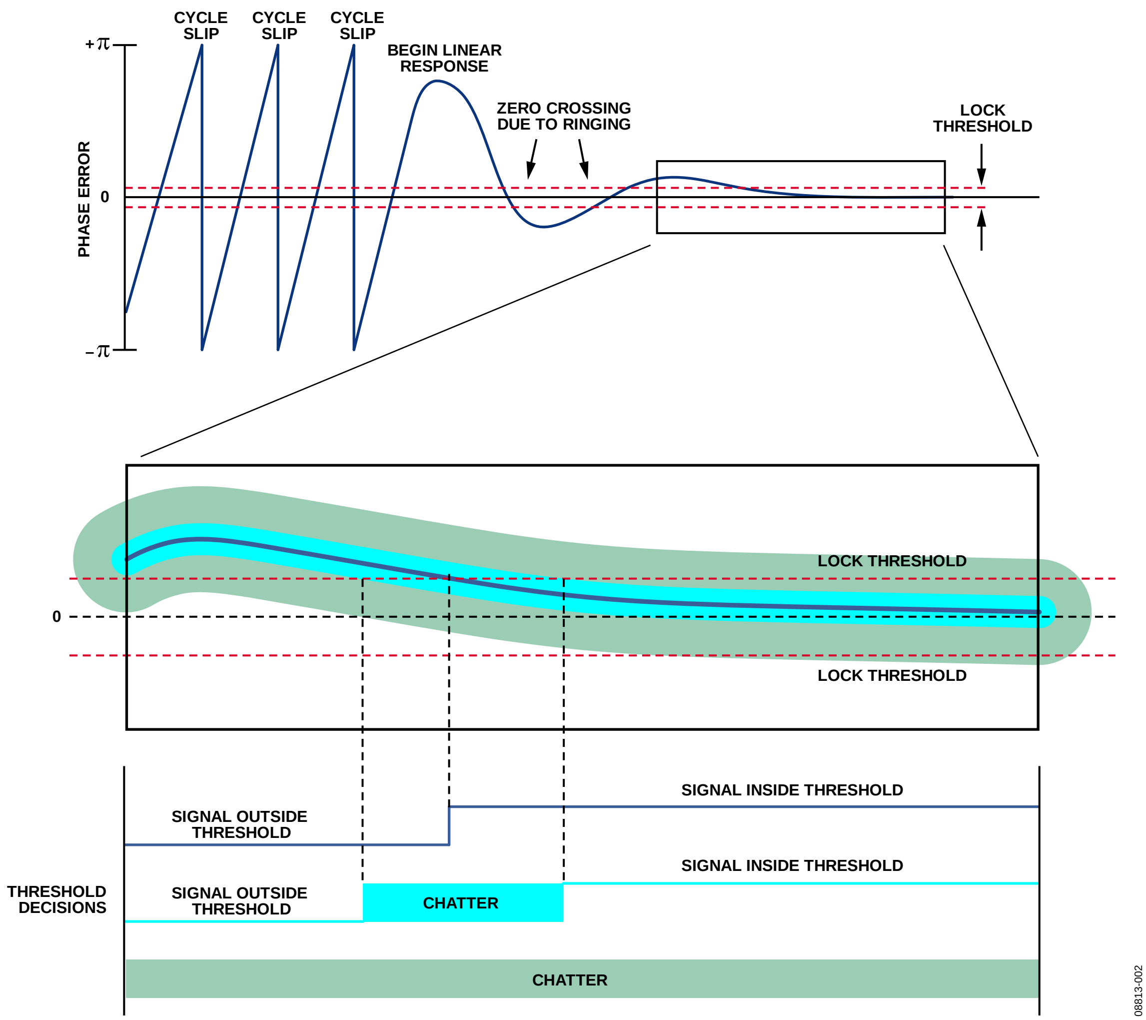

Figure 2. The Typical Phase Lock Acquisition Process

A lock detector must be able to determine the difference between a locked and an unlocked condition. This implies a phase (frequency) threshold that allows the detector to discriminate between the two conditions. That is, when the phase (frequency) error is within the threshold, the detector should indicate that the PLL is in a locked condition. Conversely, while the phase (frequency) error is outside the threshold, the detector should indicate that the PLL is in an unlocked condition. The threshold appears as the two red horizontal dashed lines in Figure 2.

The zoom box in Figure 2 shows a solid blue trace surrounded by a broad aqua trace and an even broader light green trace. The blue trace represents the response of an ideal noiseless system, while the aqua and green traces represent jitter (noise) superimposed on the ideal blue trace. Below the zoom box are traces showing the results of the detector's threshold decisions. Note the unambiguous decision associated with the noiseless (blue) trace.

The aqua trace, however, exhibits random jumps between decisions (chatter) due to noise on the input signal as the signal crosses through the threshold level. Worse still is the green trace, in which the peak value of the noise is greater than the threshold level. The result is persistent random excursions through the threshold level causing continuous decision chatter. Clearly, knowing the jitter characteristics of the input signal is a prerequisite to choosing the appropriate lock threshold value as explained in the Effect of Input Jitter on Choosing the Lock Threshold and Fill/Drain Rates section.

EFFECT OF INPUT JITTER ON CHOOSING THE LOCK THRESHOLD AND FILL/DRAIN RATES

LOCK THRESHOLD VS. JITTER

Recall that in the sections describing the Lock Threshold Detail and Fill and Drain Rate Detail, the lock threshold appears to be completely independent of the fill rate and drain rate. However, it is only possible to treat them independently if the peak input jitter is less than the lock threshold value. Otherwise, random excursions of the DPFDL signal outside of the threshold limits causes undesired out-of-threshold decisions to occur even after the PLL control loop settles to equilibrium.

Each undesired out-of-threshold decision, in turn, causes the lock detector to drain the tub. A high rate of undesired out-of-threshold decisions could cause the level in the tub to drop to the unlock level even though the PLL is tracking the input signal as it should. The unwanted result is that the lock detector will signal an unlock condition even though the PLL loop is in equilibrium. It is the undesired out-of-threshold decisions resulting from jitter exceeding the lock threshold that creates the interdependence between the lock threshold and the fill and drain rate.

The crux of the problem is that the user sets the lock threshold based on some system requirement. For example, the requirement may be for the PLL to indicate phase lock when the reference and feedback signals are within 5°. This means setting the lock threshold to a value corresponding to 5°. However, if the peak input jitter is 10°, for example, then there are bound to be a significant number of samples outside of the 5° threshold setting and each one causing an out-of-threshold decision. Preventing the lock detector from producing false unlock indications in spite of the undesired out-of-threshold decisions puts constraints on the choice of the lock threshold, fill rate and drain rate values.

Generally, one sets the lock threshold as described in the Lock Threshold Detail section based on some underlying system requirement. However, there are two scenarios to consider regarding the choice of lock threshold, which depends on the magnitude of the rms (σ) jitter relative to the lock threshold level.

Scenario 1: lock threshold > 6 σ

Scenario 2: lock threshold < 6 σ

In the first scenario, the lock threshold is wide enough so that the peak excursions of the jitter rarely exceed the threshold. In the second scenario, the lock threshold is so narrow that the peak excursions of the jitter frequently exceed the threshold. Because the value of σ is crucial to setting a viable lock threshold, it is worthwhile to review the Appendix: Normal (Gaussian) Distribution in the Context of Jitter section to learn how σ relates to the lock threshold.

FILL AND DRAIN RATE VS. JITTER

The standard deviation (σ) of the input jitter can also have an impact on the choice of the fill rate and drain rate. Usually, one selects the fill rate and drain rate based on the desired responsiveness of the detector. That is, large values make the detector very responsive, while small values make the detector sluggish. The fill rate controls the responsiveness of the detector for indicating a lock condition, while the drain rate controls the responsiveness for indicating an unlock condition. The fill/drain ratio, η, is an indication of the tendency of the detector to favor lock or unlock indications, where:

$$\eta = \textit{fill rate}/\textit{drain rate}$$

For η = 1, the detector is equally responsive to indicating lock or unlock conditions. For η > 1, the detector is more responsive to indicating a lock condition, whereas for η < 1, it is more responsive to indicating an unlock condition.

When the lock threshold is greater than 6 σ, jitter is not an issue (at least with regard to the lock detector) and one is free to choose the fill rate and drain rate solely on the desired responsiveness of the detector. When the lock threshold is less than 6 σ, however, then jitter causes random out-of-threshold decisions to occur frequently, even though the PLL may have settled to complete equilibrium. The random out-of-threshold decisions cause the tub to drain more than it otherwise should. To counteract the excess draining caused by random out-of-threshold decisions, one must increase η, but the question is, by how much?

HOW TO ADJUST THE FILL/DRAIN RATIO

Consider a sequence of M samples delivered to the lock detector. With the PLL in equilibrium and with no jitter present, the lock detector should add M fill buckets to the tub, because every sample should be within the lock threshold. Thus, the sequence of M samples produces a net volume (V) of water in the tub as follows:

$$V_{NoJitter} = (M)(\textit{fill rate})$$

However, if jitter produces random out-of-threshold decisions, then each of the out-of-threshold samples removes one drain bucket from the tub. The fraction of the M samples that cause draining to occur is P_{OUT}, which is the probability that a sample exceeds the lock threshold range. Likewise, the probability that a sample is within the lock threshold range is P_{IN}, which relates to P_{OUT} as:

$$P_{OUT} = 1 - P_{IN}$$

If the jitter characteristics (μ and σ) are known, then P_{IN} can be computed using the P(α,β) formula in the Appendix: Normal (Gaussian) Distribution in the Context of Jitter section. Knowing P_{IN} allows the formulation of the net volume produced by the M samples in the presence of jitter:

$$V_{Jitter} = (M)(P_{IN})(fill\ rate) - (M)(P_{OUT})(drain\ rate)$$
$$= (M)(P_{IN})(fill\ rate) - (M)(1 - P_{IN})(drain\ rate)$$

Note that V_{Jitter} is less than $V_{NoJitter}$, because V_{Jitter} contains a drain component, whereas $V_{NoJitter}$ does not. Clearly, one must choose a larger fill rate for V_{Jitter} in order to yield the same volume as $V_{NoJitter}$. Using the variable for new fill rate (NFR) and setting $V_{Jitter} = V_{NoJitter}$ yields

$$(M)(P_{IN})(NFR) - (M)(1 - P_{IN})(drain\ rate) = (M)(fill\ rate)$$

Dividing both sides of the above equation by (M)(drain rate) yields

$$(P_{IN})(NFR)/drain\ rate - (1 - P_{IN}) = fill\ rate/drain\ rate$$

Solving for new fill rate (NFR)

$$new\ fill\ rate = (fill\ rate/P_{IN}) + (drain\ rate)[(1/P_{IN}) - 1]$$

The preceding formula enables one to adjust the jitter-free fill rate so that the lock detector behaves in a jittery environment as though it were in a jitter-free environment. However, if the rms jitter (σ) is too large relative to the lock threshold, then the density of undesired out-of-threshold decisions becomes overwhelming to the point that the lock detector bathtub can drain down to the unlock level causing an out-of-lock indication even though the PLL has settled into complete equilibrium. Obviously, no further compensation is possible once the fill rate reaches 255. In fact, if the fill/drain ratio skews too far in either direction, the usefulness of the lock detector as a lock/unlock indicator becomes questionable.

As long as σ is not too large relative to the lock threshold, the most reasonable solution is to choose the smallest acceptable values for the fill rate and drain rate. There is no simple answer for how small σ must be in order to allow satisfactory lock detector function because lock threshold, η, μ, and σ together, all play a role in determining the fill rate and drain rate. As a general guideline, the preceding techniques should be viable as long as the lock threshold is no less than ½ σ (half the rms jitter).

SIMULATION OF A SAMPLE APPLICATION: SYNCHRONIZATION TO GPS (1 PULSE/SEC)

OVERVIEW

This section provides a sample application (GPS synchronization) and presents simulations of the lock detector for various parameter settings and conditions. Each simulation comprises a sequence of 50,000 DPFDL samples with the first 10,000 samples simulating the PLL in the process of acquisition and the remaining 40,000 samples with the PLL in equilibrium. The acquisition model is a decaying exponential that starts at twice the lock threshold value and completely decays after 10,000 samples (see Figure 3).

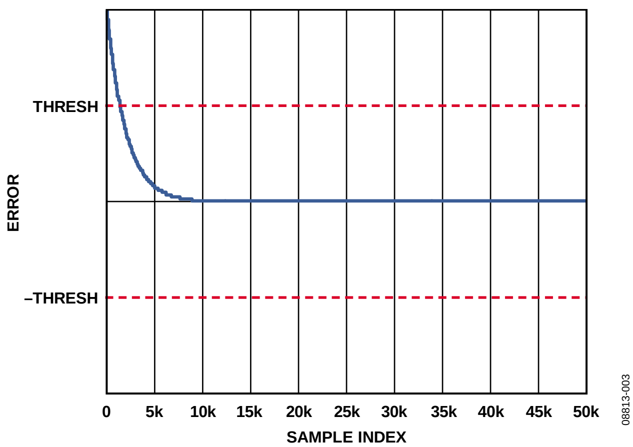

Figure 3. Acquisition Sequence

The simulation results in this section include plots of both the bathtub level and the resulting locked/unlocked indication. To provide for easy comparison, the bathtub and locked/unlocked indicator simulation plots each contain two traces, one blue and the other cyan. The blue trace shows the jitter-free simulation results while the cyan trace shows the simulation results with jitter applied.

Because a GPS receiver has an output rate of 1 pulse/sec, 50,000 samples equates to 50,000 seconds (nearly 14 hours) of lock detector operation. Note that the simulations always begin with the bathtub half full. Because the first 10,000 samples of each simulation models the PLL acquisition process, the initial samples are outside of the lock threshold for the first 1,400 samples or so. Therefore, the tub drains during this period, typically reaching a saturated unlock condition (empty). Once the acquisition process has sufficiently decayed, the tub begins filling because the signal falls within the lock threshold. This initial drain and fill process clearly shows up in the first part of each simulation.

Note that a typical GPS receiver exhibits jitter in the range of 75 ns rms (σ = 75,000 ps). In such an application, it makes sense to set the lock threshold at its maximum value of 65.535 ns (65,535 ps). This is desirable because it ensures that the lock threshold is greater than ½ σ.

Now, suppose that in a jitter-free environment a user wants the lock detector to be twice as responsive to indicating an unlock condition as to indicating a lock condition. This means that η = ½. If one wants the detector to be moderately responsive, then choose a fill rate of 25 and drain rate of 50.

BATHTUB LEVEL SIMULATION FOR ZERO-MEAN JITTER

Figure 4 shows a simulation of the bathtub level of the lock detector with fill and drain rates of 25 and 50, respectively. For the jittered trace (cyan), the jitter magnitude is 75 ns rms. Note the significant lag in the detector response when jitter is present due to the excess drain events associated with the additional jitter-induced out-of-threshold decisions.

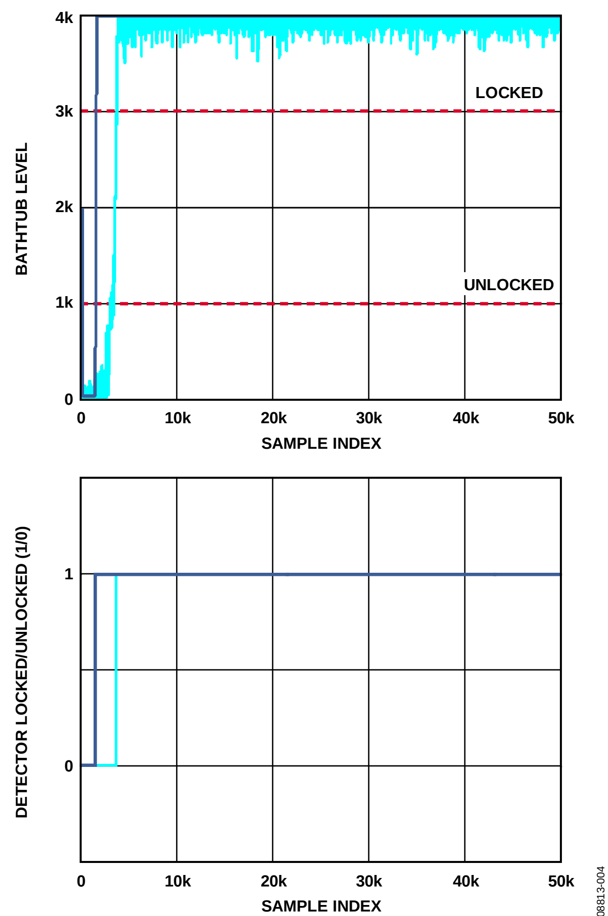

Figure 4. Simulation: Lock Threshold = 65,535 ps, σ = 75,000 ps, μ = 0 ps, Fill Rate = 25, and Drain Rate = 50

The lag in the response of the lock detector clearly indicates the need to find a new fill rate to counteract the jitter-induced draining of the tub. The first step is to determine the value of P_{IN} assuming μ = 0 ps (zero-mean jitter). Using the formula for $P(\alpha,\beta)$ in the Appendix: Normal (Gaussian) Distribution in the Context of Jitter section with α = −65,535 ps (negative lock threshold), β = +65,535 ps (positive lock threshold), σ = 75,000 ps, and μ = 0 ps, yields

$$P_{IN} = P(\alpha,\beta)$$
$$= P[(\beta - \mu)/\sigma] - P[(\alpha - \mu)/\sigma]$$
$$= P[(65535 - 0)/75000] - P[(-65535 - 0)/75000]$$
$$= 0.61777$$

This implies that undesired out-of-threshold decisions occur 38% of the time (100% – 62%). In fact, this represents a best-case scenario because of our zero-mean assumption. In a nonzero-mean scenario ($\mu \neq 0$), the percentage of undesired out-of-threshold decisions is larger than the 38% calculated.

Next, apply P_{IN} to the formula for calculating a new fill rate.

$$new\ fill\ rate = (fill\ rate/P_{IN}) + (drain\ rate)[(1/P_{IN}) - 1]$$
$$= (25/0.61777) + (50)\ [(1/0.61777) - 1]$$
$$= 72 \text{ (rounded up to the nearest integer)}$$

Figure 5 shows the same simulation, but with a fill rate of 72 and drain rate of 50.

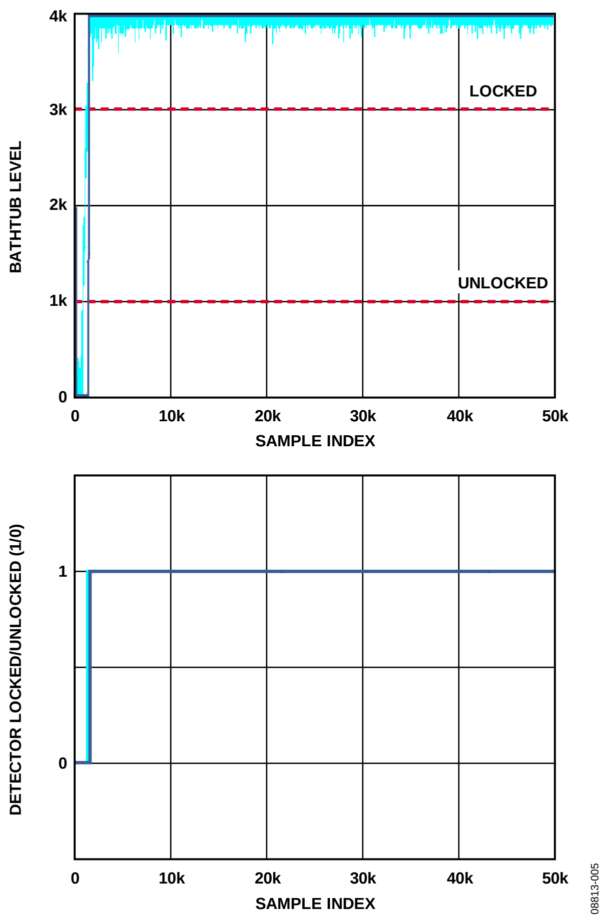

Figure 5. Simulation: Lock Threshold = 65,535 ps, σ = 75,000 ps, μ = 0 ps, Fill Rate = 72, Drain Rate = 50

Notice how the new fill value causes the detector response to approach the ideal (no jitter) response. The increased fill rate (72 vs. 25) effectively counteracts the excess drain events caused by the jitter-induced out-of-threshold decisions. Keep in mind, however, that this scenario is a best-case scenario because of the zero-mean assumption. As the PLL attempts to track a drifting input signal, the zero-mean assumption is no longer valid.

BATHTUB LEVEL SIMULATION FOR JITTER THAT IS NOT ZERO MEAN

Now, suppose that input drift causes the mean jitter to shift by as much as 50% of the lock threshold, which equates to $\mu = 32{,}768$ ps. The nonzero value of μ has a direct impact on the value of P_{IN}. Using the formula in the appendix for $P(\alpha,\beta)$ with $\alpha = -65{,}535$ ps (negative lock threshold), $\beta = +65{,}535$ ps (positive lock threshold), $\sigma = 75{,}000$ ps, and $\mu = 32{,}768$ ps yields

$$P_{IN} = P(\alpha,\beta)$$
$$= P[(\beta - \mu)/\sigma] - P[(\alpha - \mu)/\sigma]$$
$$= P[(65{,}535 - 32{,}768)/75{,}000] - P[(-65{,}535 - 32{,}768)/75{,}000]$$
$$= 0.57393$$

This implies that undesired out-of-threshold decisions occur 42% of the time (100% – 58%) as compared to 38% when $\mu = 0$ ps.

Figure 6 shows the simulation result using the original fill rate and drain rate values (25 and 50, respectively), but with $\mu = 32{,}768$ ps.

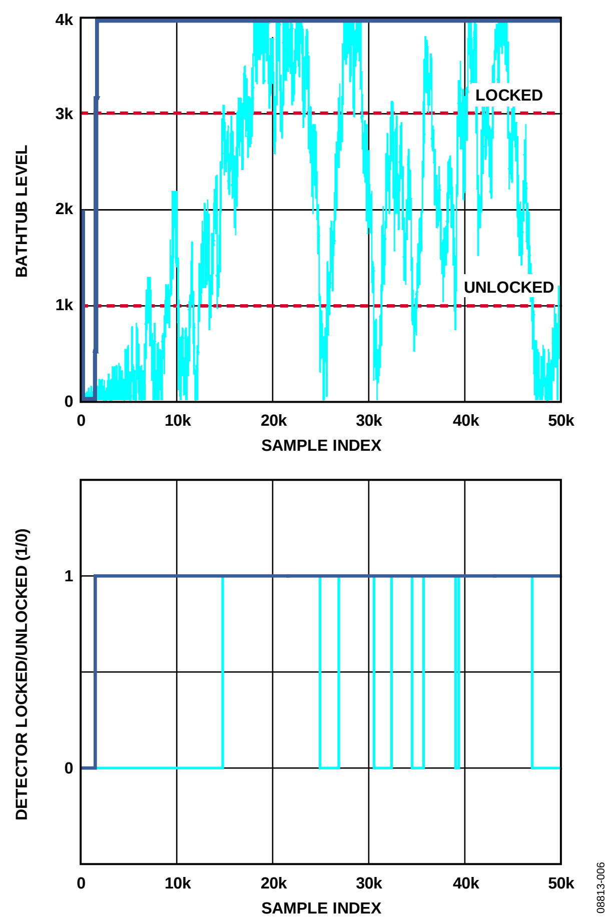

Figure 6. Simulation: Lock Threshold = 65,535 ps, σ = 75,000 ps, μ = 32,768 ps, Fill Rate = 25, Drain Rate = 50

The lock detector response changes dramatically due to the slight increase (38% to 42%) in the probability of out-of-threshold decisions. In fact, the detector displays several false unlock indications, which is a direct result of the shift in the mean value of the jitter (μ = 32,768 ps). Note that the mean signal level is still well inside the lock threshold (32,768 ps vs. 75,000 ps), but the jitter causes frequent undesired excursions outside of the lock threshold resulting in extraneous out-of-threshold decisions.

To compensate for these excess out-of threshold decisions, calculate a new fill rate (using P_{IN} = 0.57393).

$$new\ fill\ rate = (fill\ rate/P_{IN}) + (drain\ rate)[(1/P_{IN}) - 1]$$
$$= (25/0.57393) + (50)[(1/0.57393) - 1]$$
$$= 81 \text{ (rounded up to the nearest integer)}$$

Figure 7 shows the same simulation, but with fill rate = 81 and drain rate = 50. Notice once again how the new fill value causes the detector response to approach the ideal (no jitter) response. The increased fill rate (81 vs. 25) effectively counteracts the excess drain events caused by the jitter-induced out-of-threshold decisions, even with the mean value of the jitter shifted by 32,768 ps.

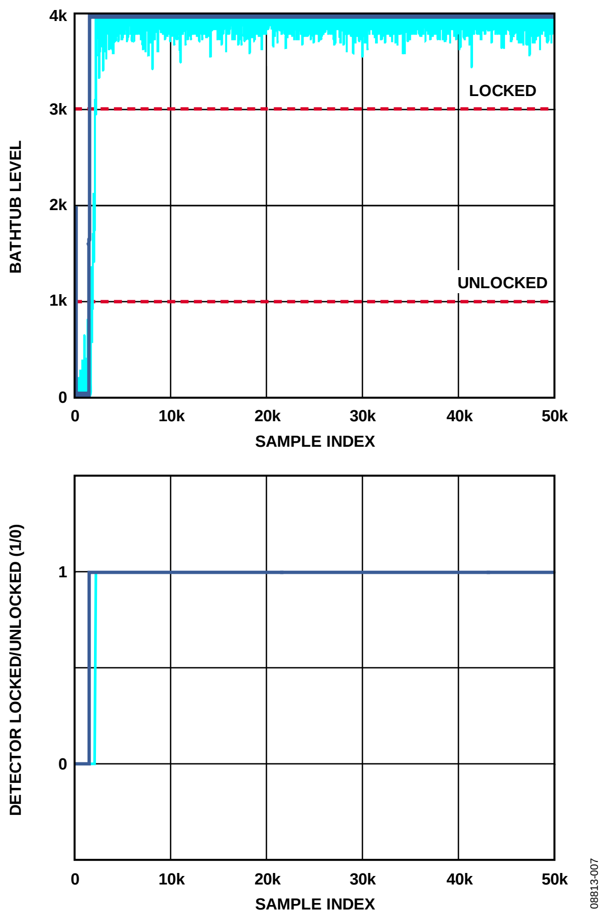

Figure 7. Simulation: Lock Threshold = 65,535 ps, σ = 75,000 ps, μ = 32,768 ps, Fill Rate = 81, Drain Rate = 50

BATHTUB LEVEL SIMULATION FOR A MINIMALLY RESPONSIVE DETECTOR WITH JITTER THAT IS NOT ZERO MEAN

A minimally responsive detector is one that uses the lowest possible values for the fill rate and drain rate. For example, consider the same parameters as in the previous simulation (an original desired ratio of η = ½, lock threshold = 65,535 ps, σ = 75,000 ps, and μ = 32,768 ps). A minimally responsive detector for which η = ½ implies a fill rate of 1 and drain rate of 2.

Figure 8 shows the resulting lock detector response. Note that with a fill rate of 1, the lock detector is unable to overcome the additional jitter-induced out-of-threshold decisions. Even after 50,000 samples, the tub level is well below the unlock indication level and the lock detector fails to indicate a lock condition.

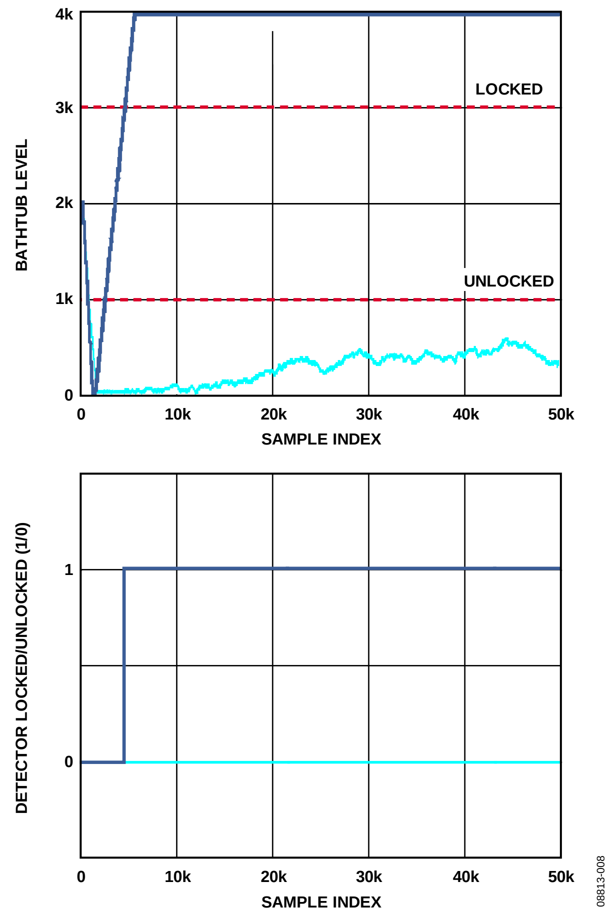

Figure 8. Simulation: Lock Threshold = 65,535 ps, σ = 75,000 ps, μ = 32,768 ps, Fill Rate = 1, Drain Rate = 2

Because μ and σ are the same (see the Bathtub Level Simulation for Jitter that is not Zero Mean section), the same value of P_{IN} applies (0.57393). Applying the formula for the new fill rate yields a fill rate of 4 (rounded up to the nearest integer).

Figure 9 shows the compensated response (for a fill rate of 4). Once again, notice that the detector response is reasonably close to the ideal response. Not only does the tub refill, it does so at a rate commensurate with the jitter-free condition (the blue trace).

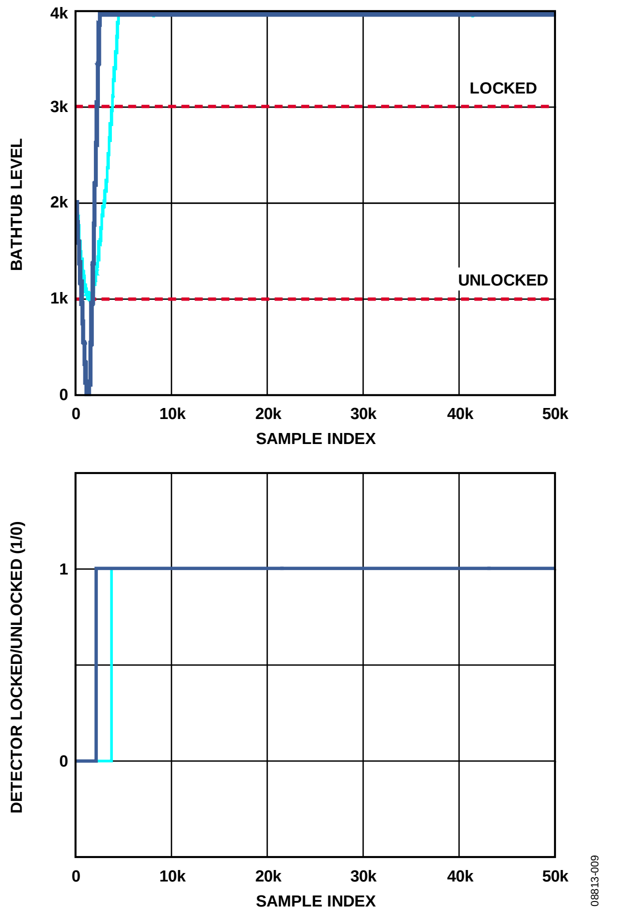

Figure 9. Simulation: Lock Threshold = 65,535 ps, σ = 75,000 ps, μ = 32,768 ps, Fill Rate = 4, Drain Rate = 2

CONCLUSION

The simulated lock detect times appear to indicate a discrepancy with respect to the jitter-free and jittery input signals in spite of the effort to compensate for jitter-induced out-of-threshold decisions. For example, the time at which the detector indicates lock is slightly different for the jitter-free vs. jittery input signals. In addition, as the bathtub level rises toward the lock level mark, the slope of the jittery trace (cyan) tends to be less steep than the jitter-free trace (blue).

The cause of the discrepancy is the fact that the first 10,000 samples constitute the PLL acquisition process, during which μ is greater than the value used in the computation of P_{IN}. Recall that the P_{IN} computation assumes that the PLL has settled to equilibrium. Therefore, the jitter statistics during the acquisition process are different from those at equilibrium, so the new fill rate value does not correctly compensate during the acquisition process. Thus, using the bathtub analogy, the net influx of water to the tub deviates from the expected amount during the acquisition process. The result is a variation in the time it takes the detector to indicate lock (usually later because the jitter statistics favor excess out-of-threshold decisions during the acquisition process). To be clear, the calculated new fill rate value properly compensates for the excess jitter-induced out-of-threshold decisions with the PLL in equilibrium. The apparent variation in lock indication time in the simulations is an artifact of the acquisition process, not a flaw in the procedure for determining the new fill rate.

The lock detector simulations presented herein are equally applicable to the phase lock detector and frequency lock detector. The only difference is the nature of the samples processed by each. The phase lock detector processes time error samples, while the frequency lock detector processes period error samples. Aside from the nature of the input samples, the phase lock and frequency lock detectors function identically.

The flexibility of the AD9548 lock detectors (phase and frequency) allows one to independently tailor the response of the detectors for both lock and unlock indication. Furthermore, with knowledge of the jitter statistics of the reference input signal, this application note demonstrates a method for modifying the detector response to compensate for the adverse affects of jitter. Although this application note focuses on jitter with a Gaussian distribution, the concepts presented herein are extendable to other distributions (uniform, for example), so long as one has knowledge of the statistical properties of the input jitter.

APPENDIX: NORMAL (GAUSSIAN) DISTRIBUTION IN THE CONTEXT OF JITTER

Generally, one can assume that random jitter follows the well-known normal distribution. This is reasonable because random jitter is typically due to the presence of additive white Gaussian noise (AWGN) in the system.

Two parameters describe a normal distribution: the mean (μ) of the distribution and its standard deviation (σ). The distribution appears in Figure 10. In the context of jitter, the x-axis represents phase error samples that deviate from their ideal noise-free values by some random amount. Frequently, a distribution has a mean of zero ($\mu = 0$), resulting in a zero-mean normal distribution (see Figure 11). The significance of σ (in either case) is that 68% of the total area under the curve resides between $x = \pm\sigma$. Note that because the normal distribution peaks at $x = \mu$, samples are much more likely to have values near μ than near the endpoints (±∞).

The concept of the normal distribution leads to the cumulative probability of the normal distribution. The cumulative probability, P(α,β), is the probability that a particular sample is between two arbitrary values (α and β) on the x-axis of the distribution, where:

$$P(\alpha,\beta) = \int_{\alpha}^{\beta} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} dx$$

This integral has no closed form solution, so it must be evaluated using numerical methods. Alternatively, one can use the function P(z), the standard normal cumulative distribution function, for which tabulated values appear extensively in the literature.

It has the form

$$P(z) = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}x^2} dx$$

P(α,β) is expressible in terms of P(z) as follows (where $\alpha < \beta$)

$$P(\alpha,\beta) = P(z_1) - P(z_2) = P(\frac{\beta-\mu}{\sigma}) - P(\frac{\alpha-\mu}{\sigma})$$

Note that unless the value of σ and μ is known, the P(α, β) formula is of no use. Fortunately, $\mu = 0$ can usually be assumed and σ, it so happens, is the same as the root-mean-square (rms) value of the jitter.

This is convenient because PLLs often operate in a system that must meet the requirements of a government or industry standard (SONET, for example). The standard, in such cases, may impose bounds on the amount of jitter present on the input signal. For example, the standard may specify input jitter as being less than 100 ps rms, which means that σ is no more than 100 ps. On the other hand, it may specify input jitter as being no more than 2 ns peak.

It is usually safe to convert a peak specification by substituting it for a value of σ that is one-sixth of the peak jitter value. The reason is that the probability of a sample being outside the ±6σ region is approximately 1 in 500,000,000 (virtually nonexistent). Thus, a 2 ns peak specification means it is safe to assume that σ is, at most, 1/6 of 2 ns, or 667 ps.

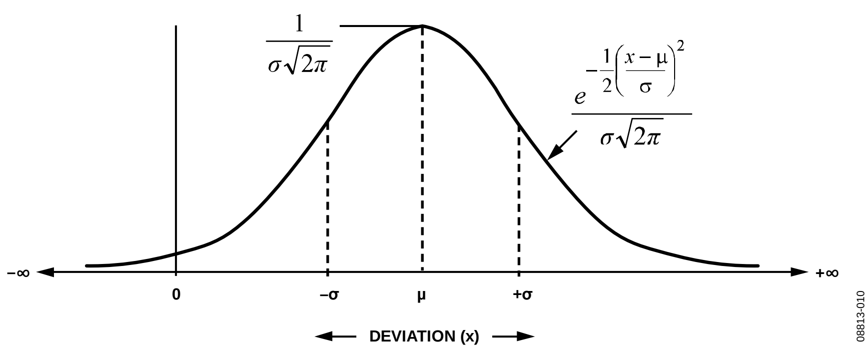

Figure 10. Normal Distribution

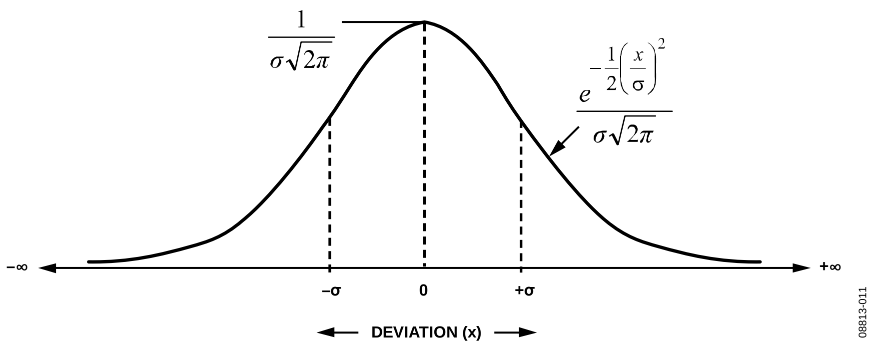

Figure 11. Zero-Mean Normal Distribution

Sometimes, one has no knowledge of the jitter characteristics of the input signal. In these situations, one must either take measurements to determine the value of σ, or simply take an educated guess.

As an illustrative example, consider the case in which the lock threshold level is set to 7.5° and it is known that the jitter is zero mean (μ = 0) with a standard deviation of 5° (σ = 5°). Note that 7.5° equates to 1.5 σ. The probability that a phase detector sample is within the lock threshold range of ±7.5° (that is, α = −7.5° and β = 7.5°) is

$$P(\alpha,\beta) = P[(\beta - \mu)/\sigma] - P[(\alpha - \mu)/\sigma]$$
$$= P[(7.5 - 0)/5] - P[(-7.5 - 0)/5]$$
$$= P(7.5/5) - P(-7.5/5)$$
$$= P(1.5) - [1 - P(1.5)]$$
$$= 0.933193 - (1 - 0.933193)$$
$$= 0.866386$$

Here, the value of P(1.5) comes from tabulated values of the standard normal cumulative distribution. Because most tables only list values for positive z, use the relationship, P(−z) = 1 − P(z) to handle negative values of z. The preceding calculation shows that there is an 87% chance of a phase detector sample being between ±7.5°. Figure 12 is a visual representation of this example, with the shaded region indicating the lock threshold range (±7.5° or ±1.5 σ). The dots represent a time series of random samples. Note how the samples tend to cluster within ±σ, as expected.

Now, consider the same example, but this time the jitter is not zero mean. Instead, it has a mean value of +7.5° (μ = 7.5°). The probability that a phase detector sample is within the lock threshold range of ±7.5° (that is, α = −7.5° and β = 7.5°) is

$$P(\alpha,\beta) = P[(\beta - \mu)/\sigma] - P[(\alpha - \mu)/\sigma]$$
$$= P[(7.5 - 7.5)/5] - P[(-7.5 - 7.5)/5)$$
$$= P(0) - P(-3)$$
$$= P(0) - [1 - P(3)]$$
$$= 0.5 - (1 - 0.998650)$$
$$= 0.498650$$

Here, the value of P(0) and P(3) come from tabulated values of the standard normal cumulative distribution (again, use the relationship, P(−z) = 1 − P(z) to handle negative values of z). Note that with μ = 7.5° there is a 50% chance of a sample being between ±7.5°, whereas with μ = 0° the probability is 87%. Figure 13 is a visual representation, again with the shaded region indicating the lock threshold range (±7.5° or ±1.5 σ). Note how the samples still tend to cluster within ±σ, as expected, but with the entire group offset from zero by μ (7.5° or 1.5 σ, in this case).

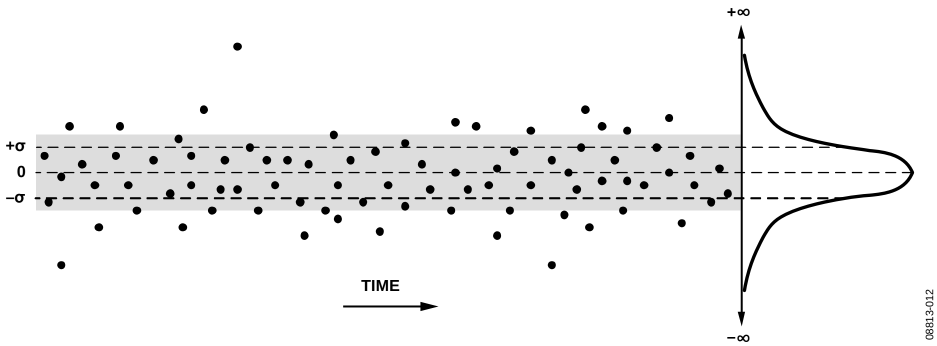

Figure 12. Zero-Mean Phase Error Sequence

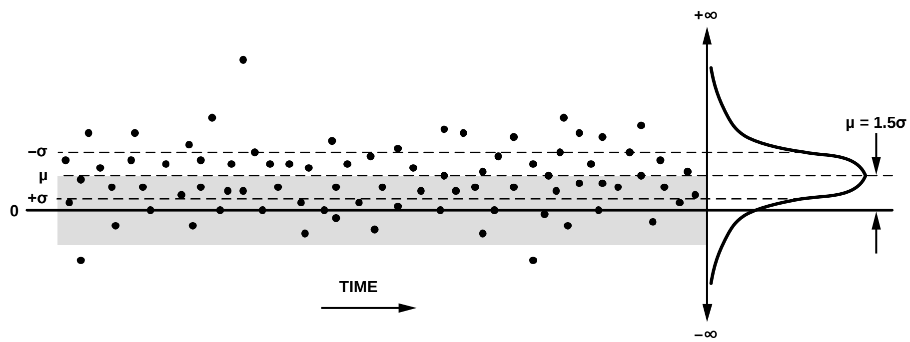

Figure 13. Phase Error Sequence Offset by μ

AN08813-0-3/10(0)