Fractional N synthesizers

Analyzing fractional N synthesizers and their ability to reduce phase noise, improve loop speed and reduce reference spur levels.

By Gary Appel

The fractional N synthesizer has been credited with the ability to decrease phase noise, provide increased loop speed for a given step size, and provide reduced reference spur levels.

This article will examine the waveforms produced by charge pumps when operating in the fractional N mode. It will use these waveforms to calculate approximate levels for the spurious sidebands created by the fractional N technique. While this article will not derive the loop equations for a phase-locked loop, it will go through some of the calculations required to get a simple loop up and running.

Finally, it will examine the multiple-modulus divider provided in the chips and develop a concise algorithm for determining a valid set of divider numbers for the three- and four-modulus prescalers. While this analysis uses Phillips devices for presentation, the developed theory can be applied to all fractional N synthesizers.

The fractional N loop

In the traditional synthesizer, the voltage-controlled oscillator (VCO) frequency is at some integer multiple of the divided down reference frequency at the phase comparator (comparison frequency, or f_{comp}). Frequency steps smaller than f_{comp} are not possible. As a result, the comparison frequency can be no higher, in frequency, than the desired channel spacing, or step size. If the VCO frequency is high, and the step size is small, this results in a large division ratio in the VCO path.

For example, with a 500 MHz VCO, and a channel spacing of 50 kHz, the VCO frequency will have to be divided by 10,000 to obtain the comparison frequency of 50 kHz. This will result in an increase of the phase detector noise of $20_{log}10000$, or 80 dB. And the 50 kHz channel spacing will require a narrow loop bandwidth, to control the amplitude of the reference spurs. The narrowband loop will exhibit a slow transient response. The fractional N loop allows us to lock the VCO to a frequency that is a fractional multiple of f_{comp}. This, in turn, allows use of a comparison frequency larger than the step size. The overall division ratio can be reduced, lowering the contribution of the phase detector noise. Because the reference spurs are now at a higher frequency, the loop can be sped up while retaining the same rejection of the reference spurs.

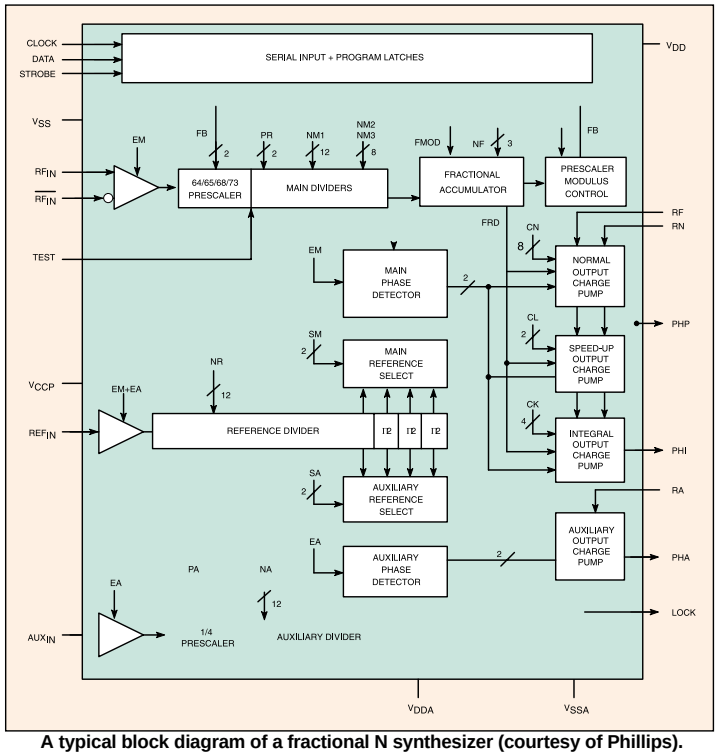

A typical block diagram of a fractional N synthesizer (courtesy of Phillips).

As an example, consider again the 500 MHz VCO, with the channel spacing of 50 kHz. The required division ratio for the standard loop, using a 50 kHz comparison frequency, is 10,000. To set the VCO to 500.05 MHz would require the next integer division ratio, 10,001. Now, increase the comparison frequency to 100 kHz, reducing the division ratio to 5,000. To set the VCO to 500.05 MHz would require a division ratio of 5000.5. The 0.5 is clearly unobtainable using a standard digital counter. The fractional N loop gets around this problem by alternating the division ratio between 5,000 and 5,001. The average division ratio is then precisely 5,000.5, just what we need to set the VCO on frequency.

A closer look

In a normal synthesizer, while locked, the phase detector produces a narrow pulse, which is required to keep the VCO on frequency. Each pulse from the phase detector should be identical to all the other pulses.

In the fractional N loop described above, the VCO is never quite on frequency. That is, it is never an exact integer multiple of the comparison frequency. In one cycle of the comparison frequency the VCO frequency will appear to be high by half the comparison frequency. In the next cycle, the VCO will appear to be low by an equal amount. The loop will therefore attempt to ramp the VCO frequency up, then down in alternate cycles of the phase detector, creating a spur at half the comparison frequency. Because this spur occurs at a fraction of the comparison frequency[1], it is known as a fractional spur. It will be shown later that the chips used to develop this technique contain some additional circuitry to mini-

mize the level of the fractional spurs.

The above example represents a fractional modulus of two. It is possible to choose either a fractional part of zero, when the divider is an integer, or a fractional part of one, when alternating the divide ratio between two adjacent integers. This allows selection of a either a fractional modulus of 5 or 8. This makes it possible to set the VCO divider to fractional multiples of 1/5 or 1/8 of the comparison frequency. This is accomplished by alternating the division ratio between two adjacent integers in a cycle that repeats over a a period of five phase detector cycles, or eight phase detector cycles. (With a fractional modulus of 8, the cycle may repeat in two or four phase detector cycles.) A fractional modulus of 5 presents a choice of a fractional part of 0/5 to 4/5. A fractional modulus of 8 offers a choice of a fractional part of 0/8 to 7/8.

It is necessary to define two other terms that will be used in the following discussion. First, the time from one phase detector pulse to the next will be addressed as a comparison cycle. The time required for the phase detector waveform to repeat will be called a fractional cycle. For example, if a fractional modulus of 5 is used, the fractional cycle will normally contain five comparison cycles.

The fractional N current waveform

Return to the example using the fractional modulus of 2 along with a fractional part of 1. The divider and charge pump waveforms for this example are illustrated in Figure 1. Suppose that as the fractional cycle starts (beginning with the division ratio of 5,000), the reference divider and the VCO divider both transition at exactly the same time. At this point the two dividers are perfectly aligned in phase. It has already been determined that the VCO is running just a bit fast, so that the next VCO divider transition occurs before the reference divider transition by exactly half a VCO cycle. The offset in time results in an error pulse from the phase detector, with a width equal to half a cycle of the VCO. At the third VCO divider transition (which terminates the fractional cycle), the VCO has moved ahead by another half cycle, or a total of one VCO cycle. But because the VCO divider is now dividing by 5,001, the two transitions again line up. Think of this extra divide as swallowing a VCO cycle. Figure 1 shows that this

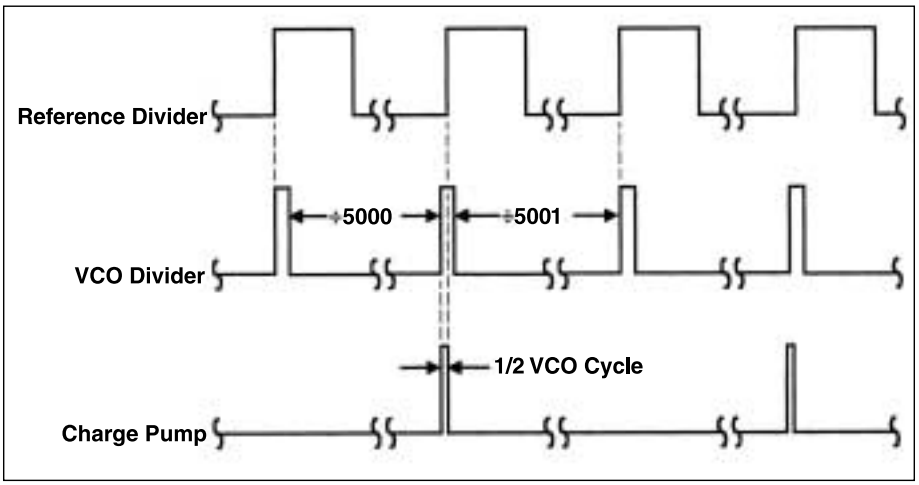

Figure 1. Divider and charge pump waveforms.

fractional N synthesizer will deliver error pulses from the phase detector on alternate comparison cycles.

This waveform would soon drive the loop out of lock, as the continuous pulses drive the tuning line up to the positive rail. It will later be shown that another charge pump in the chips provides a compensating pulse to remove the charge delivered by the phase detector charge pump, preventing the continuous rise of the tuning voltage.

If a fractional part of 1/8 is selected, the error pulse from the phase detector will start with a width of zero. Each subsequent transition of the VCO divider will result in a growth of the width of the error pulse equal to 1/8 of a VCO cycle, to a maximum width of 7/8 of a VCO cycle. On the eighth cycle of the VCO divider, the divide number is increased by one, so that the VCO divider transition and reference divider transition are again coincident, and the width of the error pulse is again zero. This waveform is shown in Figure 2.

If a fractional part of 3/8 is chosen, an error width is accumulated three times as fast. As the cycle begins, the transitions occur together. In the second cycle, the VCO transition will lead by 3/8 of a VCO cycle. In the third cycle, it will lead by 6/8 of a VCO cycle. In the fourth cycle, the VCO will again advance by 3/8 of a cycle, for a total of 9/8 of a VCO cycle. But the synthesizer keeps track of the slipping phase and detects when the accumulated phase reaches, or exceeds, a full VCO cycle. When this occurs, the VCO divide number is increased by one to swallow the extra VCO cycle. At the end of the fourth cycle, after delaying the VCO divider transition by a full cycle, the VCO divider transition will lead the reference divider transition by 9/8, - 1, or 1/8 of a VCO cycle. This slipping and correcting of the phase will continue until the end of the eighth VCO divider cycle, when the VCO divider and reference divider transitions are once again coincident. With a fractional part of three, the VCO divider will have divided by the larger divide number three times to swallow a total of three VCO cycles. This will result in a VCO frequency that has increased by 3/8 of the comparison frequency. Similarly, fractional parts of five and seven will advance the width of the error pulse even more quickly, include five and seven divisions by the larger divide number, and again return to a zero error at the conclusion of the eighth cycle of the VCO divider, while swallowing five or seven VCO cycles during the fractional cycle. In each of these cases, the fractional cycle will consist of eight comparison cycles, so a spurious signal at 1/8 of the comparison frequency will be developed.

An article by Johnathan Stillwell[2], who developed the first N synthesizers at Phillips, states that higher fractional parts raise the fractional spurious frequency, bringing it closer to the comparison frequency reference spur. This is not true. In fact, a fractional part of 1/8 will create spurs that are identical to those created using a fractional part of 7/8. The phases of the spurs will be different, however.

With a fractional part of 7/8, the waveform in Figure 2 will appear to reverse, or run backward, with the first pulse 7/8 VCO cycle wide, and the last 1/8 VCO cycle wide. With fractional parts of 3/8 or 5/8, all the pulses in

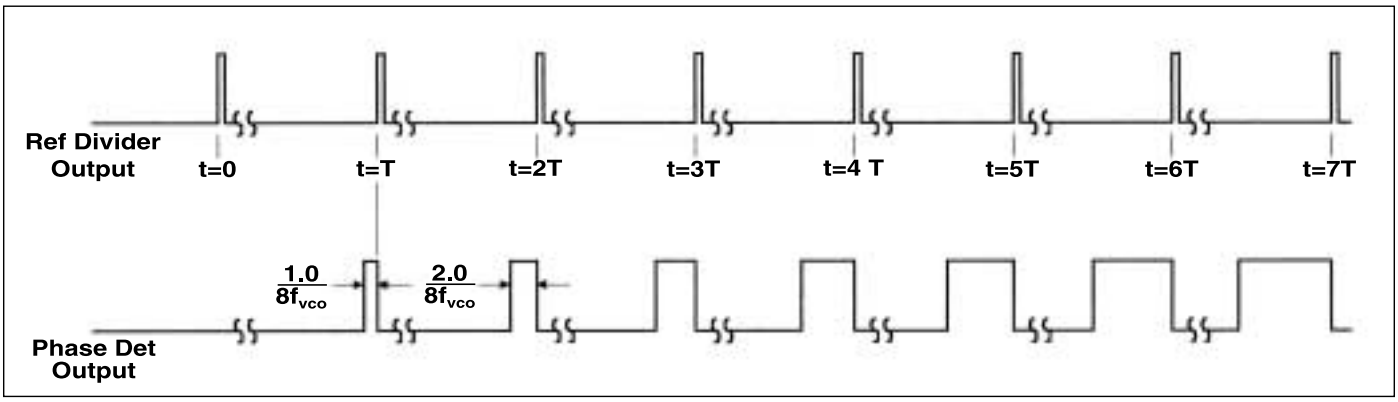

Figure 2. An example of the growth of the phase detector pulse by 1/8 of the VCO cycle as the the reference period counts up.

Figure 2 will again be present, but their order will be shuffled.

Even numbered fractional parts (2, 4, and 6) modify the cycles slightly. A fractional part of 2 or 6 will result in a zero-width error pulse after four cycles. The fractional cycle will therefore consist of just four comparison cycles, and the lowest frequency spurious signal will occur at 1/4 of the comparison frequency.

If the fractional part is set to 4, the cycle repeats after only two cycles of the VCO divider, just as in the case of the fractional modulus of 2 described above. The fractional cycle will consist of only two comparison cycles, and the lowest frequency spurious signal will be at 1/2 of the comparison frequency.

If the fractional part is set to 0, the fractional cycle and the comparison cycle are the same. In this case, the synthesizer operates just as though it were not in a fractional N mode.

If a fractional modulus of 5 is chosen, the fractional cycle will consist of five comparison cycles, producing spurs at 1/5 of the comparison frequency. This will be true for any fractional part except 0, where the synthesizer will again operate as though it were not in the fractional N mode.

The loop

Figure 3 illustrates a PSpice representation of a simple passive loop. This loop will be used to model the 500 MHz synthesizer. The circuit models frequency and phase errors are equivalent voltages. It will be necessary to examine the transfer function for the phase-locked loop as the ratio of the output frequency error (at the output of the VCO divider) to an induced frequency error at the input to the synthesizer phase detector. The input source, V_1, represents the induced frequency error. The output voltage, developed across R_4, will represent the frequency error at the output of the VCO divider.

Following the input source is an integrator with a gain of 1.0/s, which transforms the frequency error into a phase error. The charge pump, G_1, has been set to a level of 0.5 mA. This sets the phase detector gain at 0.5 mA per cycle because it will produce a DC level of 0.5 mA for a phase error of a full cycle. The loop filter, consisting of R_1, C_1, and C_2, has been designed to obtain a loop bandwidth of 10 kHz. Because the filter's transfer function is defined by the output voltage (the VCO tuning voltage) resulting from the charge pump current, the transfer function is just the impedance of the three-element network. The 500 MHz VCO is represented using a gain block with a gain of 50 MHz/V (the assumed sensitivity of the VCO). The voltage at the output of the VCO gain block represents the VCO frequency error. The ratio of the VCO output voltage to the input source (V_1) represents the transfer function for an error in the VCO output frequency caused by an error in frequency at the input of the phase detector. The VCO divider is represented by a gain block, with a gain equal to the reciprocal of the divide ratio. With a VCO frequency of 500 MHz, a step size of 50 kHz, and a fractional modulus of 8, the divide ratio is:

$$N = \frac{500\ \text{MHz}}{8 \bullet 50\ \text{kHz}} = 1250 \tag{1}$$

In this case, the divider ratio of 1,250 results in a gain of $800 \bullet 10^{-6}$. The gain is specified as negative to provide the negative feedback required if a closed loop simulation were desired.

Resistors R_2, R_3, and R_4 are present to allow PSpice to correctly perform the initial bias point calculation, or to provide a DC path to ground, as required by the simulator. It is not intended to develop the loop equations for a phase-locked loop here, or discuss the different type loops. But for completeness, a brief discussion of how the component values in Figure 3 were chosen will be provided.

As mentioned earlier, the loop shown in Figure 3 was designed to provide a loop bandwidth of 10 kHz. The loop filter was chosen to provide an open-loop gain slope of 20 dB per decade over the range of 1 kHz to 100 kHz. This 20 dB-per-decade gain slope occurs where the loop filter response (impedance) is dominated by the resistor, R_1. Below 1 kHz, the loop filter impedance is dominated by C_1, providing a 40 dB-per-decade slope. Above 100 kHz, the loop filter impedance is dominated by C_2, again providing a 40 dB-per-decade slope.

To calculate a value for R_1, the open-loop gain will be calculated at 10 kHz, and set R_1 to obtain a gain of 0 dB. The open-loop gain is given by:

$$A_v = \frac{1}{s}\left(\frac{K_\phi R_1 K_{VCO}}{N}\right) = 15.9^{-6} \bullet \frac{(0.5\ \text{ma})(R_1)(50\text{MHz/Volt})}{1250} = 0.318 \bullet 10^{-3} \bullet R_1 \tag{2}$$

where $K\phi$ is the phase detector constant (0.5 mA/Hz), and K_{VCO} is the VCO tuning sensitivity (50 MHz/V).

By setting the loop gain, A_V to 1.0, a value for R_1 can be calculated of 3.142 kΩ. The closest standard value of 3.3 kΩ has been chosen for this loop

As defined above, C_1 will just start to dominate the loop gain at a frequency of 1 kHz. This is accomplished by setting C_1 to obtain a reactance of 3.3 kΩ at 1 kHz. This requires a capacitance of

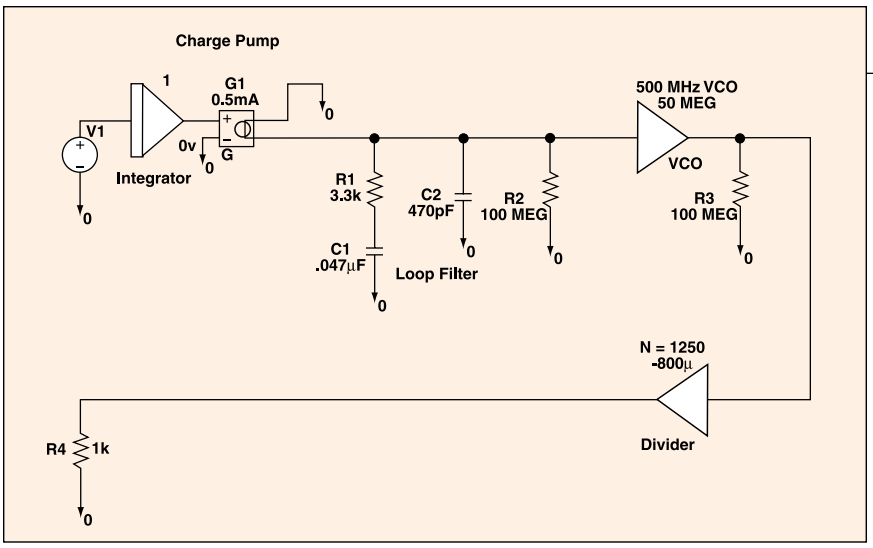

Figure 3. A simple passive loop for the model device.

0.048 µF. Again, the next closest standard value, 0.047 µF was chosen for the loop. Similarly, the value of *C2* is 470 pF, to obtain a reactance of 3.3 kΩ at a frequency of 100 kHz.

The open-loop gain for the loop in Figure 4 was calculated using PSpice. The resulting loop gain is displayed as a function of frequency in Figure 4, which shows that the desired response has been obtained.

Reference sideband calculation

Reference sidebands, whether fractional or integer, are caused by slight perturbations in the VCO tuning voltage, originating with the current pulses delivered by the synthesizer charge pump. Theoretically, the integer spurs could be eliminated by eliminating any current leakage path in the tuning line, thereby reducing the phase detector pulse width to zero. In practice, it doesn't work.

If leakage were the only cause of spurs at the comparison frequency, the spurious sideband level could be calculated by first determining the amplitude of the corresponding spurious frequency component on the tuning line because of the finite pulse width. The resultant frequency modulation of the VCO could then be calculated, followed by the sideband levels, usually by using the low-modulation index approximation. However, the synthesizer can also create reference sidebands from imbalance between the charge sink and charge source, or even from a leakage path from the reference or VCO divider outputs. These contributions make a calculation of the reference sideband level more difficult. Ignoring these additional sources, calculate the maximum allowable phase detector pulse width given a specified maximum allowable spurious level at the comparison frequency.

Frequency	Impedance
50 kHz	2.94 kΩ
100 kHz	2.53 kΩ
150 kHz	1.85 kΩ
200 kHz	1.50 kΩ
250 kHz	1.25 kΩ
300 kHz	1.06 kΩ
350 kHz	0.92 kΩ
400 kHz	0.81 kΩ

Table 1. Loop filter impedance at fractional spur frequencies.

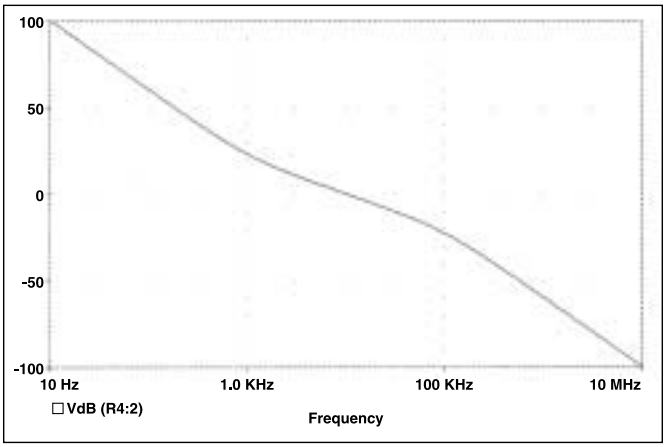

Figure 4. Open-loop gain response for the synthesizer loop.

The amplitude of a discrete frequency component on the tuning line voltage is determined by calculating the Fourier series for the pulse waveform delivered by the charge pump, and multiplying by the impedance of the loop filter. Once the Fourier coefficients are obtained, which are equivalent to the peak voltage deviation caused by each harmonic, each Fourier coefficient can be multiplied by the VCO gain constant to obtain the peak deviation of the resultant frequency modulation of the VCO, due to that harmonic. The resulting sideband level can then be calculated using the standard low-modulation index approximation:

$$spur = -20\log\left(\frac{\Delta f}{2 \bullet f_m}\right)\text{dBc} \quad (3)$$

where: Δf is the peak deviation, and f_m is the frequency of the spurious component (f_m would be equal to the comparison frequency for the fundamental reference frequency spur).

For example, to obtain a maximum sideband level of –60 dB at the 400 kHz reference frequency of the loop, Equation 2 can be rearranged to calculate the maximum allowable peak deviation (Δf) as 800 Hz. Working backward, using the VCO gain constant of 50 MHz/V, a maximum tuning line voltage amplitude of 16 µV (800 Hz/50 MHz/V) peak for the 400 kHz component on the tuning voltage waveform can be calculated.

The magnitude of the impedance of the loop filter at the first eight fractional spur frequencies is displayed in Table 1. The impedance is equal to 0.81 kΩ. Using this impedance we can calculate the maximum allowable level of the 400 kHz component of the charge pump current waveform. A maximum amplitude for the 400 kHz component of the charge pump pulse of 19.8 nA (16 µV/0.81 k) is obtained.

For a narrow current pulse, the Fourier coefficient of the k^{th} harmonic is given by:

$$I_k = 2 \bullet I_p \bullet \frac{\tau}{T} \quad (4)$$

where I_p is the peak current (equal to the charge pump current), τ is the pulse width, and T is the waveform period.

In this loop, the charge pump current is 0.5 mA, and the period (T) of the

400 kHz reference frequency is 2.5 ms. By rearranging the equation, a maximum pulse width (τ) of 0.05 ns can be calculated to obtain the desired reference level of –60 dBc above. It is not likely that an error pulse this narrow will be achieved using the actual chips. Therefore, additional filtering of the reference spurs would probably be required.

Fractional spurs

The level of the fractional spurious sidebands can be calculated in a similar manner. But, in the case of the fractional spurs, one must keep in mind that the charge pump waveform is more complicated. As illustrated earlier, because the divider ratio is being altered between two adjacent integers to accomplish the fractional division, the charge pump waveform may take as many as eight reference divider cycles before it repeats. And the charge pump waveform may contain as many as seven pulses that contribute to the Fourier series.

As a second consideration, the frequency of the spurious component is closer to the loop bandwidth, and will be subjected to less attenuation from the loop filter. With a fractional modulus of 8, and a reference frequency of 400 kHz, the fractional spur frequency can be as low as 50 kHz, while the loop bandwidth is 10 kHz (see Figure 4, which shows the phase detector output waveform expected for a fractional modulus of 8, and a fractional part of 1/8[3]).

As discussed earlier, the phase detector outputs an error pulse whose width grows with each cycle, until returning to a zero width at the end of the eight comparison cycles that form the fractional cycle. According to the device's design, the phase detector pulse always ends at the leading edge of the reference divider output pulse. The VCO is running fast, by 1/8 of a VCO cycle for each comparison cycle. So the leading edge of the phase detector output, which is triggered by the VCO divider, advances by 1/8 of a VCO cycle for every cycle of the reference divider.

Obtaining the Fourier series for this waveform is a relatively simple exercise. By using superposition, the Fourier series can be added together for the eight separate pulses, and one of the eight pulses is of zero width. The harmonic components for the Fourier series equivalent of a train of current pulses is given by:

$$I_k = 2 \bullet I_{peak} \bullet \frac{\tau}{T} \frac{\sin\left(\frac{k\pi\tau}{T}\right)}{\frac{k\pi\tau}{T}} e^{-j\phi} \quad (5)$$

where: I_k is the complex coefficient of the k^{th} harmonic, τ is the pulse width, T is the period for the fractional cycle, I_p is the charge pump current, and ϕ is the angle by which the center of the pulse is offset from t = 0.

The width (τ) of each of the eight pulses can be calculated from the expression:

$$\tau_k = \frac{k}{8(f_{VCO})} \quad (6)$$

where k takes on values from 0 to 7, and f_{VCO} is the frequency of the VCO.

The offset angle is produced by two mechanisms. First, each pulse is delayed by 45 degrees (or $\pi/4$) of the fractional cycle from the previous pulse. In addition, as can be seen in Figure 4, each pulse is advanced by an amount equal to half of its width. Thus the angle in Equation 5 can be calculated from the expression:

$$\phi_k = \frac{\pi}{4}k - \pi\frac{\tau_k}{T} \quad (7)$$

If the Fourier series for these eight pulse trains are calculated and summed (actually seven because pulse 0 is a zero width pulse), the harmonic content of the phase-detector charge pump current for the fractional spurs will be obtained. This has been done for the previous example, using Mathcad. The results are:

$$I_k = 0.1308\ \mu A\ @112.6° \quad (8)$$

for the first eight harmonics of the fractional spurious frequency. The first eight harmonic coefficients are essentially identical because of the narrow pulses. If the pulse widths are calculated using Equation 6, it is clear that they vary from 0.25 ns to 1.75 ns, with T equal to 20 µs. These are indeed narrow pulses.

Now the level of the fractional spurs can be calculated from the magnitude of the current calculated above. The calculation will be carried out for only the first fractional spur, at 50 kHz. Because of the loop filter, this will be the worst-case spurious signal.

Looking back at Table 1, the magnitude of the impedance of the loop filter at 50 kHz is 2.94 kΩ. Multiplying by the current magnitude, the level of the 50 kHz component on the tuning line can be calculated as:

$$V_{50kH} = (0.1308\ \mu A)(2.94\ k\Omega) = 0.38\ mV \quad (8)$$

The peak deviation can be calculated by multiplying by the VCO sensitivity:

$$\Delta f = (0.38\ mV)(50\ MHz/V) = 19\ kHz \quad (9)$$

Using the low-modulation index approximation for calculating the sideband level:

$$spur = -20\log\left(\frac{\Delta f}{2f_m}\right) dBc =$$
$$-20\log\left(\frac{19\ kHz}{2(50\ kHz)}\right) = -14.4\ dBc \quad (11)$$

However, in calculating the results, it is noted that the level is a bit high.

Fractional N compensation

It has been illustrated that the phase-detector charge pump will put out a series of pulses of increasing width during the fractional N cycle. An important point to note is that the output pulses from the charge pump during fractional N operation do not tune the VCO. They take no part in keeping the loop in lock. They are just a result of the constantly slipping phase. An additional circuit designed into the device provides an additional charge pump, independent of the VCO, intended to extract an equal amount of charge from the loop filter, at the same time as the phase-detector charge pump dumps it into the loop filter. This extra charge pump serves two important functions. First, it prevents the tuning line from slewing up to the positive rail by removing the charge that is pumped into the loop filter by the phase detector charge pump. Second, if the charge is extracted at the same time as the phase detector charge pump adds charge, the effects will cancel, and the fractional spurs will disappear.

Figure 5 illustrates the waveform

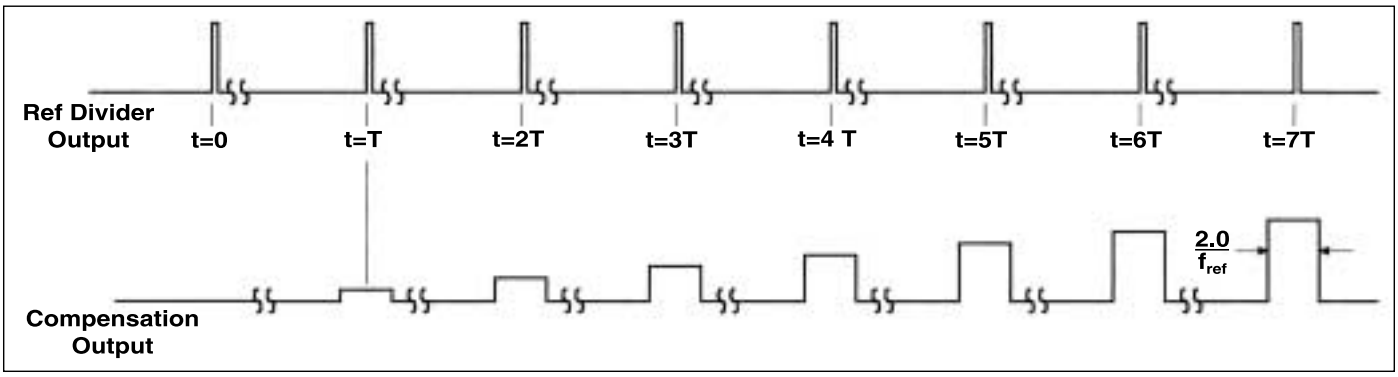

Figure 5. Example of the variation of compensation pulse to compensate for the varying error pulse widths.

produced by the compensation charge pump[4]. For each error pulse from the phase detector charge pump, there is a compensating pulse from the fractional compensation charge pump. Just as the phase detector pulses grow linearly in width from pulse to pulse, the compensation pulses similarly grow linearly in amplitude. So long as the area of each compensation pulse is equal to the area of the corresponding error pulse, the net charge transferred into the loop filter will be zero. Unlike the phase detector pulses, the seven compensation pulses are all the same width, each being twice the period of the external reference clock. For a 10 MHz clock, the pulse width is 200 ns, much wider than the corresponding phase detector pulse widths. It can also be seen in Figure 5 that each of the pulses in the compensation waveform is at an exact multiple of 45 degrees, whereas the phase detector pulses were offset slightly from 45 degrees, by half the pulse width. It can be seen that the pulse trains do not match, so some fractional spurs should be expected to remain. Knowing that the error pulse amplitudes are 0.5 mA, and calculating the first phase detector pulse width from Figure 4 as 0.25 ns, its area can be calculated:

$$Q_{err} = (0.5\text{ mA})(0.25\text{ ns}) = 0.125\text{ pC} \quad (12)$$

And, knowing that the width of the compensation pulses is 200 ns, the required amplitude for the first compensation pulse can be calculated:

$$i_{comp} = \frac{0.125\text{pC}}{200\text{ ns}} = 0.625\text{ uA} \quad (13)$$

Because increases in the charge are transferred linearly from pulse to pulse, by setting the first pulses to cancel, each of the seven pulses in the waveform have been forced to cancel. To determine the first eight Fourier series components for the compensating pulse train, it is necessary to repeat the earlier analysis. Using Mathcad, the amplitude of the first (50 kHz) component is:

$$I_1 = 0.1306\text{ uA at an angle of 112.5 degrees} \quad (14)$$

providing a substantial cancellation of the 50 kHz component of the phase detector output. The eighth harmonic (which is equal to the normal reference spur frequency) is:

$$I_8 = 0.1293\text{ uA at an angle of 112.5 degrees} \quad (15)$$

The amplitude of the eighth harmonic component has dropped slightly. This is a result of the relatively wide compensation pulse. Even with a 200 ns pulse width, the drop is slight, and the cancellation of the phase detector pulse will be good.

In practice, the phase detector pulses will not be cancelled perfectly. The major problem will be the inability to accurately set the current of the compensation pulses. If trimmer adjustment is not desirable, the two current pumps will have to be set using 5%, or perhaps 1% resistors[5]. Assuming 5% resistors are used, the best cancellation guaranteed (with one charge pump 5% high, and the other 5% low) is about 90%, or 20 dB. With an additional 20 dB of spurious rejection due to the cancellation, the spurious level in the example (referring back to equation 11) would be expected to become –34.4 dBc. Using 1% resistors, the best to hope for is a reduction of about 34 dB, assuming there are no other sources of error.

Raising the frequency of the reference source will create a narrower cancellation pulse. This will allow a better cancellation of the higher order harmonics of the fractional spur. It was illustrated in the example that the eighth harmonic, which is at the same frequency as the normal reference spur, was only slightly affected by the wide pulse width, certainly less than the degradation caused by using 5% resistors. Low reference frequencies, however, will result in degraded cancellation of the higher frequency harmonics.

Changing the VCO frequency

In the example, it was shown that the width of the error pulses start at 1/8 of a VCO cycle, growing to 7/8 of a VCO cycle before returning to a width of zero. This implies that the charge dumped into the loop filter by the phase detector charge pump, which is proportional to the pulse width, is inversely proportional to the VCO frequency.

There is no digital control over the compensation charge pump. It is set once by an external resistor. Its amplitude, width, and therefore its charge contribution, are fixed. Thus, if cancellation of the phase detector pulses is to be maintained across some VCO frequency range, it must be accomplished by adjusting the main phase detector charge pump current as the VCO frequency is changed. Because the phase detector pulse width is inversely proportional to the VCO frequency, the phase-detector charge pump current will have to be made proportional to the VCO frequency to compensate for the decreased pulse width. This is accomplished by adjusting the CN register in the synthesizer chip.

Fortunately, this is a profitable way to maintain a constant loop gain as the VCO frequency changes. As the VCO frequency increases, the loop gain drops inversely proportional to the

VCO frequency, as the VCO divider number increases.

Increasing the phase-detector charge pump current, proportional to the VCO frequency, will exactly compensate for the reduced loop gain. It is known that the energy contained in these pulse waveforms will increase as the width of the pulses increase. This implies that a higher VCO frequency will result in less energy in the fractional spurs. We should see a 6 dB decrease in the fractional spur level every time the VCO frequency doubles.

Calculating the divider numbers

For the devices used as models in this article, the manufacturer has proposed a rather complicated algorithm for calculating possible divider combinations for achieving a desired main divide ratio, using up to a three- and four-modulus divider. In this section, a simple algorithm for selecting a divide number will be examined.

By including a three- or four-modulus prescaler, it is possible to obtain a wider range of divide ratios than are possible with a dual-modulus prescaler. For example, a 64/65 dual-modulus prescaler can obtain any divide number greater than 4,031. Depending upon the device, using a three-modulus divider, one can obtain any divide ratio above 1,023. Using a four-modulus divider, one can obtain any divide ratio above 765.

The standard 64/65 dual-modulus prescaler will divide an input signal by a fixed number of 64s, plus a fixed number of 65s. It will be assumed that any number of divisions by 64 and/or 65 is legal, including zero, so long as at least one of the two divide numbers is non-zero.

The limitations on the divide number for this prescaler should be examined. The first divide ratio that can be obtained is just 64, by dividing with the 64 divider once. By switching to the other divider, 65 can also be divided by. No other numbers can be divided by until we reach the number 128, which can be obtained by dividing by 64 twice. Thus, a "hole" of 62 numbers is located between 65 and 128. 129 can also be divided by, by dividing once by 64, and once by 65. And 130 can be divided by dividing by 65 twice. A pattern is developing. There is a range of divide ratios that can be obtained, with each range starting at an integer multiple of 64. For the n^{th} multiple of 64, it is possible to obtain that divide ratio, and the next n divide ratios.

For example, at the 3^{rd} multiple of 64, or 192, it will be possible to hit four divide numbers–192 to 195. By the time the 63rd multiple of 64 (4,032) has been reached, it will be possible to obtain the full range of 64 divide ratios that separate this from the next multiple of 64, leaving no hole in between these two multiples of 64. Thus the last unobtainable divide ratio with a 64/65 divider is 4,031. Many divide ratios below 4,031 can be hit, but there will be holes. These holes will grow longer (contain more unobtainable, sequential divide numbers) as the divide ratios drop.

Look again at obtaining successive divide ratios. Start with some multiple of 64. This divide number is reached by using only the 64 divider. Then move up to the next divide number by exchanging a 64 for a 65. This works so

long as the remaining number of divisions by 64 is non-zero. Once zero divisions by 64 have been reached, they can't be exchanged for 65s any more. At this point, it is necessary to jump up to the next multiple of 64. How does this affect the obtainable divide ratios?

Let's consider a triple modulus prescaler where division by 72 is also possible. The device programs the 64 modulus divider by loading the total number of divisions, minus two, into the control register. Thus, the smallest divide ratio obtainable corresponds to programming a zero in this register, obtaining a minimum divide ratio of 128. To obtain a divide ratio of 129, it would be necessary to divide once by 64, and once by 65. But because the minimum number of divisions by 64 has already been found, this is not possible. The next achievable divide number is three times 64, or 192. But this time the number of divisions by the 64 divider can be dropped by one. Therefore it is possible to achieve a divide ratio of 193 by setting the 64 divider to two, and the 65 divider to one. But it is also possible to trade off the extra divide-by-64 for a divide-by-72, achieving a divide ratio of 200. There aren't any more extra divide-by-64s, so divide ratios are out until the next multiple of 64.

Similarly, at the fourth multiple of 64, it will be possible to hit divide ratios of 256, 257, and 258. By trading one of the divide-by-64s for a divide-by-72, divide ratios of 264 and 265 can be achieved. With both of the extra 64s traded for divide-by-72s, it will be possible to achieve a divide ratio of 272. Again, there is nowhere to go until the next multiple of 64, because there are no more extra 64s to exchange for a divide by 65 or 72.

It can be seen from this pattern that the divide numbers are filling in more rapidly with the triple modulus divider. But there is something else that should be noticed in this pattern. Not just one region between multiples of 64 is being filled, but multiple regions, with each region starting eight positions above the previous region. The reason for the existence of these multiple regions is that a single divide-by-64 is being traded for a divide-by-72, giving us an offset of eight, at a cost of a single divide-by -64. With the dual modulus prescaler, this would cost eight divide-by-64s. The triple modulus prescaler then, due to the divide-by-72, can be more thrifty in the consumption of divisions by 64. Look at an algorithm for calculating the divide numbers for the triple modulus divider.

Given a desired divide ratio (N), start by finding the largest multiple of 64 (M_{64}) that is less than or equal to N. The multiplier will be called N_{Base}.

$$M_{64} = N_{Base}(64) \tag{16}$$

Then take the difference (D):

$$D = N - M_{64} \tag{17}$$

As seen above, this difference between the desired divide number, and the next lowest multiple of 64, must be made up by using the other divider ratios.

Next, take the largest multiple of eight (M_8) that is less than or equal to D. Call the multiplier N_{72}.

$$M_8 = N_{72}(8) \tag{18}$$

As you might guess, this multiplier is the number to program into the 72 modulus divider. If M_8 is equal to zero, the dual modulus mode will be selected so as not to divide by 72.

Again, take the difference (N_{65}).

$$N_{65} = D - M_8 \tag{19}$$

N_6 represents what is left after the divisions by 64s and 72s, and is the number to program into the 65 modulus divider.

Finally, the required number of divisions by 64 can be calculated:

$$N_{64} = N_{Base} - N_{72} - N_{65} \tag{20}$$

The value of N_{64} will now tell if N is a valid divide number. If N_{64} is less than the minimum number of divisions that can be programmed into the device (2), then this divide number is not obtainable with the triple modulus prescaler. If N_{64} is larger than the maximum number of divisions by 64 that can be programmed into the device (4,097), it may still be possible to obtain the divide number by using additional divisions by 72 and 65, but the algorithm will fail. However, this is already an enormous divide number. When dealing with a device that uses a four-modulus prescaler, with divide ratios of 64, 65, 68 and 73, the technique is similar, but offers a bit more flexibility.

With a four-modulus prescaler, more divide numbers can be obtained, and the maximum divide number that cannot be obtained drops. With a three-modulus prescaler, start by finding the largest multiple of 64 that is less than or equal to the required divide number. Then pull out as many divide-by-72s as possible, making up the difference with divide-by-65s, reducing the number of divisions by 64 by an amount equal to the total number of divisions by 65 and 72. With the four-modulus prescaler, the algorithm is similar. In this case, calculate again the maximum multiple of 64 less than or equal to the required divide number. Then pull out as many divide-by-73s as possible. In this case, pull factors of nine out of the difference, rather than factors of eight. But with the four-modulus prescaler, one can also trade off a division-by-64 for a division-by-68, consuming an additional difference of four with the loss of a single division-by-64. Thus the four-modulus prescaler is even more thrifty in its consumption of divide-by-64s. Divide-by-65s are still required to make up any difference less than four.

Conclusions

Two main topics have been discussed. First, the generation of fractional spurs in fractional N synthesizer chips. While the normal reference spurs can be reduced by reducing tuning line leakage, and by proper design of the charge pump, the presence of fractional spurs is guaranteed by the finite and well-defined pulse widths of the phase detector charge pump during fractional N operation. It has been illustrated how to calculate the fractional spur levels resulting from these finite width pulses. It has also been shown that the model chips provide a mechanism for cancelling the fractional spurs. It was determined that a reduction of about 20 dB should be expected using 5% components to set the charge pump currents. It was illustrated that a fractional modulus of 8 will generate fractional spurs at multiples of 1/8 of the comparison frequency for any fractional part that is odd. A fractional part of two or six will generate fractional spurs that are multiples of 1/4 of the comparison frequency. A fractional part of four will generate fractional spurs that are multiples of 1/2 of the comparison frequency. A fractional part of zero will not generate any fractional spurs. If a fractional modulus of 5 is used, all fractional parts will

generate fractional spurs at multiples of 1/5 of the comparison frequency, except a fractional part of zero, which will not generate any fractional spurs. A calculation of the expected fractional spur level for a 500 MHz synthesizer was gone through using a fractional modulus of 8. Also seen was how to calculate the current required in the compensation charge pump to cancel the effect of the phase detector error pulses. The article examined how to compensate as the VCO frequency changes and showed that fractional spur levels are not insignificant. If the fractional spurs lie too close to the loop bandwidth, they will not be filtered adequately.

Also discussed was the selection of the divide numbers for the three- and four-modulus prescalers. An algorithm was detailed for the triple modulus prescaler, and the algorithm for the four-modulus prescaler was sketched out. With these algorithms, it can be determined if a desired divide ratio is possible. If so, an optimum set of divide numbers can be calculated that maximizes the number of divisions by 64.

With the information presented in this article, the engineer will be better prepared to determine the tradeoffs involved in choosing the fractional N loop.

RF

Notes

1. It would actually look more like the tuning voltage steps up and down, since the entire ramp would occur during the short duration of the error pulse.

2. "A Fractional N Frequency Synthesizer for Digital RF Communications". *RF Design*, February 1993.

3. Again assuming that the cancellation techniques used by the Phillips models for this article synthesizers will result in a net charge into the loop filter of zero with this waveform.

4. Again illustrated the waveform for a fractional part of one, where the pulse widths increase linearly from each to the next.

5. The necessity of adjusting the level of the phase detector charge pump as the frequency changes will soon be shown, if it is desirable to maintain the best performance across a range of VCO frequencies. The smallest step size for this adjustment is about 0.5%

About the author

Gary Appel has been working with RF circuit and system design since 1976. He has been an independent consultant since 1993. He is presently working with Robidart & Associates, a consulting company specializing in the design of RF products. Appel also enjoys writing software for modeling RF systems and circuits. He may be contacted at Robidart & Associates, 18430 Technology Dr Suite A, Morgan Hill, CA 95037. His e-mail address is: *gappel@truevine.net*.