Integrating The PLL System On A Chip

by Dean Banerjee, National Semiconductor

The ability to produce a set of many different frequencies is a critical need for virtually all wireless communication devices. One of the earlier methods of doing this was to use an array of crystals and select the one that would produce the desired output frequency. This method was very cumbersome and expensive. It also had the disadvantage that crystals were not available at higher frequencies. The introduction of the PLL (Phased Locked Loop) system addresses both of these issues as it allows one to generate a very broad array of frequencies from a single crystal. It is also capable of producing higher frequencies than a crystal, which means that it is often used even if only a single frequency is desired.

There is a continuing drive to integrate more and more of the PLL system on silicon. With the exception of the crystal reference this has been done This TechNote discusses these building blocks and some of the issues that silicon vendors face integrating them on silicon, with special attention given to the VCO.

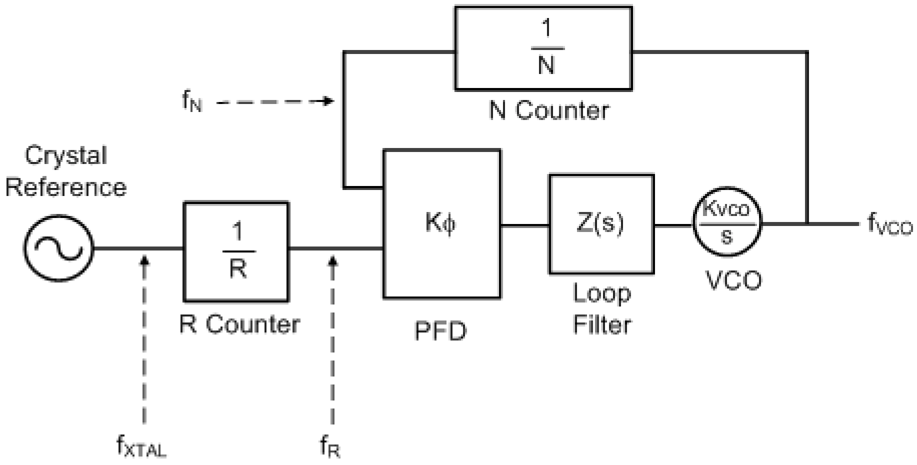

Fig. 1: Typical PLL System

PLL System Operation

The PLL System achieves its output frequency by using negative feedback in order to steer the frequency of a VCO (voltage-controlled oscillator) to the correct frequency. The VCO is a voltage-to-frequency converter that has an output frequency that can be changed based on a tuning voltage, but the frequency accuracy is not very good. Because of this a feedback loop is necessary, and in order to create this loop the crystal reference frequency is divided by an R counter in order to generate the frequency f_R -- also known as the comparison frequency. The VCO frequency is divided by an N counter in order to generate the frequency f_N. The phase/frequency detector converts the phase difference between the f_R and f_N signals into a voltage which in turn is converted by the charge pump into a correction current. If the frequencies are not equal, then there will be a non-zero correction current. For the purposes of simplicity, the term phase/frequency detector

(PFD) will be used to refer to both devices and will be considered to convert a phase error into a correction current. The loop filter is a low-pass filter that converts the correction current from the PFD into a voltage, which is the tuning voltage to adjust the VCO frequency. By using this feedback loop the PLL system achieves a tunable frequency at the output of the VCO with the same frequency accuracy as the crystal reference. For a given application the R counter is typically a fixed value and the N counter can be programmed to different values in order to achieve all the desired output frequencies. The output frequency, f_{VCO}, is derived from the crystal frequency, f_{XTAL} in accordance with:

$$f_{VCO} = f_{XTAL} \times \frac{N}{R}$$

The Crystal Reference

The crystal reference generates a very accurate (one part in a million, or better, are typical) and stable frequency. Capacitors, resistors, and inductors can be formed with silicon and bond wires, but any oscillator created using these as a resonant circuit would have poor frequency accuracy. It is for this reason that the crystal reference strongly resists integration into silicon although it is certainly possible to make modules with quartz crystals on them.

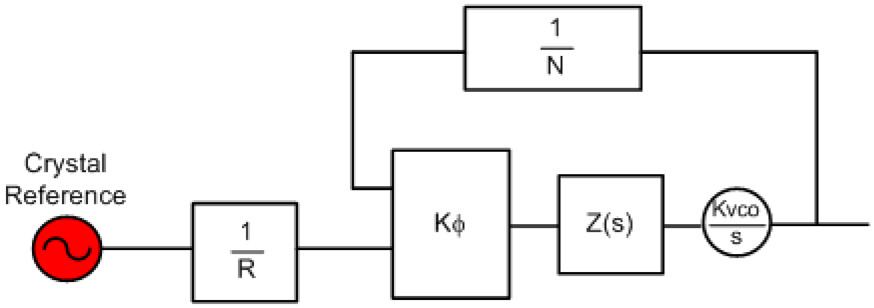

Fig. 2: Crystal Reference

The PFD, R Counter And N Counter

Silicon vendors have been integrating the PFD and dividers on-chip for a long time and selling them as *PLLs*, although it is understood that the crystal, loop filter, and VCO need to be supplied externally by the user.

Phase noise performance is typically very important, and the PFD is typically the dominant contributor. There are many design issues and tricks associated with the PFD, but many of these are proprietary, and beyond the scope of this article.

The R counter is typically a programmable counter that can be implemented as a series of flip-flops and the architecture is rather straightforward. In some cases, the R counter can even be omitted.

The architecture of the N counter is typically more complicated because it needs to be able to handle the potentially higher frequencies of the VCO. High-frequency processes tend to be more expensive and consume more current, yet are necessary in this case. In order to achieve a healthy compromise the N counter is often implemented with a set of counters. The first, called the prescaler, is implemented with a high-frequency process and divides the VCO frequency down to a manageable one. Once the frequency is divided down a lower-frequency process can be used for the rest of the chip. This divided frequency from the prescaler can now be run through a low-frequency counter in order to achieve the complete N counter value. This is a single modulus prescaler and has the limitation that the N counter value must be a multiple of the prescaler value. Because of this dual modulus architecture has become much more popular. This architecture uses two prescalers and two low-frequency counters that work together in order to allow the N counter to be changed in increments of one. However, there are limitations on how low the N counter value can go. The quadruple modulus architecture further expands on the dual modulus approach by combining four prescalers and three low-frequency counters that reduce this lower limit on the N counter. A lower N counter value is desirable because the N counter value multiplies the noise inside the loop bandwidth of the system, which is roughly the frequency where the closed-loop gain of the system starts to roll off. Once the prescaler limitation has been overcome the channel spacing is typically the factor that limits the N counter value. The fractional-N architecture overcomes this limitation by allowing the N counter to be a fraction. The improved resolution allows the comparison frequency (f_R) to be higher, which in turn allows a lower N counter value. This has the ultimate consequence of theoretically lower phase noise. Older implementations of fractional-N PLLs had fractional-N circuitry that added significant noise to the system. So although these fractional PLLs did ultimately have better phase noise it was not as much as theoretically expected due to the added noise of the fractional circuitry. More recent fractional PLLs use delta-sigma architectures that add much less noise and allow much higher fractional values. In general fractional PLLs have much better noise but the spurious performance can be better, or worse, depending the fractional chip used and the application.

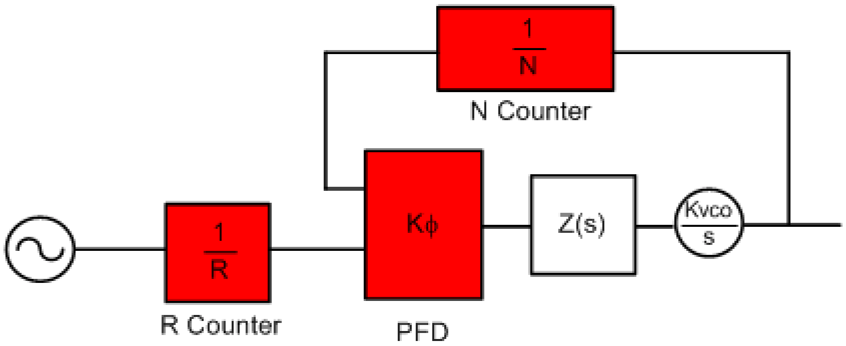

Fig. 3: PFD And Dividers

The Loop Filter

The loop filter is a low-pass filter typically implemented with external components. This is an area where the user of the PLL has the most influence on system performance. The loop filter can change the loop bandwidth -- a critical performance characteristic of the entire PLL system with significant impact on switching speed and noise performance. Fig. 4 shows a typical loop filter but not all the components are always necessary. C1, C2, and R2 are critical but the others are optional to improve spurious performance. An excellent reference discussing loop filter design is included at the end of this TechNote.

There are two big issues in integrating the loop filter. The first is that there is no single loop filter optimal for every application and integration limits the market to which the PLL chip can be sold. The LMX2502 and LMX2505 ICs integrate the entire loop filter, but they target very specific high-volume markets and cellular standards. The second issue is the implementation of capacitance in silicon. Capacitors and resistors can be implemented but tolerance is only fair and larger capacitors take up considerable die area. A limitation on how large the loop filter capacitors can be limits how narrow the loop bandwidth can be. A compromise to these issues is to partially integrate the loop filter while keeping some of the larger capacitors, like C2, external. The LMX2531 PLL/VCO uses this type of approach. Apart from simplifying the design and reducing the external component count integrated loop filters also have the advantage that they can filter spurs and crosstalk that may be much more difficult to filter with an external loop filter. This is especially true when the VCO is integrated on the chip.

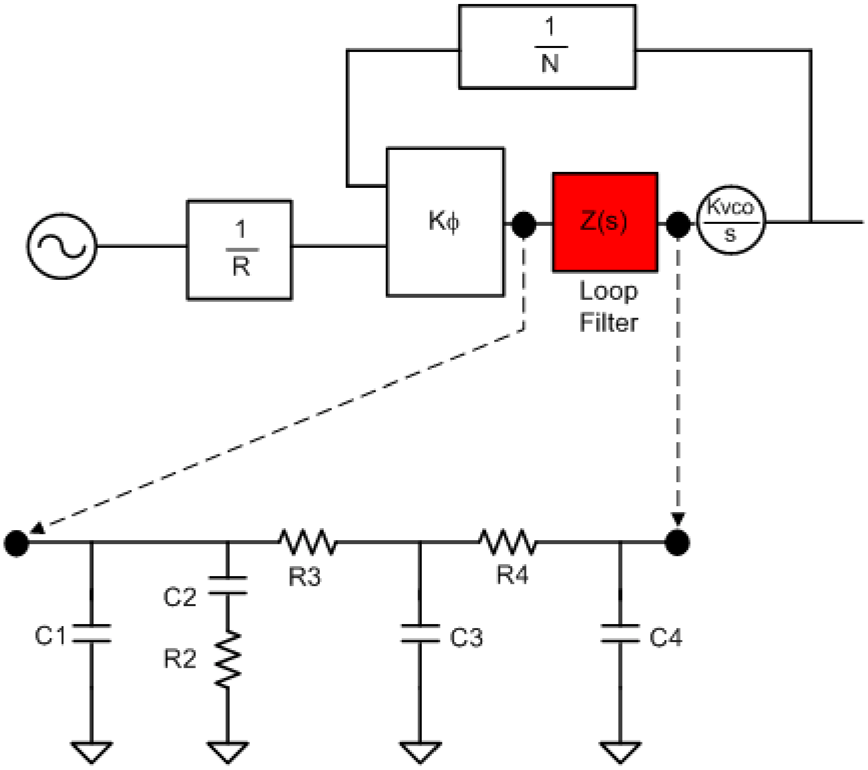

Fig. 4: Loop Filter

Designing The VCO

The VCO converts the tuning voltage from the loop filter into a frequency and typically limits the tuning range of the entire PLL system. There is a conversion gain associated with the VCO usually expressed in MHz/V. If this gain is made higher the VCO can cover a wider frequency range with the same tuning voltage, but the phase noise is normally degraded.

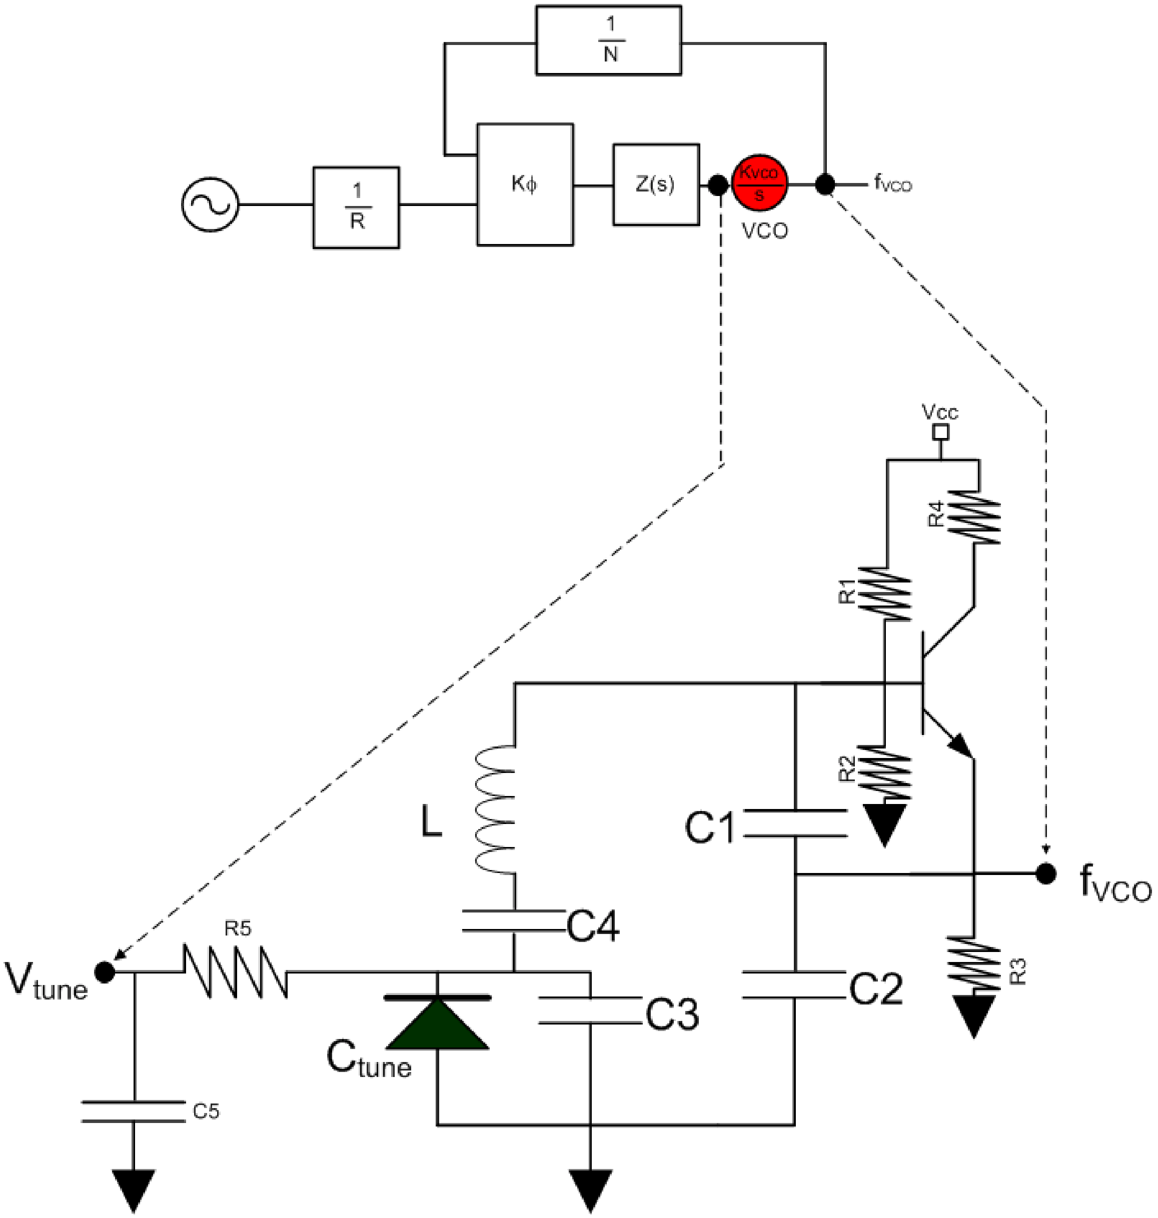

Fig. 5: VCO

A Simple VCO Design

The design of a VCO is actually quite involved. In theory all that is required is an inductor, a capacitor, and an amplifier. Fig. 5 shows a simple VCO. The inductor, L, and network of capacitors, C_{tune}, C1, C2, C3, and C4, form a resonant circuit called the tank. The equivalent capacitance of this circuit, C_{Total}, satisfies:

$$\frac{1}{C_{Total}} = \frac{1}{C1} + \frac{1}{C2} + \frac{1}{C3 + C_{tune}} + \frac{1}{C4}$$

Typically, parasitic capacitances decrease the oscillation frequency of the VCO, but neglecting these, the theoretical output frequency of the circuit is:

$$f = \frac{1}{2\pi\sqrt{L \bullet C_{Total}}}$$

C1 and C2 couple the energy from the tank circuit back to the active device. The ratio is important. If C2 is much larger than C1 then a large amount of energy is being sent back to the active device, which results in increased noise. However, if C2 is too small relative to C1 then there can be issues with start up times. C4 is a dc blocking capacitor to prevent the tuning voltage from fighting the bias of the transistor.

The capacitor, C_{tune}, allows the frequency to be tuned with a voltage. This is implemented with a diode that is reverse biased, called a varactor diode. As more voltage is applied the junction capacitance of this diode decreases, thus decreasing the total capacitance, and increasing the output frequency. The capacitor, C3, limits the tuning range but also improves the phase noise of the VCO, so there are trade-offs involved in the choice of this component.

Implementing The VCO On A Single Chip

As one can see even a simple VCO design has many components to tweak. Low noise VCO design is often treated as “black magic” and can be a very addictive and time consuming activity. Trace lengths, matching issues, non-ideal component behaviors, and harmonics are all considerations when trying to connect a discrete VCO to the high-frequency input pin of a PLL. In addition to these issues, cost and size considerations can make it even more motivating for one to consider a PLL chip that includes the VCO as well. In order for silicon manufacturers to do this it is necessary to implement the VCO tank inductance and capacitance.

Implementing The Tank Inductance For On-Chip VCOs

One way to deal with the VCO tank inductor is to make a pin that the user shunts with an external inductor or creates an inductance by using PCB board traces. This makes a flexible product and works well at lower frequencies, although it does place the burden on the user to supply the external inductance. At high frequencies (1 GHz and higher) the inductance of the bond wires (on the order of 1 nH) and traces on the board can be a significant contributor and often needs to be considered. Parts with an external inductor typically allow the user to short the pin to get the highest frequency by using only the bond wires themselves as an inductor. At higher frequencies parasitic inductances from tests sockets can make the maximum output frequency seem lower than it really is and that needs to be taken into consideration.

Another technique involves intentionally using the inductance of bond wires in order to create an tank inductor. Since bond wires and packaging are necessary to use the silicon die one could consider this on-chip inductor. If done correctly the output frequency of the VCO is not impacted by how the part is contacted to the board. There are many tricks that can be done to influence the bond wire inductance, but many of these techniques may be considered proprietary, so this will not be discussed in more detail. That being said the use of bond wires in order to form an inductor for the VCO is a very popular technique.

Implementing The Tank Capacitance For On-Chip VCOs

For an oscillator to work capacitance is required. The integrated capacitance on silicon can be split into three categories: varactor diode capacitance, switched capacitance, and fixed capacitance.

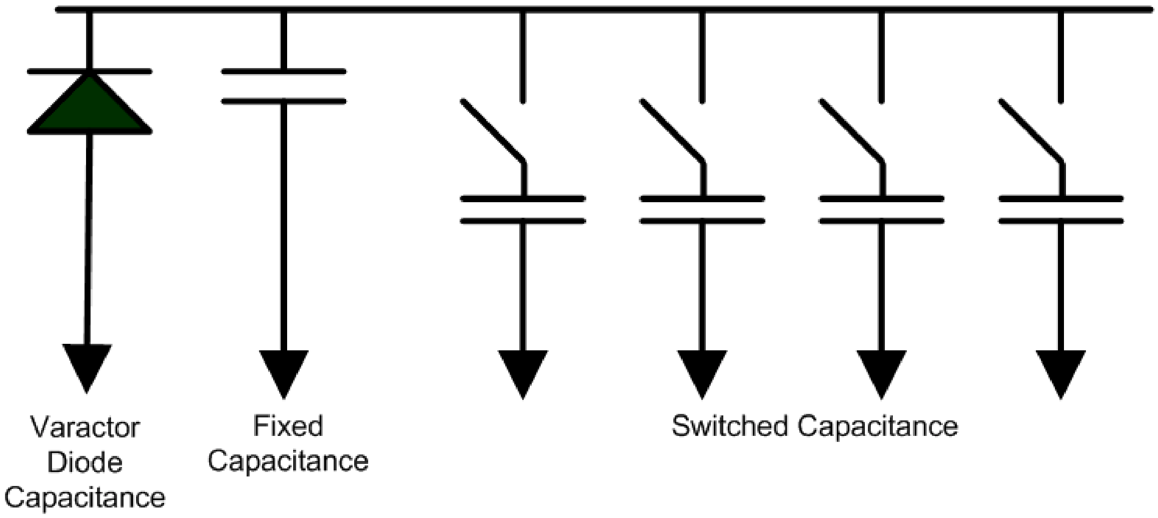

Fig. 6: More Detailed Look At Simple VCO

In Fig. 5, there is a fixed capacitance, C3, and a varactor diode capacitance, C_{tune}. On silicon these can be implemented in a similar fashion. However, a silicon VCO typically has a third capacitance, called the switched capacitance. This is actually a bank of capacitors that the chip can switch in and out. A simple method to implement this on silicon might be to require the user to switch the capacitors in and out with voltages at pins, but most modern on-chip VCOs include the logic circuit to automatically do this. If the PFD and counters are integrated with the VCO then the capacitor switching logic is much easier to implement, although it has been done for a standalone VCO based on tuning voltage. For the case where the PFD is integrated with the VCO, there is typically a frequency calibration routine that is run whenever the N counter value is programmed.

There are distinct advantages to using this switched capacitance. One is that the lock time is typically faster because the switching logic speeds up the system. The exception to this rule is when the system has very fast lock time which is on the order of the switching logic speed. Another advantage is that the VCO can be designed with a much lower VCO gain, which results in lower phase noise. Typically the best phase noise is at the lowest frequency because all the capacitors are switched in and the VCO gain is the least.

VCO Phase Noise Performance And Lesson’s Equation

Lesson’s equation shows up in many discussions concerning VCO phase noise. A paper by Jim Carlini referenced in the back of this TechNote presents the following expanded form of Lesson’s equation:

$$L(f) = 10 \bullet \log\left(\frac{1}{2} \bullet \left[\left(\frac{f_{vco}}{2 \bullet Q_L \bullet f} \right)^2 + 1 \right] \bullet \left[\frac{f_{1/f^3}}{f} + 1 \right] \bullet \left[\frac{F \bullet k \bullet T}{P} \right] + \frac{2 \bullet k \bullet T \bullet R_{var} \bullet Kvco^2}{f^2} \right)$$

where,

$L(f)$ = Phase noise in dBc/Hz

f = Offset frequency where phase noise is measured

F = Noise figure of active device

k = Boltzman’s constant = 1.380658×10^{-23} J/°K

T = Temperature in degrees Kelvin

P = RF power at input of active device

f_{vco} = Operating frequency of the VCO

Q_L = Loaded quality factor of the inductor = XL/RL

f_{1/f^3} = $1/f^3$ noise (flicker noise) corner frequency

R_{var} = Noise resistance of the varactor diode

K_{vco} = VCO gain in MHz/V

The most important factor in this equation is Q_L, the quality factor of the inductor, which is the imaginary reactance divided by the real resistance at the frequency of interest. Assuming that everything has been done to maximize this other factors also contribute. Note that improving the noise figure of the active device, increasing the output power, using varactors with lower noise resistance, and reducing the VCO gain all are steps that can be taken to improve the phase noise of the VCO. From this formula it can also be seen that reducing the VCO gain, K_{vco}, improves phase noise to an extent.

Phase Noise Enhancement Techniques For Silicon VCOs

The first technique for improving the phase noise is reducing the tuning gain of the VCO by splitting it into several different bands. This effectively reduces the tuning gain considerably while still allowing the VCO to tune over a broad range. When the switched capacitance is used it is important to keep the on-resistance of the switches low in order to have the highest possible Q factor for the tank circuit.

Another technique is to digitally tweak key parameters in to get the best phase noise. The advantage of this is that these algorithms can be run every time the frequency is changed to have it optimized over all frequencies. The methods that are used for this are proprietary but it can be said that they are effective. Consider trying to do this with a discrete VCO. One could tweak the capacitors, C1 and C2, in Fig. 5 to determine the optimum phase noise. Or perhaps you could tinker with biasing of the transistor, or other components. Substantial time could be spent trying to squeeze the last dBc of phase noise out of a discrete. Even if the optimal parameters are found they might not be optimal over component variations, frequency, voltage, or temperature. This combined with the lower tuning gain that is achieved with the switched capacitor bank can make silicon VCOs have potentially very low phase noise. These digital phase noise optimization techniques, as well as the reduced tuning gains, for silicon VCOs can make them very competitive in phase noise when compared to a discrete VCO.

Comparing On-Chip VCOs to Discrete VCOs

Every VCO is different but there are some trends that are common. In general the switched capacitor architecture for silicon VCOs allows them to achieve a very wide tuning range for the power supply voltages they are given. Also, if they come with a PLL then the overall reliability is improved because both components can be tested together. On-chip VCOs also have the advantage of small size, low profile, and cost effectiveness. They also often contain many other useful features, like a divide by 2 for the frequency output, lock detect, and variable power control.

When higher power supply voltages are available this tends to be an advantage for discrete VCOs, since there are maximum process voltage restrictions that often can prevent a silicon VCO from taking advantage of this. Although great for digital semiconductors the drive towards smaller process geometries and lower operating voltages can present challenges for low-noise silicon VCO design.

The LMX2531 family from National Semiconductor covers a wide range of frequencies and has low phase noise. These parts have the tank inductor integrated, provide an automatic algorithm that optimizes the phase noise, and has a partially-integrated and adjustable loop filter. The components R3, R4, C3, and C4 are integrated and the components C1, C2, and R2 are supplied by the user. This partially-integrated loop filter has the advantage that it has some on-chip components to filter out any potential crosstalk issues on the chip, yet it also has some external components for added flexibility.

One additional feature is that this part includes integrated LDOs. The power supply lines of a VCO are typically very sensitive to noise, but with these integrated LDOs there is much higher noise immunity on these pins.

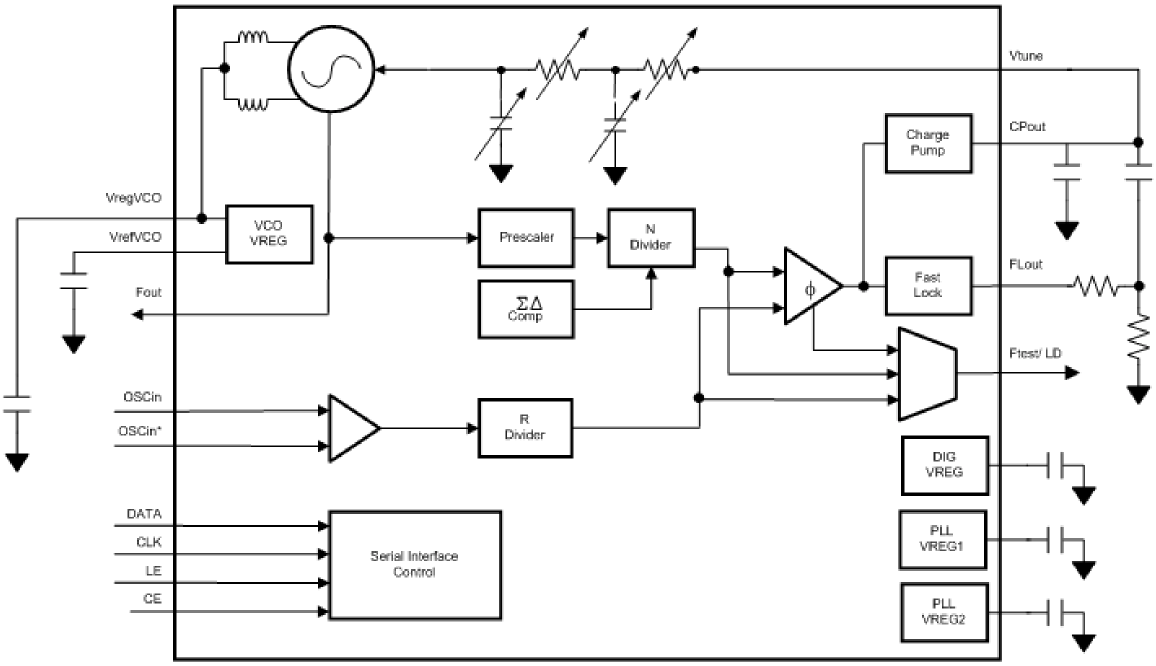

Fig. 7: LMX2531 Block Diagram

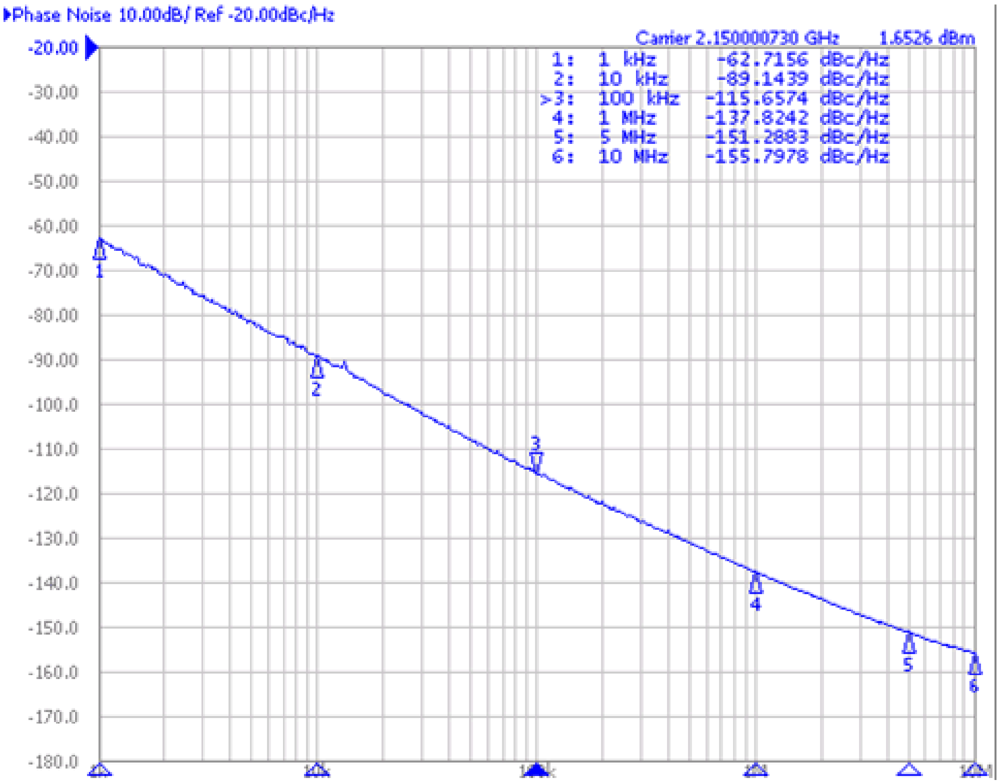

Fig. 8: Typical LMX2531 Phase Noise At 2150 MHz

Conclusion

With the exception of the crystal reference the PLL system is slowly becoming integrated. Aside from cost and size considerations integrating more of the PLL system allows the silicon vendors to specify more of the PLL performance, because they have more control of it and reliability is increased. Perhaps the greatest design challenge is getting a good quality VCO but this has now been achieved. Due to process voltage limitations it is hardest for silicon VCOs to compete with discrete VCOs that can use higher voltage supplies.

About The Author

Dean Banerjee is an applications engineer working for National Semiconductor and has been working with PLL synthesizer ICs since 1996. He has extensive experience supporting customers and has played a significant role in the creation of the EasyPLL tool at http://wireless.national.com He has written technical papers, the most significant being the book entitled, *PLL Performance, Simulation, and Design*, which covers many PLL design issues in great detail. Dean earned MS degrees in Electrical Engineering from Southern Illinois University and Applied Mathematics from the University of Illinois.

References

Banerjee, Dean, *PLL Performance, Simulation, and Design*, 4th Edition. This is available at http://wireless.national.com
Carlini, Jim, *A 2.45 GHz Low Cost*, *High Performance VCO*, Microwave Journal Technical feature http://www.ansoft.con/news/articles/MWJ.04.00.pdf

as published in...
analog ZONE