
A PRIMER ON JITTER, JITTER MEASUREMENT AND PHASE-LOCKED LOOPS

1. Introduction

As clock speeds and communication channels run at ever higher frequencies, engineers who have previously had little need to consider clock jitter and phase noise are finding that they need to increase their knowledge of these subjects. This primer provides an overview of jitter, offers practical assistance in making jitter measurements, and examines the role phase-locked loops (PLLs) have in this field.

2. Classes of Jitter

Let's begin by defining a few classes of jitter, followed by a discussion of how and why the different classes are important for some applications. Jitter falls into two broad categories: random jitter and deterministic jitter (see Figure 1).

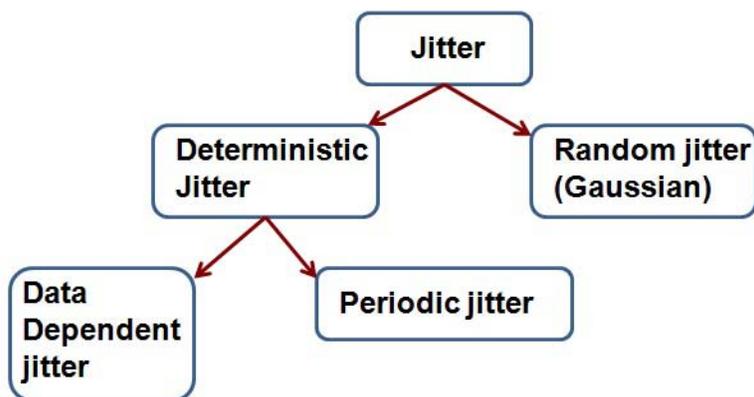


Figure 1. Classes of Jitter

Random jitter is a broadband stochastic Gaussian process that is sometimes referred to as intrinsic noise because it is always present. In the various phase noise plots shown later in this primer, the relatively smooth sections at the bottom of the curves represent the noise floor and are indications of random jitter. Though random jitter is often referred to as background or thermal noise, it can easily be the most significant contributor to overall jitter. Because random jitter is modeled as a Gaussian distribution, its instantaneous noise value is, strictly speaking, mathematically unbounded. Looking at the Gaussian curve in Figure 2, the two tails extending away from the center of the curve asymptotically approach (but never fully reach) zero. Although the probability of some values is very small, it is not zero, and, therefore, Gaussian noise is considered to be unbounded.



all-electronics.de
ENTWICKLUNG. FERTIGUNG. AUTOMATISIERUNG



Entdecken Sie weitere interessante Artikel und News zum Thema auf [all-electronics.de](https://www.all-electronics.de)!

Hier klicken & informieren!



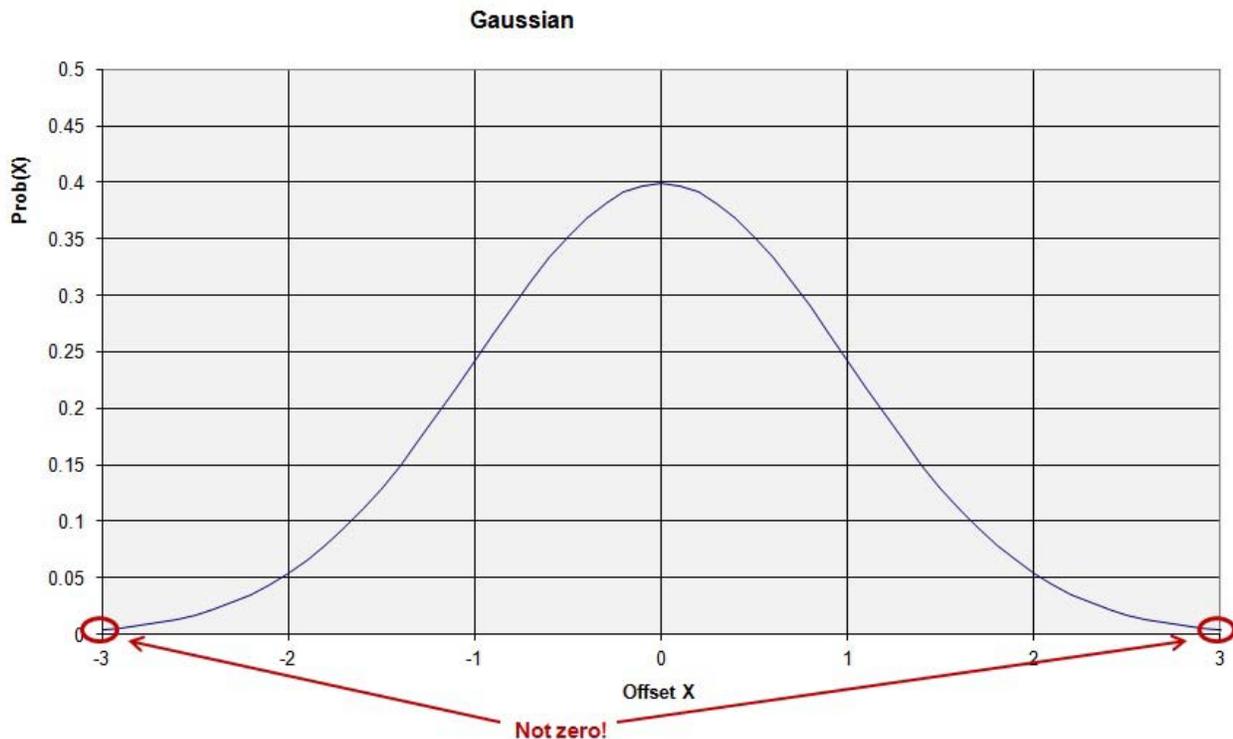


Figure 2. Gaussian Distribution

Deterministic jitter is not intrinsic or random and has a specific cause, though it may be very difficult to determine the specific cause. It is often periodic and narrowband. For example, deterministic noise may be caused by crosstalk due to an interfering periodic signal, either from a switching power supply or simultaneous switching of high-speed outputs on an IC (for example).

Deterministic noise can be further subclassified into periodic jitter and data-dependent jitter. The example of an interfering noise coming from a switching power supply is periodic because the noise will have the same frequency as the switching power supply. In contrast, an example of quasiperiodic, data-dependent jitter is intersymbol interference (ISI) caused by an isochronous 8B/10B coded serial data stream (e.g., Ethernet PCI Express). Such serial data coding involves dynamically changing duty cycles and irregular clock edges due to serial bit patterns, both of which contribute to overall jitter. Given that clock signals do not exhibit any data-dependent jitter (this primer is solely concerned with clock signals), there will be no further mention of it in this article.

Jitter can also be classified as correlated or uncorrelated. When inquiring about correlation, it is always important to ask “Correlated to what?” Correlation is between two events or distributions. Periodic jitter is always correlated to (or caused by) something. An example of correlated aperiodic jitter is when a serial data line interferes with a clock signal. The interference is correlated to the serial data line, but the serial data line may be either periodic or aperiodic. In all likelihood, a serial data transmission will have both periodic and aperiodic components. Two events are uncorrelated if they are statistically unrelated. Uncorrelated jitter is not statistically connected to an identifiable cause or influence. Random jitter is always uncorrelated. However, uncorrelated jitter is not always random, and it is possible to have two jittered signals both with periodic jitter that is uncorrelated. In these cases, the two uncorrelated sources of jitter do not exhibit any influence over one another.

A key point is that both correlated and periodic jitter have causes. Once the causes have been identified, it is sometimes possible to take remedial actions to reduce a portion of the overall jitter.

3. Types of Jitter Measurements

Different types of jitter measurements are commonly used, and it is important to use the jitter measurement that best fits the particular application at hand.

As shown in Figure 3, peak cycle-to-cycle jitter is the maximum difference between consecutive, adjacent clock periods measured over a fixed number of cycles, typically 1 k or 10 k. Cycle-to-cycle jitter is used whenever there is a need to limit the size of a sudden jump in frequency. The term peak-to-peak is defined as the difference between the smallest and the largest period value sampled during a measurement. For example, peak-to-peak jitter is useful when driving a PLL. In this situation, it can be desirable to limit the size of an instantaneous change in frequency to ensure that downstream PLLs remain in lock.



Figure 3. Cycle-to-Cycle Jitter

As shown in Figure 4, peak-to-peak period jitter is the difference between the largest clock period and the smallest clock period for all individual clock periods within an observation window (typically 1 ku or 10 ku cycles). It is a very useful specification for guaranteeing the setup and hold time of flip flops in digital systems.

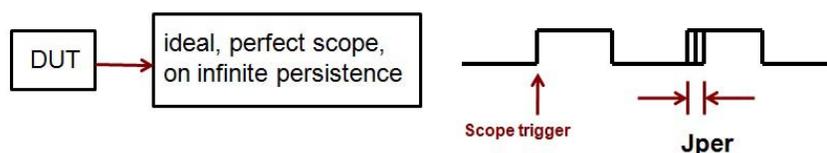


Figure 4. Period Jitter

As shown in Figure 5, time interval error (TIE) jitter, also known as accumulated jitter or phase jitter, is the actual deviation from the ideal clock period over all clock periods. It includes jitter at all jitter modulation frequencies and is commonly used in wide area network timing applications, such as SONET, Synchronous Ethernet (SyncE), and optical transport networking (OTN).

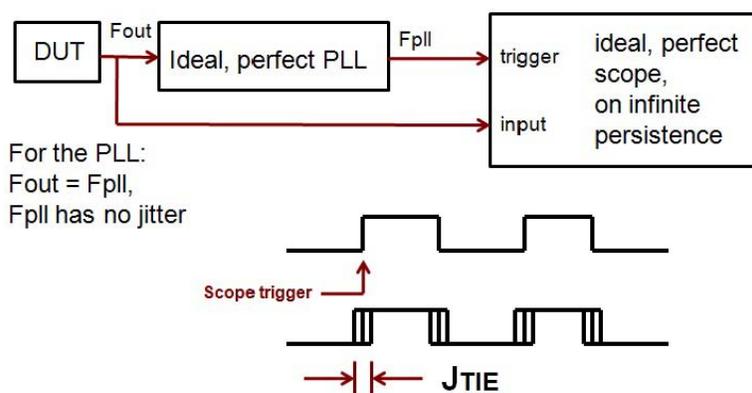


Figure 5. TIE Jitter

AN687

It is intuitive that the frequency jitter content of cycle-to-cycle jitter is high because the time frame affecting a given clock period extends no further out than one clock period. Similarly, it would seem that period jitter would have more high-frequency content than the TIE (or accumulated) jitter. This can be seen in Figure 6, which shows the results for three types of jitter measurements for a purely uniform, random, white-noise jittered clock signal. As can be seen, the TIE jitter is the same for all frequencies (a property of white noise), whereas the period and cycle-to-cycle jitter show different degrees of having greater high-pass content.

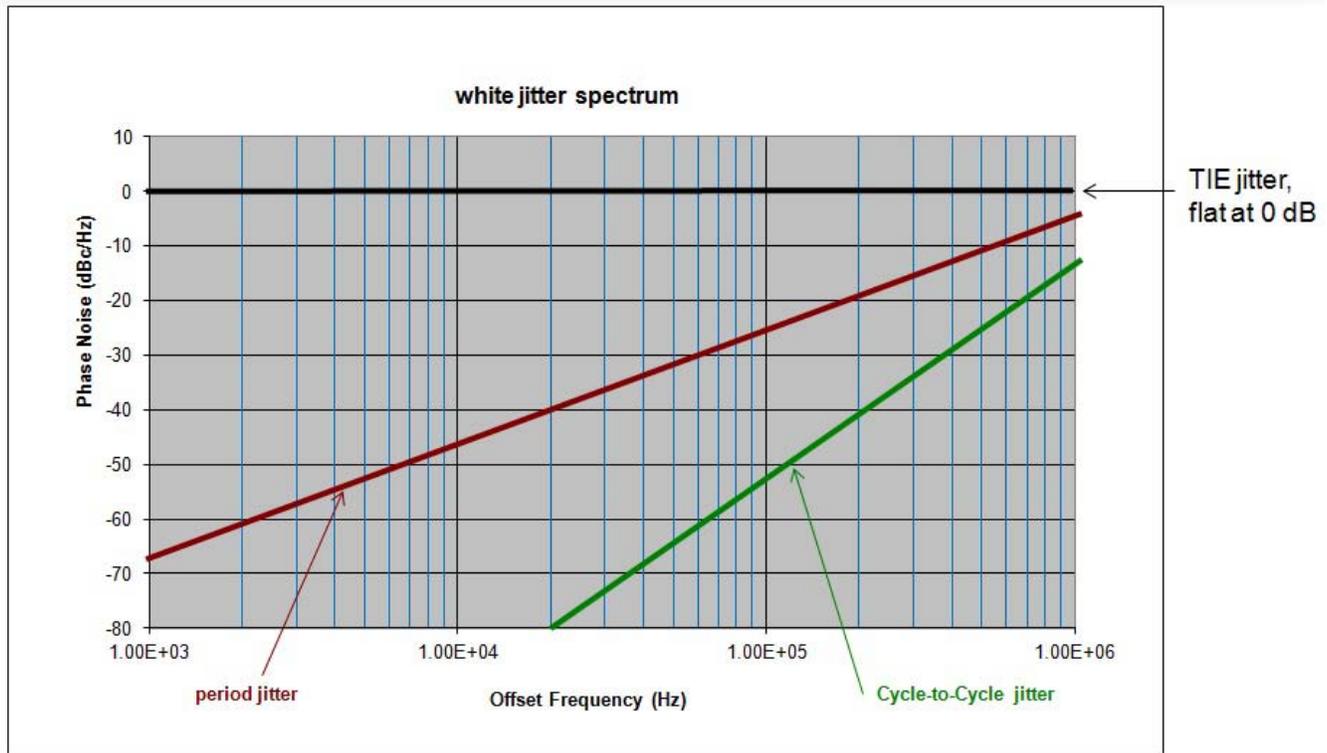


Figure 6. Frequency Content of Cycle-to-Cycle, Period and TIE Jitter

Different statistics can be taken for all types of jitter. That is, root mean squared (RMS), peak-to-peak, and other statistical values exist for cycle-to-cycle, period, and TIE jitter, although some are in more common use than others. Whenever peak-to-peak statistics are used, the number of samples taken needs to be large enough to ensure confidence in the measurement. Typically, such sample sizes range from 1,000 to 10,000.

It is often desirable to convert an RMS jitter value to a peak-to-peak number and vice versa. A common approach is to use a crest factor, which assumes a Gaussian noise model. A tolerable bit error rate (BER) value is either given or assumed, and the resulting crest factor is used to convert between the two. For example, with a BER of 10⁻¹², the RMS crest factor is 14. Therefore, a clock signal with 1 ps RMS of jitter would have a peak-peak jitter value of 14 ps. For more on this, see Reference 2, section 12.11.2.2.

4. Wander

By definition, wander is jitter that is within 10 Hz offset of the carrier. If the jitter is less than 10 Hz off of the carrier, it is called wander. If it is more than 10 Hz off of the carrier, it is called jitter. As will be seen later, it takes a very low bandwidth PLL to attenuate wander. In most systems, wander is not an issue because the system PLLs can easily track low-frequency wander, enabling them to wander together. Any differences in frequency between the segments are typically not significant because their first-in-first-out (FIFO) buffers can absorb the momentary discrepancies. Since frequency differences are so low, the FIFO does not need to be very deep.

However, wander can accumulate for some systems, and the wander can eventually become large enough to cause problems. Figure 7 shows two different systems. The first is a typical SONET application that implements centralized timing with a very low loop BW PLL to avoid wander accumulation. The second example is a Synchronous Ethernet (SyncE) line timing application that allows wander to accumulate. Both wander and jitter can accumulate in cascaded PLL systems.

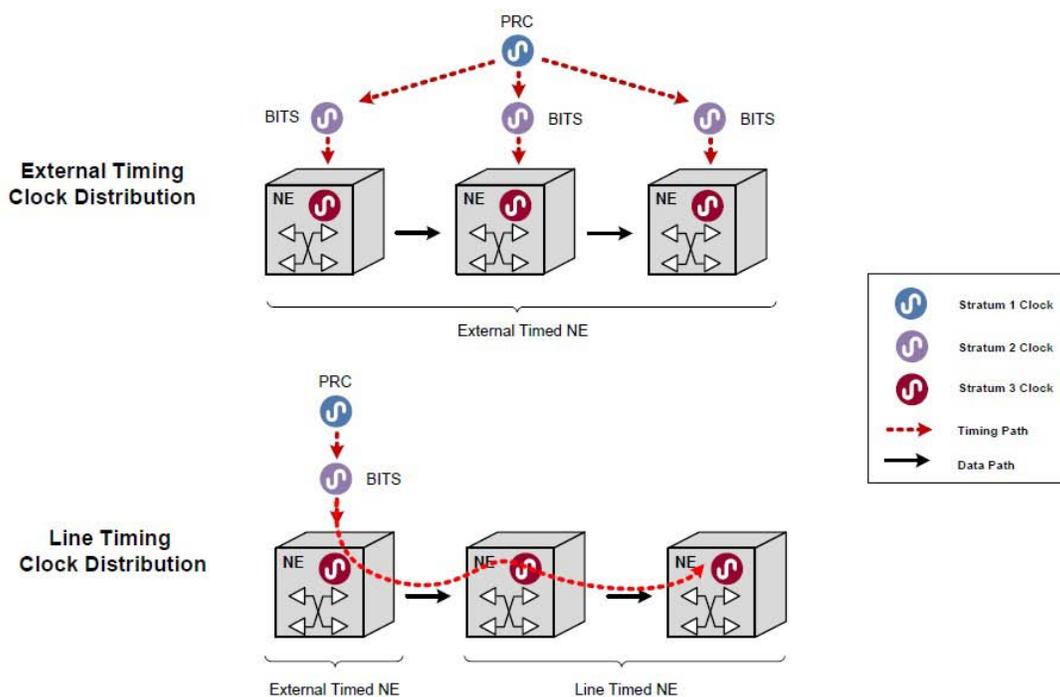


Figure 7. External Timing Clock Distribution

5. Time vs. Frequency Domain Measurements

Jitter is usually a time domain term, while phase noise is typically a frequency domain term. However, it is common for the terms to be used loosely, with the result that they are sometimes used interchangeably. Theoretically, with perfect measuring equipment, if one were to integrate the phase noise from the carrier out to an infinite offset, the phase noise and the jitter would have the same numeric value. However, given real-world and practical test equipment, this is simply not possible, and there will always be a discrepancy between the two measurements. In spite of this, it is important to remember that the time and frequency domain approaches both measure the same phenomena, even though the exact results depend on the details of the measurement process and the equipment used.

Most modern equipment that measures jitter falls into one of two broad categories: time domain and frequency domain. Time domain equipment typically comes in the form of a high-speed digitizing oscilloscope with high single-shot sampling bandwidth. Frequency domain equipment usually comes in the form of a spectrum analyzer, a spectrum analyzer with phase noise measurement capability, or a phase noise analyzer. Each of these two categories of equipment has its advantages and disadvantages. Let's take a closer look at the key differences between these two approaches to measuring jitter.

Time domain equipment has the virtue of being able to directly measure peak-to-peak, cycle-to-cycle, period and TIE jitter. This measurement approach permits the measurement of jitter of very low frequency clock (or carrier) signals. By post-processing the data with techniques, such as FFTs and digital filters, it is possible to integrate the phase noise value over a specific band of frequencies to generate RMS phase jitter values. Only time domain equipment can measure all of the jitter frequency components. Another key point is that time domain equipment is much better at measuring data-dependent jitter, which makes it very useful for high-speed serial links that use serializer/deserializer (SERDES) technology.

Frequency domain equipment cannot directly measure peak-to-peak, cycle-to-cycle, or period jitter because its native capability is to measure the RMS power of signals in a given frequency band. Frequency domain equipment is also awkward for measuring data-dependent jitter. However, the best frequency domain instruments have a lower noise floor than the best time domain instruments. This fact makes frequency domain instruments the first choice for ultra-low phase noise clock signal measurements that are free of data-dependent jitter. Table 1 summarizes the differences between time and frequency domain instruments.

Table 1. Time vs. Frequency Domain Instrument Differences

	Time Domain	Frequency Domain
Native Measurements	Peak-Peak Jitter Cycle-to-Cycle Jitter Period Jitter	RMS Phase Jitter Phase Noise Jitter Frequency Information
Advantages	Good with Low-Frequency Clocks Good with Data-Dependent Jitter	Lower Noise Floor Easy Detection of Spurs vs. Random Jitter

Various mathematical estimation and translation approaches can be used to switch from one type of jitter measurement to another. For example, as mentioned earlier, it is possible to use a crest factor and a desired BER to go back and forth between peak-to-peak and RMS jitter. Another option is to use a fast Fourier transform (FFT) of time domain data to provide frequency domain information and filtering. However, it should be remembered that most of these techniques rely on mathematical models that, while good approximations in most situations, have their limitations and should only be used with care and thought.

Figures 8 and 9 show examples of time and frequency domain measurements of the same 2.488 GHz sine wave and serve to highlight the strengths and weaknesses of both approaches. In Figure 8, the first two images were taken with a digitizing scope. The top trace is a sine wave from an RF generator without any modulation. The lower trace has a large amount of frequency modulation (FM). The small blue sections at the tops of the plots are histograms of the period of the signal shown by the red points. The upper histogram without the modulation indicates that the sine wave is comprised of mostly a single frequency. The lower histogram with modulation clearly shows that the sine wave spends much of its time at the extreme ends of the excursion in the clock period.

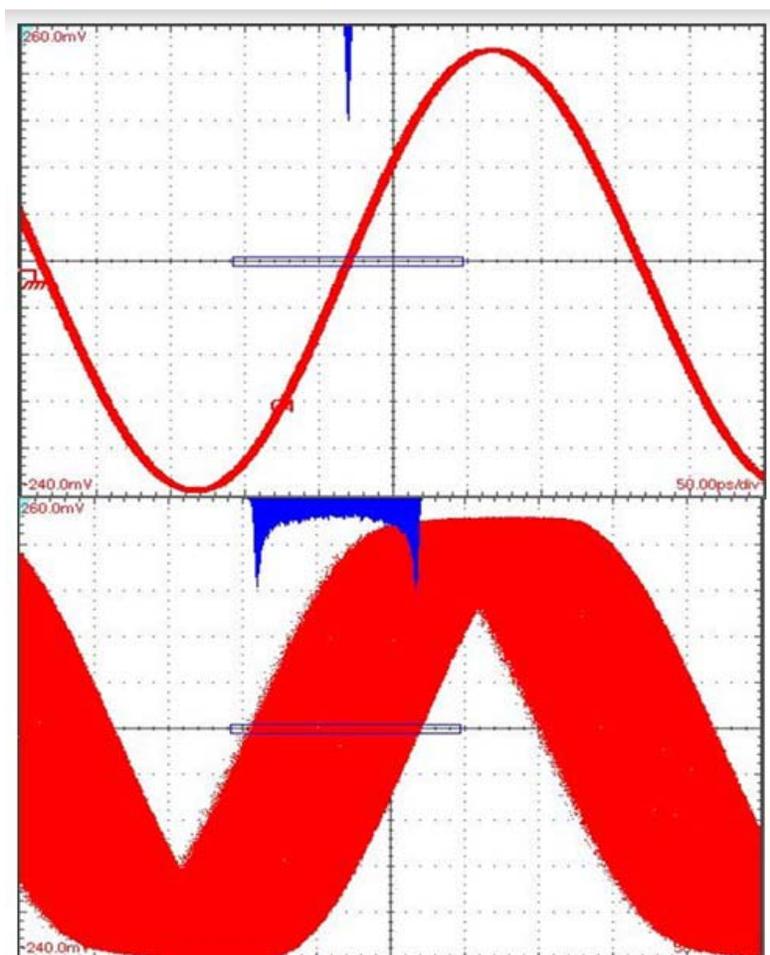


Figure 8. 2.488 GHz Clean and Jittered Scope Shots

Now, compare these two measurements to the two measurements in Figure 9 of the same signal taken with a spectrum analyzer. Again, the top plot is without modulation, and the lower trace has significant FM.

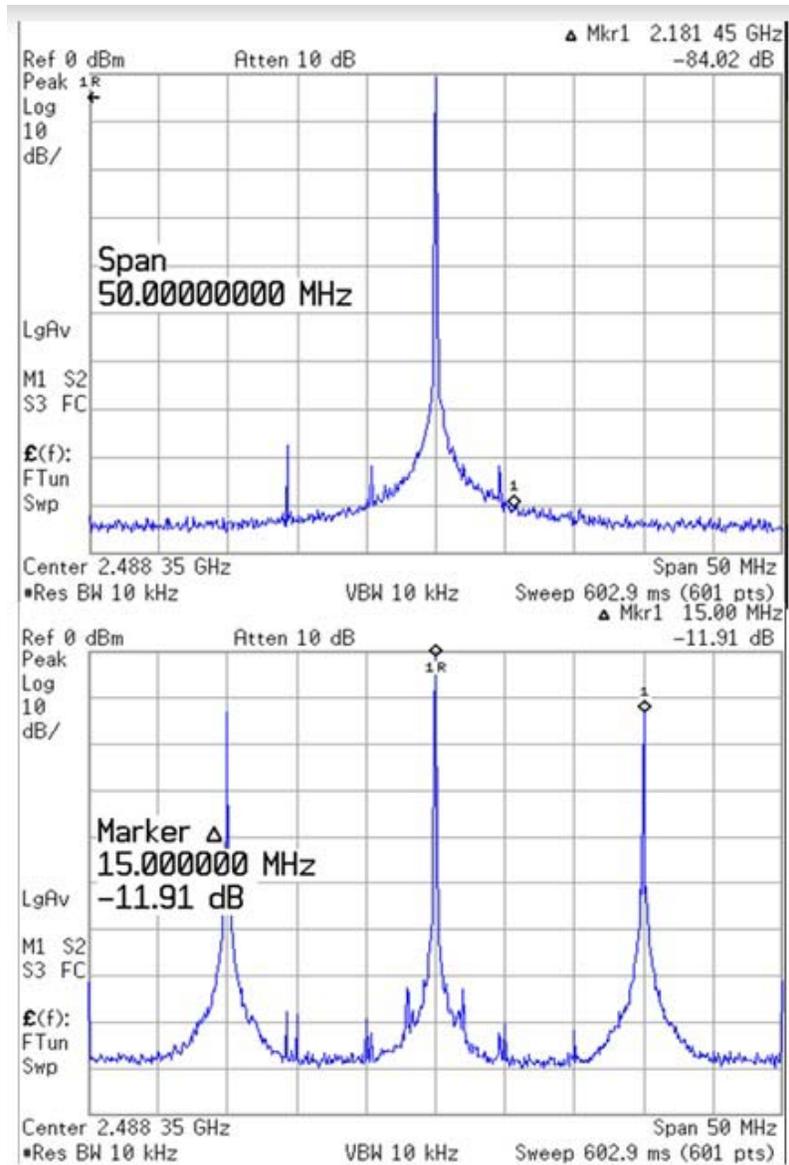


Figure 9. 2.488 GHz Clean and Jittered Spectra

The top spectrum shows that the sine wave consists mostly of a single frequency with its largest spurs below -75 dBc (75 dB below the carrier). The lower plot shows two symmetric and equal amplitude sidebands, which is an indication of FM. An advantage of a spectrum analyzer over an oscilloscope is that the modulation frequency of 15 MHz (in this case) is immediately apparent, as shown by the marker value.

6. Phase Noise Plots

High-end spectrum analyzers usually have phase noise plot capability. Test equipment specifically designed for phase noise plots is also available. Figure 10 is the spectrum of a 312.5 MHz carrier. As expected, the plot is mostly symmetric around the carrier with spurs located at both the plus and minus locations. From the marker, it can be seen that the largest spur is +1.13 MHz off of the carrier and is at -62 dBc (dB off of the carrier). Another spur that is roughly the same size is located at -1.13 MHz off of the carrier. The other spurs are similarly paired.

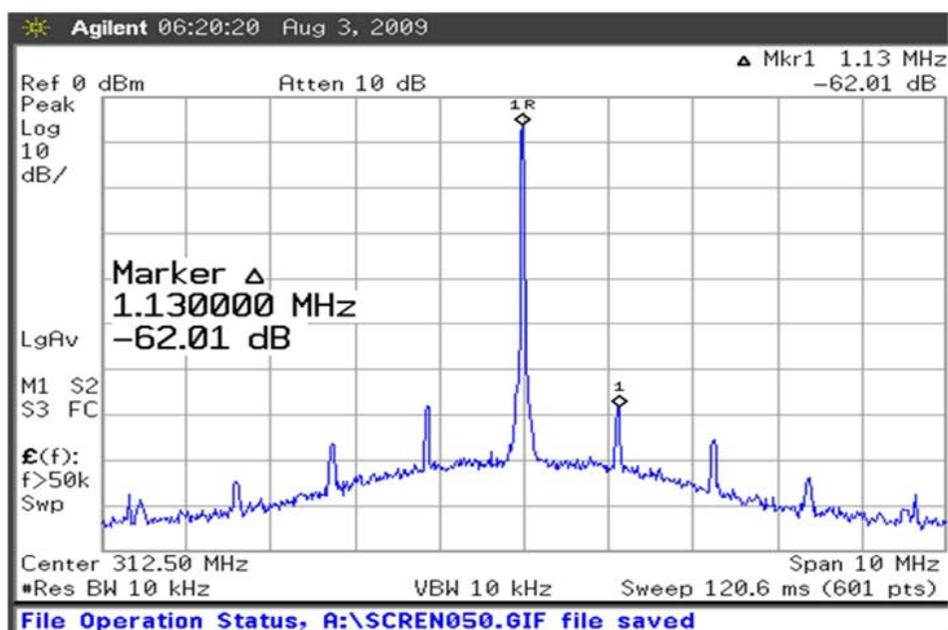


Figure 10. 312.5 MHz Clock Spectrum Example

Figure 11 shows a phase noise plot of the same signal. The carrier is located to the left of the plot, with the left-hand edge of the plot having a 100 Hz offset from the carrier and the right-hand edge having a 100 MHz offset from the carrier. The plot can get close to the carrier but will never include the carrier itself. In practice, the lowest offset from the carrier is usually anywhere from 1 Hz to 1 kHz. The phase noise plot itself is a single-sideband process in that only one half of the spectrum is taken, and the other half is assumed to be symmetric. Note that the largest spur is -62 dBc and is located at 1.13 MHz off of the carrier, just like the spectrum in Figure 9. When the RMS noise value is calculated, the RMS value is mathematically adjusted to include the contribution from the missing half of the spectrum. In this manner, a phase noise plot with a symmetric spectrum effectively combines the carrier frequency offsets that are both above and below the carrier into a single plot. Another noteworthy difference between spectrum plots and a phase noise plots is that the horizontal scale of phase noise plots is logarithmic, while a linear scale is typically used for spectrum plots.

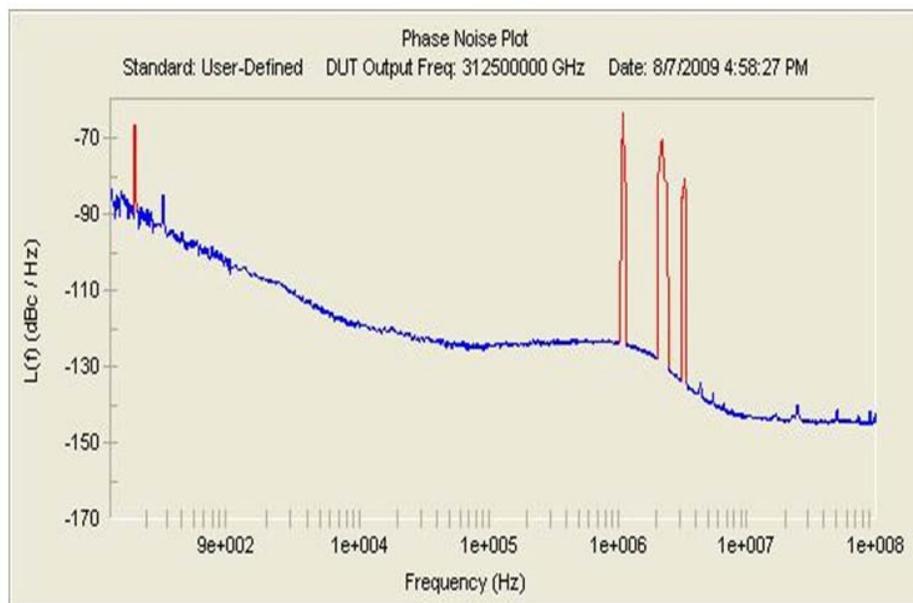


Figure 11. 312.5 MHz Phase Noise Plot Example

7. Spurs

The word “spur” is a shortened term for periodic spurious noise. See the red sections of the Figure 11 phase noise plot. Because spurs are periodic (and typically stationary), they are correlated to something (e.g., circuit or PLL artifacts), although it is often difficult to determine the source of the correlation. They are usually modeled as sine wave modulations of the carrier, and they are considered to be jitter at a single offset frequency.

As stated earlier, it is difficult to convert RMS to/from peak-to-peak values. While this is generally true, the specific case of conversion for spurs that are modeled as sine waves is well known. To convert from RMS to peak-to-peak, multiply by 2.2.

Typically, only large spurs contribute significantly to the RMS noise. For $S = \text{dBc}$ of a spur, the peak-to-peak jitter in UI (unit interval) is:

$$J_{\text{pk-pk in UI}} = \left(\frac{2}{\pi}\right) \times 10^{S/20}$$

For example, the phase noise contribution for the “A” spur will be calculated for the typical phase noise plot shown in Figure 12. The carrier frequency of 490 MHz appears in the upper right section of the plot at one of the green arrows. The size of the “A” spur is -79 dBc , and the total RMS noise integrated from 12 kHz to 20 MHz (which includes the “A” spur) is 339 femtoseconds.

$$J_{\text{pk-pk}} = \left(\frac{2}{\pi}\right) \times 10^{S/20} = \left(\frac{2}{\pi}\right) \times 10^{-79/20} = .000071 \text{ UI}$$

$$1 \text{ UI} = (490 \text{ MHz})^{-1} = 2.04 \text{ ns};$$

$$.000071 \text{ UI} \times 2.04 \text{ ns} = 146 \text{ fs pk-pk}$$

$$146 \text{ fs pk-pk} \Rightarrow \frac{146 \text{ fs}}{2 \times \sqrt{2}} = 51.6 \text{ fs RMS}$$

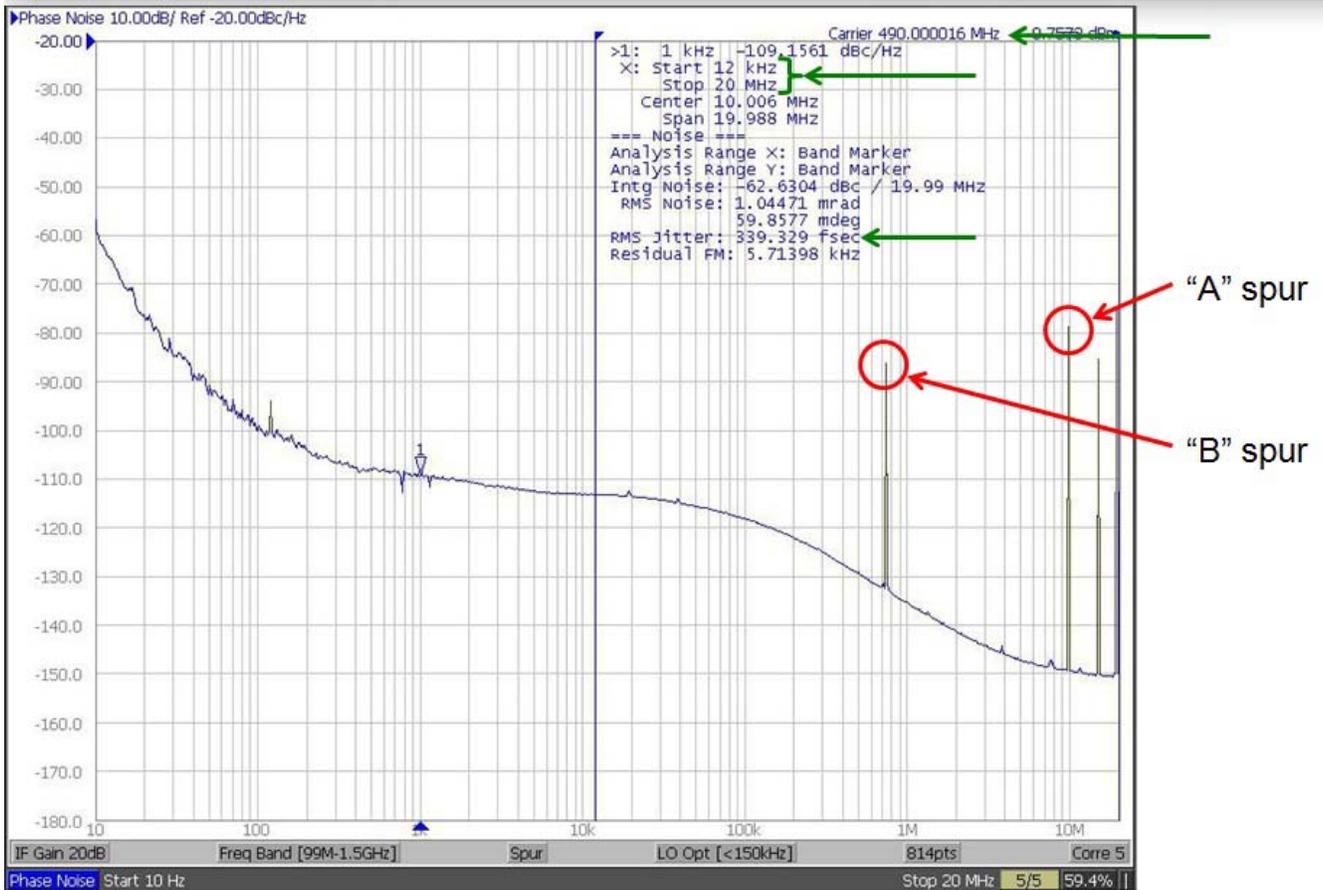


Figure 12. “A” and “B” Spur Phase Noise Plot in dBc

Using the above equations, the jitter contribution for the phase noise plot from the “A” and “B” spurs is shown in Table 2. This table shows that only large spurs make a significant contribution to RMS jitter.

Table 2. “A” and “B” RMS Spur Value

Spur	Pk-pk jitter	RMS Jitter
A	146 fs	51.6 fs
B	58 fs	20.5 fs

Spectrum and phase noise analyzers have a resolution bandwidth, which is an indication of the size of the frequency window that the instrument is using to sample the signal at a given point in the sweep. The output of the instrument at any given point on the curve is the average signal energy over a window that is the size of the resolution bandwidth of the instrument. If the resolution bandwidth is coarse (i.e., the resolution bandwidth has a large value), the instrument can quickly scan across a given frequency band. If the resolution bandwidth is fine (small), the time required for the same scan is longer. Since offset frequencies that are of interest can be more than 100 MHz off of the carrier, it is not practical for a phase noise plot to sweep the entire 100 MHz band with the same resolution for all frequencies. If the resolution bandwidth is uniformly high, the resolution for offsets that are close to the carrier will be insufficient. If the resolution bandwidth is uniformly low, it takes far too long to produce a phase noise plot. For this reason, phase noise equipment dynamically alters the resolution bandwidth as the scan progresses. The resolution bandwidth is very high when the offset frequency is close to the carrier and is lowered with greater offset values, which fits very well with a logarithmic horizontal axis.

The vertical unit of a phase noise plot is not dB or dBc, but dBc/Hz. As a result, all of the samples taken with different resolution bandwidth values are normalized to a common value of 1 Hz. This approach allows them to be placed side-by-side, creating what looks like a uniform sweep of data. Because spurs are sometimes modeled as sine waves (single tone modulation), the large-resolution bandwidth values that occur far from the carrier effectively spread the spur energy over a wider frequency range, misrepresenting the true size of the spurs on a dBc/Hz vertical axis. While this will not significantly affect the resulting RMS jitter value, some applications are more sensitive to spur levels than they are to RMS noise. In these situations, the strict use of a dBc/Hz vertical scale can be somewhat misleading. Spur levels are particularly important in wireless applications, which commonly use undersampling analog-to-digital converters. In contrast, wide area networks, such as SONET/SDH, OTN, and carrier Ethernet, are typically sensitive only to RMS jitter and usually are not concerned with representing spurs with dBc values.

Advanced phase noise equipment can detect the presence of a spur, remove it from the dBc/Hz phase noise plot, put the spur back into the phase noise plot with a vertical scale of dBc and then add its RMS value back in, with the RMS value modeled as sine wave modulation. Two different plots of the same phase noise data illustrate this. Figure 13 is purely in dBc/Hz and reports 292 femtoseconds RMS of jitter integrated from 12 kHz to 20 MHz.

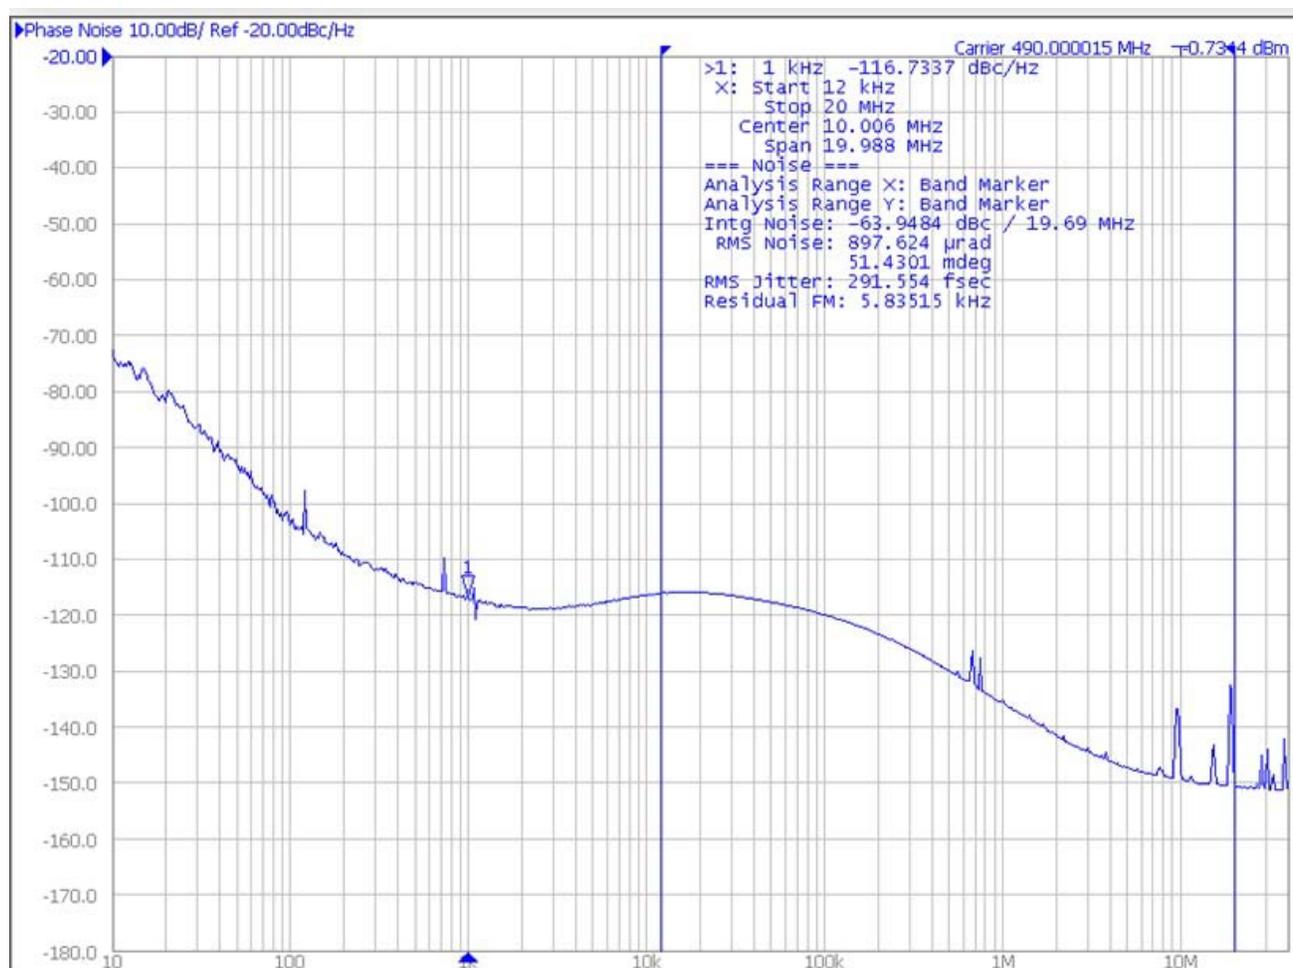


Figure 13. dBc/Hz Phase Noise Plot

AN687

Figure 14 shows a plot made with exactly the same raw data, but the spurs were removed from the plot and re-entered into the final RMS jitter value using the spur-to-RMS conversion process described above. The RMS jitter is now calculated to be 290 fsec, which is slightly less than the pure dBc/Hz integration value. Note that the spurs themselves were put back into the display of the curve with a different color because the vertical scale for these spurs is not dBc/Hz, but rather dBc. The non-spur data is blue and has a vertical scale of dBc/Hz, while the spur data is brown with a vertical scale of dBc.

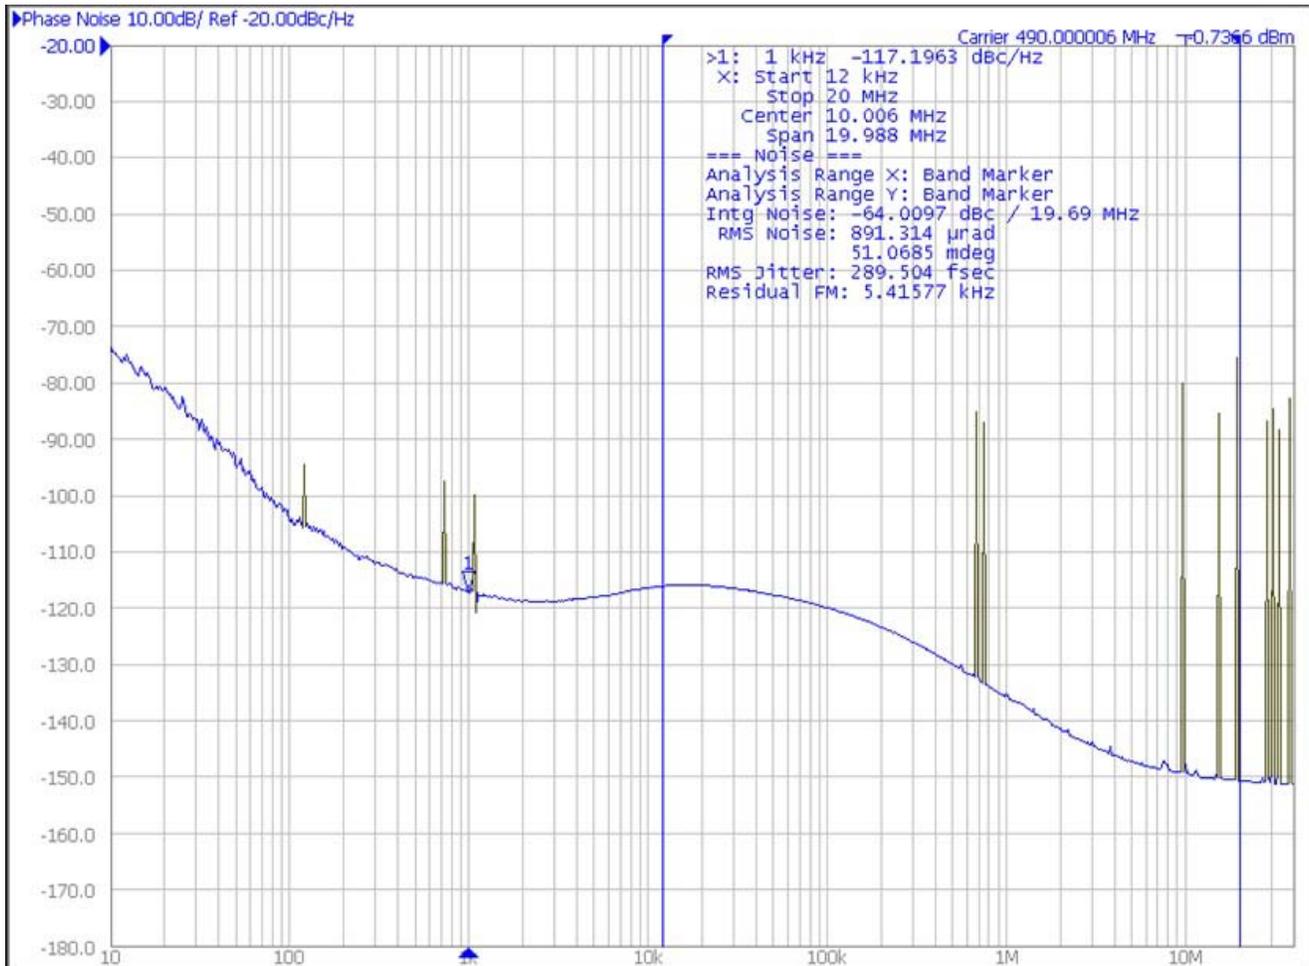


Figure 14. dBc Phase Noise Plot

8. PLL Characteristics

PLLs typically have two basic functions:

- Translate one frequency value to another
- Reduce or attenuate the jitter of a signal

Referring to the Figure 15 PLL block diagram, a jittered clock is applied to the phase detector whose output is proportional to the phase difference between the f_{OUT} feedback path and input clock. The first function of frequency translation is achieved by placing clock dividers with different divisor values at the two inputs to the phase detector. However, PLLs are often used because of their ability to perform the all-important jitter attenuation function. The following discussion will concentrate on the characteristics of jitter attenuation.

The output of the phase detector is low-pass filtered by the loop filter, which is an important element that shapes the PLL's behavior when presented with a jittered input. The loop filter drives the control voltage of a voltage-controlled oscillator (VCO). If the VCO frequency is too high, the low-pass phase detector output goes low, which reduces the VCO control voltage and hence the VCO frequency. In this manner, the VCO is forced to settle on a value that matches the long-term average of the input frequency. The value of "long term" is determined by the specifics of the loop components, including the low-pass filter.

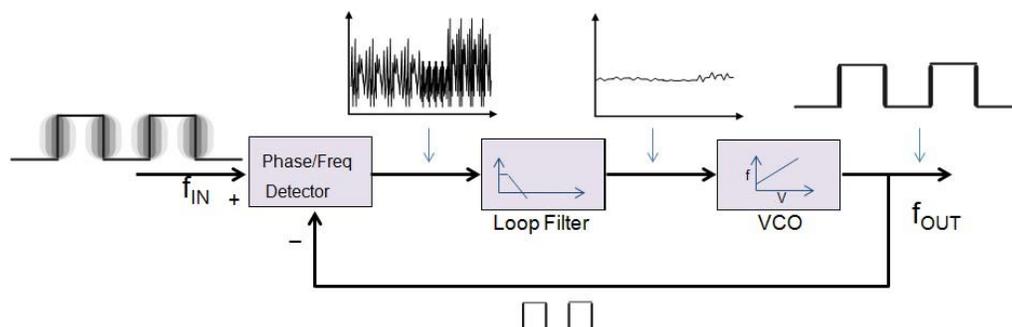


Figure 15. PLL Block Diagram

There are three main characteristics of a PLL that must be considered for PLL applications involving clock jitter:

- Jitter generation
- Jitter transfer function
- Jitter tolerance

Jitter generation is the amount of jitter at the output of the PLL when a "zero jitter" clock is applied to its input. Of course, there is no such thing as a zero jitter clock, and, in practice, this means that the input clock should have sufficiently low jitter for it to have no appreciable effect on the output jitter. Needless to say, there are some circumstances for which finding such an input clock is not a trivial matter. The jitter generation value varies significantly with the frequency offset from the carrier. Typically, jitter generation falls off as the offset value gets larger, although this is not always the case. The jitter generation curve of a PLL establishes its noise floor, and the PLL's performance at any given offset frequency will never be better than its jitter generation value at that particular offset frequency.

Figure 16 shows the phase noise versus offset frequency for three sources. The first is for a Rohde and Schwarz model SML 03 RF generator connected to the clock inputs of two different Silicon Labs monolithic PLLs. The two remaining curves are from the outputs of the PLLs when driven by the RF generator. The red curve has the best phase noise performance close to the 622.08 MHz carrier. For larger offset frequencies (i.e., after the crossover point indicated by the yellow arrow), the RF generator surpasses the performance of both PLLs. At these larger offset frequencies, the phase noise for the red and blue PLL curves is purely jitter generation because the phase noise at the PLL clock inputs is lower than the phase noise at the PLL outputs. That is, the phase noise at the PLL

AN687

outputs is due to the PLLs and not their input clocks. Close to the carrier, the green and red curves show that the RF generator's phase noise is being attenuated by both PLLs, though to a greater degree by the red PLL.

Accordingly, for points that are to the right of the tip of the yellow arrow, jitter generation is being displayed. The points to the left of the tip of the yellow arrow show jitter attenuation.



Figure 16. Jitter Generation Phase Noise Plots

The jitter transfer function gives the amount of jitter attenuation that will occur at a particular offset frequency. It is strongly affected by the loop bandwidth of the PLL, which is strongly affected by the PLL's low-pass loop filter. Accordingly, a PLL's response to jitter can be characterized as low-pass. That is, jitter at an offset frequency that is above the loop bandwidth cutoff frequency will be attenuated by the PLL, and jitter that is below the loop bandwidth cutoff frequency will not be attenuated and will, therefore, pass through the PLL. For this reason, the lowest loop bandwidth value is desired in many circumstances. However, low loop bandwidth comes at a price because the lower the loop bandwidth of a PLL, the slower it is to respond to changes at its input. At some point, very slow response time becomes a significant problem.

Figure 17 shows the jitter transfer function for a second-order PLL with a loop bandwidth of 100 Hz. Notice that the attenuation at 100 Hz is 3 dB, which would be expected from a low-pass filter with a cutoff frequency of 100 Hz.

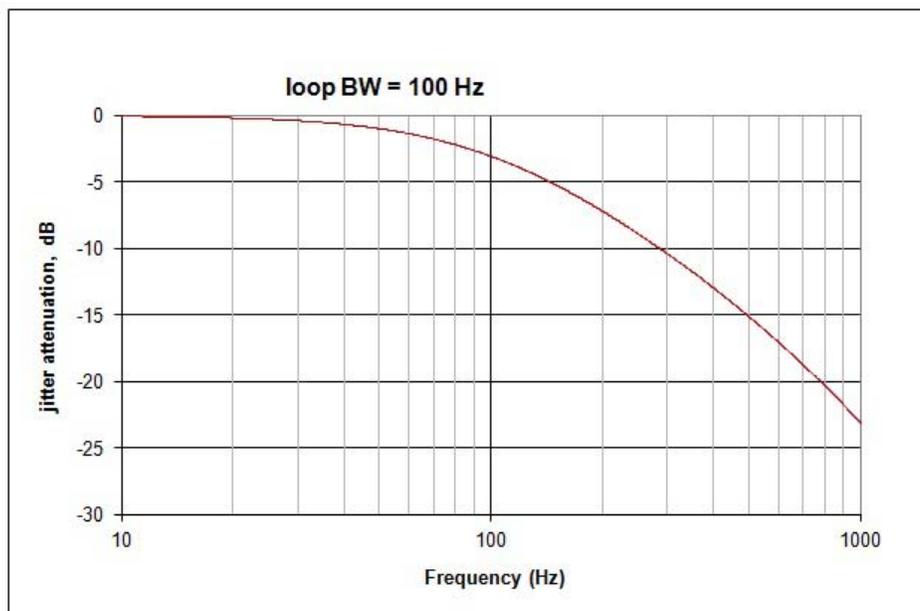
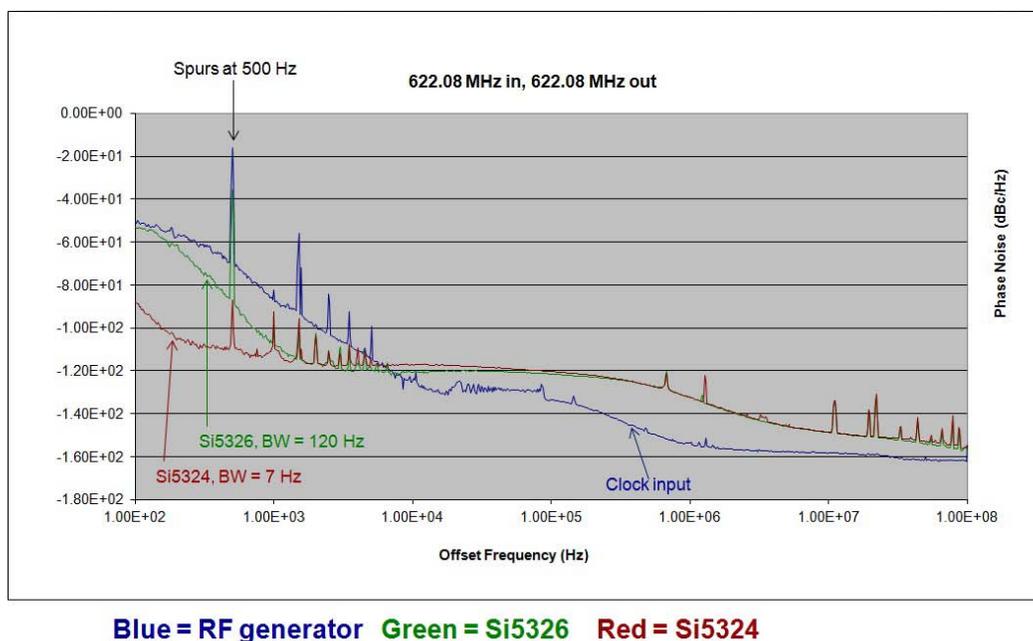


Figure 17. Jitter Attenuation vs. Jitter Offset Frequency

Next, we will examine phase noise plots created with an FM modulated clock input signal to see how the loop bandwidth affects the jitter attenuation. These plots show the jitter of a Rohde and Schwarz SML 03 RF generator with sine wave FM modulation that is connected to the clock inputs of an Si5324 PLL with a loop bandwidth of 7 Hz and an Si5326 PLL with a loop bandwidth of 120 Hz. The output jitter of the RF generator and the two PLLs are superimposed upon one another. For the “no modulation” case, see Figure 16. The phase noise plot in Figure 18 has FM modulation at 500 Hz, and the plot in Figure 19 has FM modulation at 10 kHz.

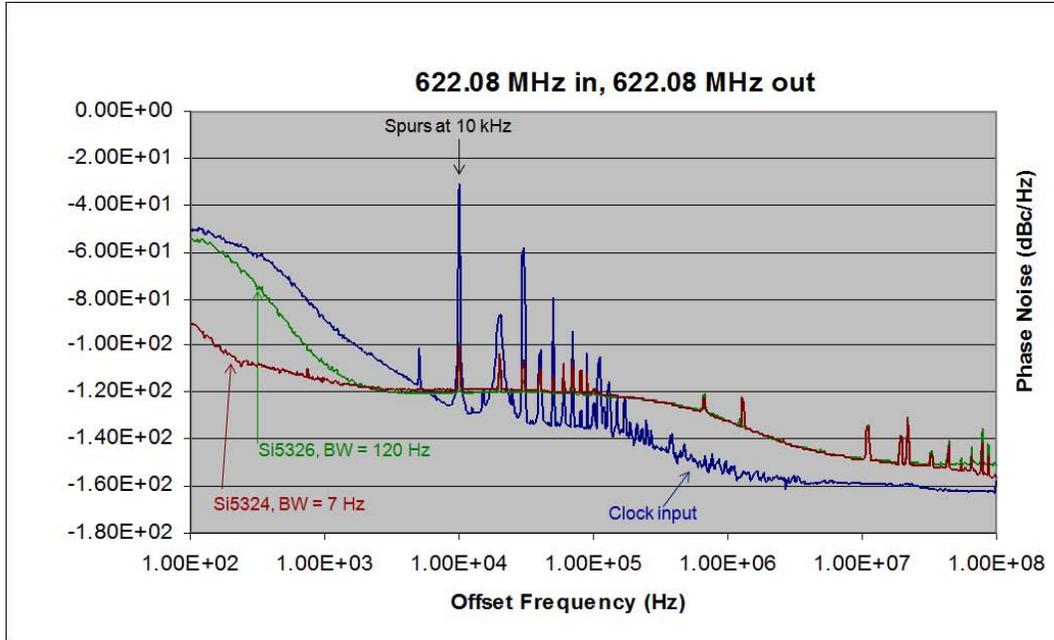


Blue = RF generator Green = Si5326 Red = Si5324

Figure 18. Jitter Attenuation at 500 Hz Offset

AN687

Notice that the prominent spur at 500 Hz is only somewhat attenuated by the Si5326 (with loop bandwidth = 120 Hz) but is significantly attenuated by the Si5324 (with loop bandwidth = 7 Hz). The increase in jitter attenuation at low offset frequencies is due to the lower loop bandwidth. For the plot in Figure 19 with jitter at 10 kHz, the jitter attenuation of the Si5324 and the Si5326 is roughly the same because 10 kHz is much larger than both of their loop bandwidth values.



Blue = RF generator Green = Si5326 Red = Si5324

Figure 19. Jitter Attenuation at 10 kHz Offset

Table 3 shows the RMS jitter for each of the cases. The integration band for the 500 Hz case was changed from 100 Hz to 50 MHz so that it would include the modulation frequency of 500 Hz. Notice that, although the Si5326 jitter at 500 Hz is much higher than the Si5324 jitter at 500 Hz, the two devices have very similar values at a 10 kHz offset.

Table 3. Jitter Values for Phase Noise Plots, (Picoseconds, RMS)

Jitter Frequency	Source Jitter	Si5326, 120 Hz	Si5324 , 7 Hz	Integration Band
0	1.18	0.293	0.281	500 Hz to 50 MHz
500 Hz	165	18.6	0.295	100 Hz to 50 MHz
10 kHz	139	0.302	0.304	500 Hz to 50 MHz

The jitter transfer curves for the Si5324 and the Si5326 are shown in Figure 20, along with the jitter transfer of an Si5338, which is a wideband PLL with a loop bandwidth of 1 MHz. As such, the Si5338 will attenuate very little jitter. As a consequence, if a wideband PLL (i.e., a PLL with a very large loop bandwidth) needs to have low output jitter, it must be driven by a low jitter clock source.

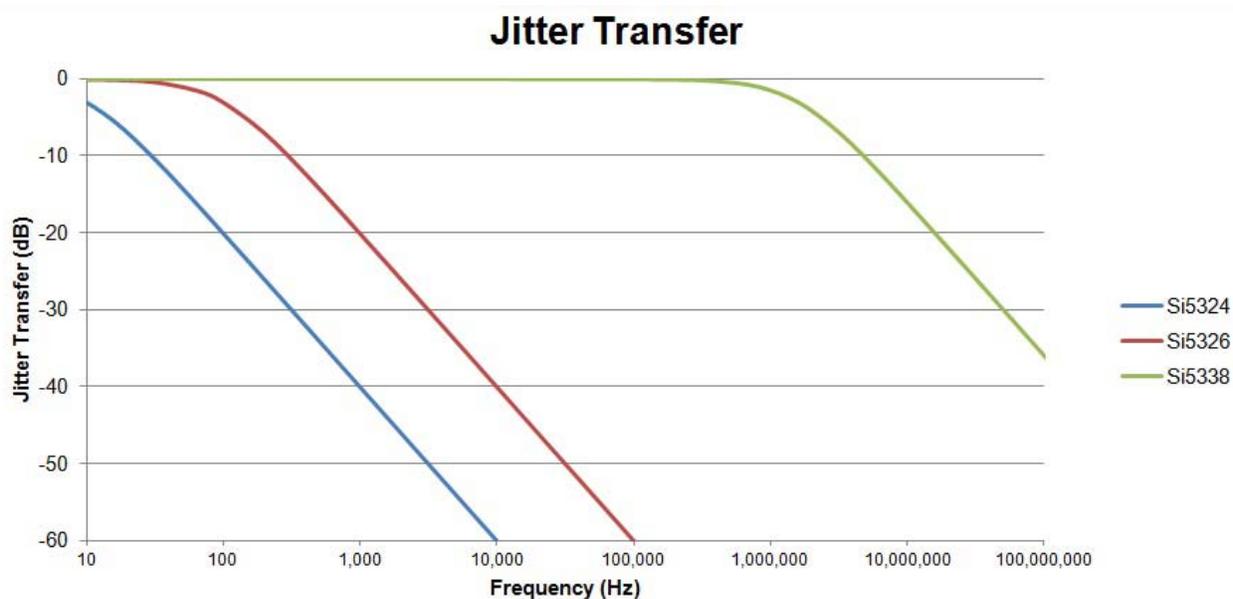


Figure 20. Jitter Transfer Function for Narrowband and Wideband PLLs

It is important to understand that the selection of an appropriate loop bandwidth value is application-dependent. If the clock source already has low jitter, then a higher low bandwidth might be suitable. However, if the input clock has significant jitter, a lower loop bandwidth should be considered. In these situations, the system-level jitter requirements need to be considered. Some networking standards establish a lower limit to the jitter offset frequency that needs to be taken into account, and, therefore, jitter below a certain offset frequency is either simply not considered or discounted. Examples of such standards include SONET, OTN, and Ethernet. However, some standards require loop bandwidth values as low as 0.1 Hz (e.g., SyncE ITU-T G.8262, option 2.) Video networks typically care about jitter that is down to 30 Hz because of video frame rates, which call for a loop bandwidth below 10 Hz.

Jitter tolerance is the amount of peak-to-peak jitter that can be applied to a PLL before it loses the ability to lock to or maintain lock with an incoming clock signal. The jitter tolerance is affected by both the loop bandwidth and the offset frequency. Clearly, the jitter tolerance for offset frequencies below the loop bandwidth of the PLL will be high because the PLL simply passes the jitter through with little attenuation. Similarly, a PLL with a low loop bandwidth will have less tolerance for jitter than a wideband PLL (i.e., one with a high loop bandwidth). A typical jitter tolerance curve is shown in Figure 21, with bandwidth referring to the loop bandwidth of the PLL. As the graph shows, there is a break in the jitter tolerance curve at the loop bandwidth.

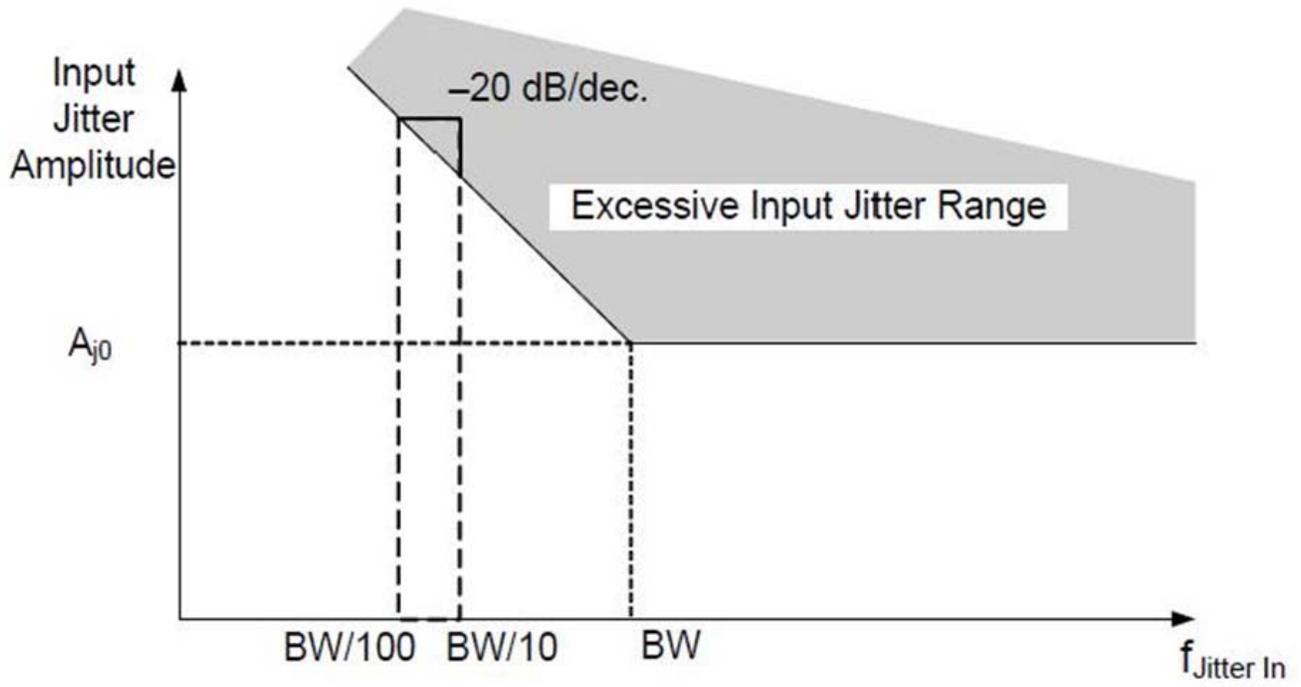


Figure 21. Typical Jitter Tolerance Curve

9. Clock Buffers

Once a low-jitter clock source has been created, it is often necessary to send this clock source to a number of different destinations. A fan-out buffer is usually used to perform this function, which leaves the designer with the question of how much jitter will be added by the fan-out buffer. In Figure 22, assume that there is a clock source with jitter, J_1 , that is driving a fan-out buffer with additive jitter of J_2 . What will the final output jitter be? Mathematically, the output jitter can be anywhere from $|J_1+J_2|$ to $|J_1 - J_2|$. In other words, it is possible for the two jitter sources to either constructively add or to subtract from one another. It all depends on whether J_1 and J_2 are correlated (or anti-correlated, as the case may be).

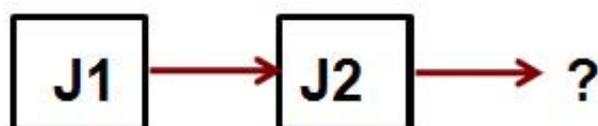


Figure 22. Clock Buffer Jitter

A simplifying assumption that is often made is that J_1 and J_2 are not correlated. As discussed earlier, this would be a very good assumption if both J_1 and J_2 were composed of purely random jitter. However, as the amount of correlated jitter increases, this becomes a less valid assumption. If J_1 and J_2 are truly uncorrelated, then the output jitter will be:

$$J_{out} = \sqrt{(J_1^2 + J_2^2)}$$

If the jitter at both the output and input of clock buffer J_2 is measured, then J_2 's additive jitter is defined as the root-mean-squared difference between J_2 's input and output jitter:

$$J_2 \text{ additive jitter} = \sqrt{(J_{out}^2 - J_1^2)}$$

This assumes that $J_{out} > J_1$, which will hold true if a reasonable percentage of the jitter involved is uncorrelated.

Consider the following example in Figure 23, which uses an Si5330 four-to-one fan-out buffer with an additive jitter specification of 0.150 picoseconds (integrated from 12 kHz to 20 MHz):

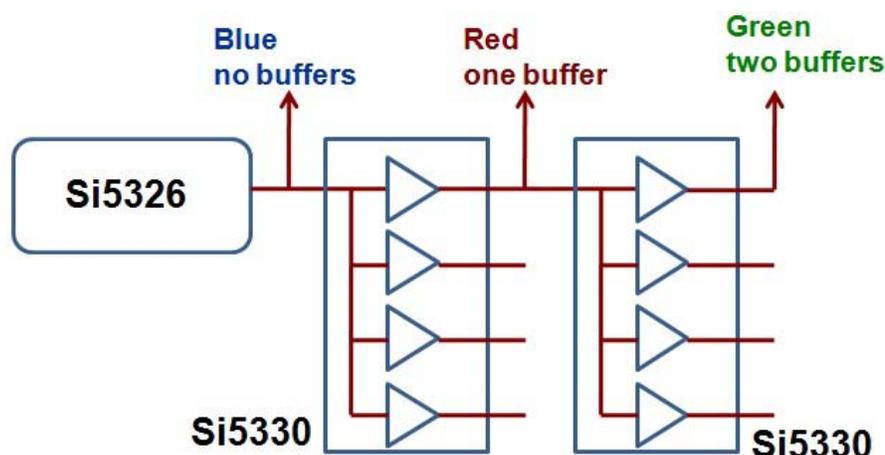


Figure 23. Fanout Buffer Measurement Configuration

AN687

The phase noise plots for the configuration in Figure 23 are shown in Figure 24.

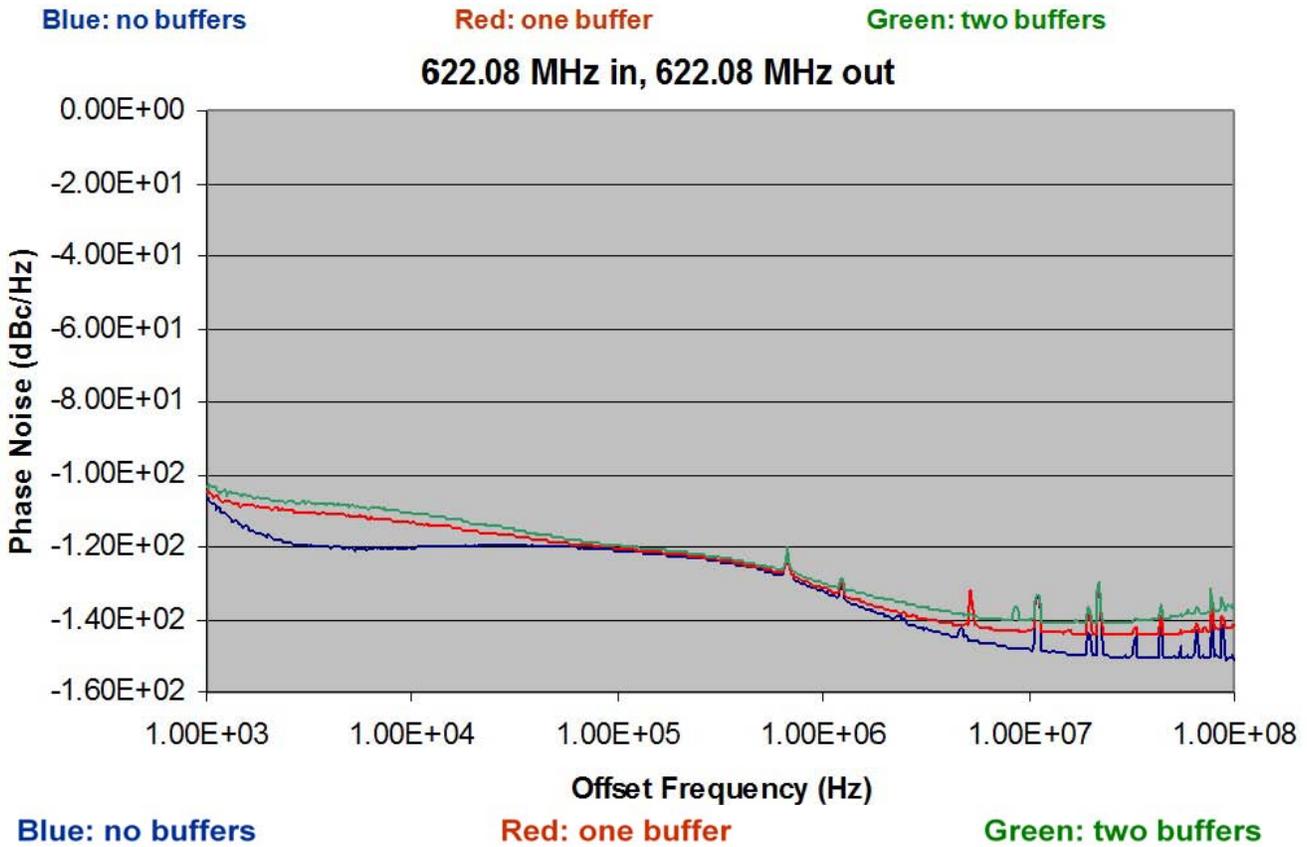


Figure 24. Fan-Out Buffer Phase Noise Plots

As expected, the lowest jitter is without any fan-out buffer, and adding a buffer increases the jitter. However, adding a second buffer does not incrementally increase the jitter as much as adding the first buffer. This is because of the square-root-of-sums-squared value that is associated with uncorrelated, random jitter. Table 4 has the corresponding RMS jitter values showing the measured and calculated values. The red values are calculated from RSS (root of sum squared), while the black values are all measured.

Table 4. Fanout Buffer RMS Jitter Values

Jitter Band	Rms Jitter	Rms Jitter	Rms Jitter	Rms Jitter	Rms Jitter
	No buffer	One buffer, measured	One buffer, RSS	Two buffers, measured	Two buffers, RSS
12 kHz to 20 MHz	258 fs	304 fs	298 fs	357 fs	334 fs

Notice that the RSS values are close to their corresponding measured values, but the RSS values are nonetheless both slightly less than the measured values. This indicates that some of the noise is indeed correlated. Here is the detailed calculation for one buffer, using the Si5330 clock specification of 150 fs:

$$J = \sqrt{(258^2 + 150^2)} = 298 \text{ femtoseconds}$$

10. References

1. Jitter, Noise and Signal Integrity at High-Speed, Mike Peng Li, Prentice Hall 2008
2. High Speed Signal Propagation Advanced Black Magic, Howard Johnson and Martin Graham, Prentice Hall 2003

CONTACT INFORMATION

Silicon Laboratories Inc.
400 West Cesar Chavez
Austin, TX 78701
Tel: 1+(512) 416-8500
Fax: 1+(512) 416-9669
Toll Free: 1+(877) 444-3032
Email: FMinfo@silabs.com
Internet: www.silabs.com

The information in this document is believed to be accurate in all respects at the time of publication but is subject to change without notice. Silicon Laboratories assumes no responsibility for errors and omissions, and disclaims responsibility for any consequences resulting from the use of information included herein. Additionally, Silicon Laboratories assumes no responsibility for the functioning of undescribed features or parameters. Silicon Laboratories reserves the right to make changes without further notice. Silicon Laboratories makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Silicon Laboratories assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. Silicon Laboratories products are not designed, intended, or authorized for use in applications intended to support or sustain life, or for any other application in which the failure of the Silicon Laboratories product could create a situation where personal injury or death may occur. Should Buyer purchase or use Silicon Laboratories products for any such unintended or unauthorized application, Buyer shall indemnify and hold Silicon Laboratories harmless against all claims and damages.

Silicon Laboratories and Silicon Labs are trademarks of Silicon Laboratories Inc.

Other products or brandnames mentioned herein are trademarks or registered trademarks of their respective holders.