JOHN WILLISON / VICE-PRESIDENT OF ENGINEERING
STANFORD RESEARCH SYSTEMS, INC., 1290-D REAMWOOD AVE., SUNNYVALE, CA 94089;
(408) 744-9040, FAX: (408) 744-9049, E-MAIL: Info@thinkSRS.com, www.thinkSRS.com.

Presenting A Novel Synthesis Approach

Rational approximation frequency synthesis is a new technique for generating RF and microwave signals with extremely low phase noise and almost infinite frequency resolution.

Frequency-synthesis techniques provide stabilized RF and microwave signals to a wide range of defense-electronics systems, including electronic-warfare (EW) systems, signal-intelligence (SIGINT) receivers, and tactical radios. While a variety of analog and digital synthesis architectures have been applied to both voltage- and current-tuned oscillators, few new synthesis methods have appeared in recent years.

Until now, that is. A new approach developed by Stanford Research Systems (www.thinkSRS.com), known as Rational Approximation Frequency Synthesis, offers output signals with low phase noise that can be tuned with extremely fine resolution—and it can be implemented with relatively low cost. The technique has already been demonstrated in a commercial product, the company's model SG384 signal generator (see this issue's Cover Feature).

One of the most popular frequency-synthesis techniques involves the use of phase-lock-loop (PLL) devices to stabilize the phase of a higher-frequency oscillator to that of a lower-frequency reference source, such as a crystal oscillator. The classical PLL block diagram is shown in **Fig 1**.

A PLL synthesizer generates precise outputs locked to a reference frequency. As shown in **Fig 1**, the reference frequency, f_{REF}, is divided by the integer R. The voltage-controlled oscillator (VCO) output, f_{OUT}, is divided by the integer N. A phase detector compares the phase of the divided frequencies. The phase detector's output is passed through a lowpass filter and used to control the frequency of the VCO so that f_{OUT}/N is equal to f_{REF}/R, hence, $f_{OUT} = N(f_{REF}/R)$.

A numerical example will help to illustrate the operation and design tradeoffs of a PLL synthesizer. Suppose f_{REF} = 10 MHz and R = 1000. If N = 10,000, then the output frequency, fOUT = $N(f_{REF}/R)$ = 100 MHz. As N is changed from 10,000 to 10,001 to 10,002, f_{OUT} will change from 100.00 MHz to 100.01 MHz to 100.02 MHz. This PLL synthesizer has a phase comparison frequency, and a channel spacing, of fREF/R = 10 kHz.

Figure 2 shows a typical phase-noise plot for a 100 MHz PLL synthesizer. The phase-noise plot shows the noise power in a 1-Hz sideband as a function of frequency offset from the carrier. There are three dominant sources of phase noise: the reference oscillator, the phase detector (and its dividers), and the VCO.

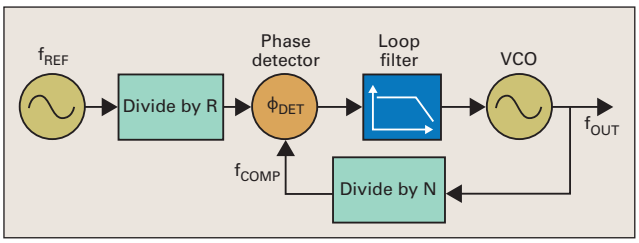

1. This block diagram shows a classical "integer-N" phase-lock-loop (PLL) frequency synthesizer.

The frequency reference dominates the noise close to the carrier but falls off quickly at large offsets from the carrier frequency. The phase-detector noise floor is relatively flat versus frequency but decreases with increasing phase comparison frequency. In fact, the phase-detector noise decreases by about 10 dB/decade; hence, it is about 30 dB lower for phase comparison frequencies at 10 MHz versus those at 10 kHz. Finally, the VCO phase noise will dominate at offset frequencies beyond the loop bandwidth. A high phase-comparison frequency, and low R and N divisors, are required for a low-phase-noise design.

In a properly designed PLL, the output noise tracks the reference at low offsets, matches the phase-detector noise at intermediate offsets, and is equal to the VCO noise at offsets beyond the PLL loop bandwidth. Careful attention to the loop filter design is also required to achieve the total noise characteristic shown in **Fig. 2**.

In addition to broadband noise there will be discrete spurious frequencies in the phase noise spectrum. A dominant spur is often seen at the phase comparison frequency. It is easier to reduce this spurious product in the PLL loop filter when the phase comparison frequency is high.

Frequency resolution of 10 kHz is adequate in many communications applications, even for systems with narrowly spaced channels, but a higher resolution is needed in test and measurement applications. The simplest way to increase the frequency resolution is to increase the value of the R divider.

In the above example, if R increases from 1000 to 10,000, the frequency resolution (channel spacing) improves from 10 kHz to 1 kHz. However, there are several serious drawbacks to this strategy. As the R divider is increased, the phase comparison frequency is decreased, leading to higher phase-detector noise, a reduction in the loop bandwidth, and increased settling times. Increasing R will

achieve high frequency resolution at the cost of a noisy output that takes a long time to settle.

Another strategy to increase resolution without decreasing the phase comparison frequency is to use a fractional-N synthesi-zer. In these synthesizers, the value of N is modulated so that it can have a noninteger value. If N averages a value of 10,000.1, then the output frequency, $f_{OUT} = N(f_{REF}/R) = 100.001$ MHz. The frequency resolution has been improved to 1 kHz.

However, modulating the N value creates spurious products in the VCO output. Dithering techniques can spread most of the spurious energy out as broadband noise, but the remaining noise and spurious products can be problematic in some applications.

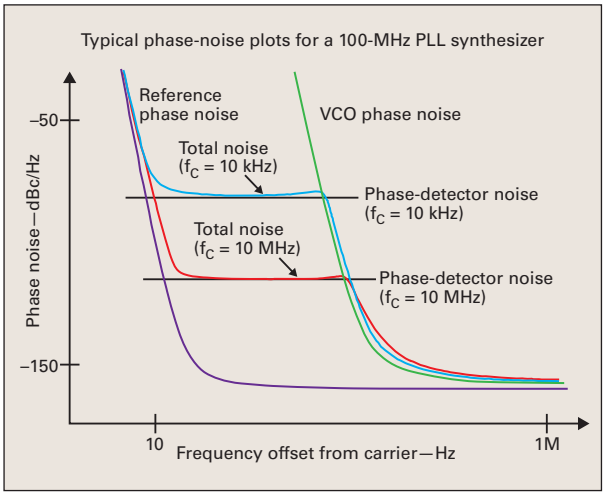

2. This plot shows a typical phase-noise spectrum for a 100-MHz PLL frequency synthesizer.

Yet another workaround to the tradeoff between high resolution and reduced phase comparison frequency (and thus, higher phase noise) is to use an oscillator based on a yttrium-iron-garnet (YIG) resonator. YIG oscillators provide wide tuning ranges with low-noise output signals due to the high quality factor (Q) of their resonator—a YIG sphere tuned by a magnetic field. They have numerous drawbacks, however, including limited tuning speed, high power consumption, relatively high cost, and susceptibility to external magnetic fields.

A new approach to synthesizer design provides high frequency resolution, fast settling, and low phase noise. As noted earlier, it is called Rational Approximation Frequency Synthesis and is based on rational numbers, which are the ratio of two integers. The approach has been overlooked as it relies on some surprising results of rather quirky arithmetic, which abandons neat channel spacing in exchange for a much better performing PLL synthesizer.

To demonstrate the approach, a numerical example might be helpful. Using a target frequency of 132.86 MHz for a PLL synthesizer, the question is how to tune to that frequency. It can be done by setting R to 1000 and N to 13,286. With f_{REF} = 10 MHz, $f_{OUT} = N(f_{REF}/R) = 132.86$ MHz. The phase comparison frequency is 10 kHz and the PLL loop bandwidth, which is typically 1/20 of the phase-comparison frequency, would be only about 500 Hz.

But there is another way to synthesizer a frequency of 132.86 MHz, or at least fairly close to it. The procedure starts by setting R to 7 and N to 93. Then, $f_{OUT} = N(f_{REF}/R) = 132.857142$ MHz, which is only 21.5 ppm below from the target frequency (and hence the term "rational approximation"). Of course, increasing the reference frequency by +21.5 ppm will produce the target frequency exactly.

But the positive benefit of this approach is clear: The phase-comparison frequency is 10 MHz/7 MHz = 1.42 MHz, or 142 times higher than that of a classical PLL with 10-kHz channel spacing. This allows a PLL bandwidth which is also 142 times wider. The higher comparison frequency of this PLL will allow wideband modulation, provide faster settling, lower phase noise, and results in an easily removed (by filtering) reference spurious product which is 1.42 MHz away from the carrier.

Is this approach applicable, which is to say, can sufficiently small values for R and N always be found to produce an output frequency that is close enough to a desired frequency? And is there a way to find the smallest possible values for R and N? Can the output frequency be made to match precisely (not just "close to") the desired frequency? As a real-world example will show, the answers to all of these questions is "yes."

Figure 3 shows a PLL synthesizer that can generate outputs anywhere in the octave between 2 and 4 GHz. Lower frequencies are easily generated by binary division of these outputs. This example uses a model ADF4108 dual-modulus PLL frequency synthesizer integrated circuit (IC) from Analog Devices (www.analog.com).

A dual modulus N counter is a high-speed divider which divides by a prescaler value, P, or by P + 1 under the control of two registers named A and B. The dual-modulus N-divider adds some numerological quirkiness since there are restrictions on the allowed values for A and B, as detailed in **Fig. 3**. The ADF4108 also requires that the phase-comparison frequency be less than 104 MHz. The reference frequency input in this example is 200 MHz.

A curious aspect of Rational Approximation Frequency Syn-thesis is that it is not obvious how to choose the values for R and N. Mathematical techniques exist for rational fraction approximation; brute enumeration of the many possibilities may also be applied. For example, R and N can be found by starting with the lowest allowed value for R and testing to see if there is an allowed value for N which gives a result, $f_{OUT} = N(f_{REF}/R)$ which is within some error band (say, ±100 ppm) of the desired frequency. Luckily, these computational requirements are modest. The required calculations can be performed on a typical microcontroller in under a millisecond.

The largest phase-comparison frequencies are achieved when there are many numeric choices available to improve the chance that a particular ratio of integers can be found which will be within the error band of the desired result. This is done three ways. First, allow a large error band. (An error band of ±100 ppm is typical because a fundamental mode crystal oscillator, which is used to clean-up the reference source, can be tuned over ±100 ppm.) Second, use a high-frequency reference oscillator. Third, provide a second reference, detuned slightly from the first, to provide additional numeric choices.

To ascertain how well Rational Approximation Frequency Synthesis works for

the example in **Fig. 3**, a computer program was written to compute the R and N values for 10,000 random frequencies in the octave band between 2 and 4 GHz. Using a single reference source at 200 MHz and an allowed error band of ±100 ppm, the average phase-comparison frequency was 9.79 MHz and the worst-case phase comparison frequency was 400 kHz.

When a second reference frequency was available (at 201.6 MHz, as determined by trial and error while searching for the highest worst-case phase comparison frequency), the average phase comparison frequency increased to 12.94 MHz and the worse-case phase comparison frequency increased to 2.35 MHz (a six-fold increase).

Rational Approximation Frequency Synthesis provides fast-settling, low-phase-noise, and spurious-free outputs, but with a troubling "error band" of typically ±100 ppm. To eliminate this error, it will be necessary to provide a low-noise reference to the PLL, which is tunable with high resolution over a range of ±100 ppm.

A 48-b direct-digital-synthesizer (DDS) frequency source was used to as the tunable reference to fix the frequency errors described above. The DDS cannot be used as the reference to the RF PLL directly, because the DDS spurious outputs, when multiplied up the required frequency, would greatly exceed the target specifications. Several measures were taken to improve the DDS reference to the RF PLL:

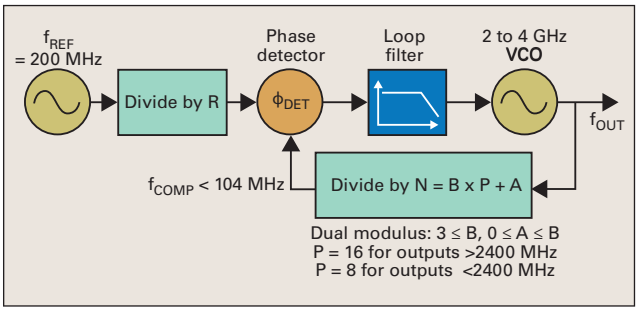

3. This block diagram shows an example of a Rational Approximation Frequency Synthesizer based on a model ADF4108 dual-modulus PLL IC from Analog Devices.

1. The frequency resolution was extended to 64-b by controlling the duty cycle of the DDS least-significant bit (LSB).

2. The DDS spurious content was reduced by dithering the DDS frequency with a pseudorandom binary sequence with a repetition rate of 100 kHz.

3. The DDS's spurious signal products were filtered by phase locking a voltage-controlled crystal oscillator (VCXO) to the DDS within a 100-Hz bandwidth.

4. The VCXO, which can be pulled over ±100 ppm with a frequency resolution of 64-b, was used as the reference to the RF PLL.

This new approach to frequency synthesis allows for operation at much higher phase comparison rates than the classical PLL approach—allowing for a wider PLL bandwidth to provide wideband modulation, faster settling, lower phase noise, and low spur outputs with virtually infinite frequency resolution. **DE**

TRACKING PLL HISTORY

PLL-based frequency synthesizers are currently supplied by a wide range of high-frequency companies that typically also produce oscillators. These PLL synthesizers are constructed around a number of different oscillator approaches—including current-tuned oscillators incorporating yttrium-iron-garnet (YIG) resonators and voltage-controlled oscillators (VCO)—that typically employ a varactor diode with capacitance that changes as a function of applied voltage to affect the tuning range of the oscillator.

The two oscillator approaches different somewhat in spectral purity, with YIG oscillators typically capable of excellent phase noise performance compared to VCOs. But they are also typically more expensive than VCOs and consume more power, with considerably slower frequency settling time (tuning speed). For designing a frequency synthesizer with a PLL, however, either typical of oscillator can be used.

The concept of a PLL actually dates as far back as the 1930s. PLLs were developed as the basis for a receiver architecture that would serve as an alternative to the superheterodyne receiver created by Edwin Armstrong. The PLL-based approach, which was known as a homodyne receiver, employed direct frequency conversion of received signals. The technique tuned a local oscillator (LO) to a desired frequency and then multiplied it with the received input signal.

It was a simpler approach to the superheterdyne receiver, which might incorporate several mixers and frequency downconversion stages before translating input signals to an intermediate-frequency (IF) range where the modulation information could be extracted or demodulated. There was at least one problem with the homodyne approach, however, and that was the lack of frequency stability in the LO sources of that time.

Any form of available tuned oscillator at that time would fluctuate in frequency over any extended time period and with any changes in temperature. It could not be tuned precisely enough to provide reliable reception and demodulation of received signals. But by using a correction loop, in which the phase of the tunable oscillator could be tied back to the phase of a more-stable reference oscillator, usually a lower-frequency, fixed-frequency crystal oscillator, the frequency variations of the tunable oscillator could be sufficiently stabilized.

Over the years, PLLs have been developed with analog and digital components. Essentially, a PLL uses a feedback loop and loop filter around a tunable frequency oscillator, along with some form of phase detector to provide a feedback signal to the tunable oscillator. In an analog PLL, an analog multiplier is used as the phase detector. In a digital PLL, a logic gate or a phase frequency detector (PFD) is used as the phase detector. The loop filter, which is typically a lowpass filter (LPF), can be passive or active.

One of the inputs to the phase detector is from the tunable frequency oscillator, the other is from a reference oscillator. In a typical negative feedback configuration, the phase difference between the two inputs is kept constant, stabilizing the tunable oscillator. Although a reference oscillator supplies the stable input signal for phase comparison, it is not considered part of a PLL, although it is a part of a PLL-stabilized frequency synthesizer.

RF Signal Generators

There's always room for improvement

The Agilent N5181A-503 is an excellent 100 kHz to 3 GHz signal generator.

$8,749
100 kHz to 3 GHz ...

Still, we saw room for improvement ...

- Extending the frequency range to 4 GHz
- Going all the way down to DC
- Adding a low-noise OCXO timebase
- Making AM, FM, ΦM & PM all standard
- And, cutting the price (about in half)

Introducing the SG384 RF Signal Generator ... only from SRS

DC to 4 GHz ...
$4,600

You'll even like the fine print: **US List Price $4,600. Standard features:** DC to 4.05 GHz with 1 µHz resolution. Output power from +13 dBm to –110 dBm. Phase noise of –116 dBc/Hz at 20 kHz offset from 1 GHz. A 1 s Allan variance of 1×10^{-11}. 10 MHz timebase input and output. AM, FM, ΦM, pulse modulation and sweeps from internal or external sources. Ethernet, GPIB, and RS-232 interfaces. **Option 1:** Differential clock outputs on SMAs with 35 ps transition times ($750). **Option 2:** Rear-panel SMA output for 4.05 GHz to 8.10 GHz ($750). **Option 3:** I/Q modulator with external BNC inputs ($750). **Option 4:** Rubidium timebase for 0.001 ppm/yr aging ($1500). Please visit **www.thinkSRS.com** for complete specifications.